

Perception and Memory

Beyond Representationalism and Relationalism

André Sant'Anna

Perception and Memory
Beyond Representationalism and Relationalism

André Sant'Anna

A thesis submitted for the degree of

Doctor of Philosophy at

The University of Otago

June 2018

To my parents, Sidnei and Teresa.

Acknowledgments

First and foremost, I am extremely grateful to my supervisors Kirk Michaelian, Andrew Moore, and Cathy Legg. Their guidance during these three years have helped me enormously. Thanks to Andrew for taking interest in my project when I first applied to Otago and for providing insightful feedback to my thesis all the way through. Thanks also for always reminding me of the importance of thinking about the ‘big picture’ of my project. Thanks to Cathy for accepting to join my supervisory panel halfway through and for making yourself available to discuss my thesis at any time. Our conversations about pragmatism, Peirce, and perception were fundamental in shaping my way of thinking about these topics.

I’m especially grateful to my primary supervisor, Kirk Michaelian, who has not only been an exemplary supervisor, but also a mentor and a friend. Thanks for being patient and for always being available to read many of my drafts, even the rough ones. Your guidance and willingness to always help me to improve were essential to my development as a philosopher.

I’d also like to thank Ruth Millikan, who was extremely kind to help me, an unknown undergraduate student from Brazil, during the process of applying to graduate school, and particularly for telling me about Otago and for first putting me in touch with the department. Thanks also to Colin Cheyne for taking interest in my project and for handling my application at Otago.

I’m grateful to the Philosophy Department at the University of Otago for all their support during these years. The friendly and extremely professional atmosphere provided by the department were essential to the development of my thesis. Thanks also to Chloe Wall, Ali Hossein Khani, Manuel Lechthaler, Ünsal Çimen, and Guillermo Badia for their friend-

ship and stimulating conversations.

I would like to thank my parents, Sidnei and Teresa, for their ongoing support from the time I first decided to study philosophy to my decision to move to New Zealand to do my PhD. None of this would have been possible if it were not for their support.

I'm also immensely grateful to my wife, Ligia, for all her support and partnership during these three years of hard work. Besides providing valuable feedback to the thesis itself, her love and understanding were essential to overcome the inevitable hardships of writing a PhD thesis.

I'd also like to express my gratitude to the University of Otago for granting me a Doctoral Scholarship to write my thesis, to the German Academic Exchange Service (DAAD) for funding my visit, from October to December 2017, to the Ruhr-University Bochum to work on my thesis, and to Markus Werning for being a wonderful host.

Parts of this thesis have been previously published in academic journals. Chapter 2, "Perception pragmatized", was published in *Philosophia* 46(2): 411–432, 2018. Chapter 3, "Thinking about events", is forthcoming in *Review of Philosophy and Psychology* and Chapter 5, "The hybrid contents of memory" is forthcoming in *Synthese*. I'm grateful to Springer Nature for granting me permission to reproduce them here. Finally, the Appendix Chapter, "Mental time travel and the philosophy of memory", is forthcoming in *Unisinos Journal of Philosophy*.

Abstract

This thesis is a collection of five self-standing articles dealing with different issues relating to representationalism and relationalism in contemporary philosophy of perception and contemporary philosophy of memory. The main goal is to motivate a hybrid approach, where insights from representationalism and relationalism are reconciled, to current debates in both domains. The thesis is divided in two parts. Part I, which deals with perception, starts by seeking alternative relational views of perception by relying on ideas from classical pragmatism. These attempts further result in the development of a hybrid theory of perception, grounded on the pragmatist theory of perception offered by Charles Peirce, that is more sympathetic to “austere” versions of relationalism. Part II, which deals with memory, starts by exploring the prospects of representationalism and relationalism to account for the objects of memory and the objects of episodic hypothetical thought. Relational accounts are further explored in this context for, despite their historical importance, they have been unpopular recently due to the difficulty they face in dealing with memory errors. Finally, as a positive proposal, I offer a hybrid theory of memory that is influenced by hybrid theories of perception. Besides calling the attention of philosophers of memory to the importance of taking relationalism seriously, this hybrid theory provides a novel theoretical framework for contemporary philosophers of memory.

Contents

Introduction	1
1. Indirect and direct realism	3
2. Perception	7
2.1 Representationalism	7
2.2 Relationalism	10
3. Memory	12
3.1 Representationalism	12
3.2 Relationalism	14
4. Towards a hybrid approach	16
5. The thesis	19
I Perception	22
1 Contextualizing perception	23
Introduction	23
Outline	24
1.1 Perceptual error and the nature of perceptual experiences	25
1.1.1 Perceptual error	25
1.1.2 Relationalism and naïve realism	28
1.2 Perceptual contextualism	32
1.2.1 The resistance view of experiences	32
1.2.2 Pragmatic contexts	36

1.2.3	Perceptual contextualism and disjunctivism	45
1.3	Objections	48
1.3.1	Objection 1: Veridical hallucinations	48
1.3.2	Objection 2: Illusions	51
1.3.3	Objection 3: Representationalism by the back door?	52
2	Perception pragmatized	55
	Introduction	55
	Outline	57
2.1	The nature of perception	58
2.2	Peirce's theory of perception	63
2.3	The pragmatic view	69
2.3.1	The representational aspect of perception	69
2.3.2	The relational aspect of perception	72
2.3.3	The veridicality of perception	78
2.4	Conclusion	84
II	Memory	85
3	Thinking about events	86
	Introduction	86
	Outline	88
3.1	From the objects of memory to the objects of episodic hypothetical thought	89
3.2	Direct realism and extreme disjunctivism	92
3.2.1	Motivating direct realism	92
3.2.2	Extreme disjunctivism	93
3.3	Representationalism and the silence of episodic hypothetical thought	102
3.3.1	Motivating representationalism	102
3.3.2	The silence of episodic hypothetical thought	103
3.4	The pragmatist alternative	106

3.4.1	Motivating pragmatism	107
3.4.2	Applying pragmatism	111
3.4.3	Avoiding disjunctivism	116
3.5	Pragmatism and auto-noesis	125
4	Failing to remember	129
	Introduction	129
	Outline	131
4.1	Introduction	131
4.2	The relational analysis of remembering	135
4.3	Unsuccessful remembering	140
4.3.1	The problem of confabulation	142
4.3.2	The problem of misremembering	149
4.4	Conclusion	154
5	The hybrid contents of memory	156
	Introduction	156
	Outline	159
5.1	The objects of episodic memory	159
5.2	Towards a hybrid view of memory	164
5.2.1	Hybrid contents	169
5.2.2	Error revisited	172
5.2.3	Indistinguishability revisited	177
5.2.4	Epistemic particularity revisited	177
5.3	Hybridism and the philosophy of memory	178
5.3.1	The co-temporality problem	179
5.3.2	Continuism vs. discontinuism	182
5.4	Objections	189
5.4.1	The character of memory vs. the objects of memory	189
5.4.2	The temporality of the objects of perception and memory	190

5.4.3	Autonoetic consciousness	191
5.4.4	Disjunctivism again?	192
Conclusion		194
Appendix		198
A	Mental time travel and the philosophy of memory	198
	Introduction	198
	Outline	199
A.1	Introduction (paper)	199
A.2	Episodic memory and mental time travel	200
A.3	Mental time travel and the philosophy of memory	204
	A.3.1 The causal theory of memory	205
	A.3.2 Mental time travel and our knowledge of the past	209
	A.3.3 The objects of mental time travel	212
	A.3.4 The metaphysics of mental time travel	214
A.4	Conclusion	217
Bibliography		232

Introduction

This thesis explores the disputes between representationalist and relationalist views of perception and memory in contemporary philosophy of perception and philosophy of memory. The main goal is to motivate a hybrid approach, where insights from representationalism and relationalism are reconciled, to current debates in the area.

As a first introductory step to this project, I should say a few words about the nature of the chapters and the structure of the thesis. In terms of structure, the thesis is organized as a collection of publishable papers. It is composed of five self-standing articles and an appendix corresponding to a short discussion article. The articles were produced with the purpose of being submitted to specialist journals. Three of them have been accepted for publication (Chapters 2, 3 and 5, plus the discussion article in the Appendix A), one of them is under review (Chapter 4), and one of them is currently in the process of being revised for resubmission (Chapter 1). Since they are self-standing pieces of work, the issue of how they fit together to compose a PhD dissertation becomes central. One natural strategy to address this would be to re-write the chapters and try to arrange the seemingly unrelated pieces in a way that would produce a convincing narrative addressing a certain problem or question. I decided *not* to pursue this strategy here.

The main reason for not pursuing it is that doing so would produce an artificial and inaccurate picture of how my thinking about the relevant issues evolved during my PhD. As will become clear later, because the chapters are self-standing pieces, the formulation of some ideas at different chapters varies, sometimes significantly, thus resulting in some theoretical conflicts. While, as a general rule, avoiding conflicts is a desirable thing, not all conflicts are harmful to a project. This is because the development of an idea in different contexts,

even when the idea itself is presented in conflicting ways, can do more to clarify its implications than an overly rigorous formulation that is context-insensitive. This capacity to apply the same or similar ideas in different contexts, and hence to enhance my understanding of them, was crucial to the completion of this project. Thus, instead of working towards a unified project with a definite purpose in mind, my efforts have been dedicated to exploring the implications of broadly related ideas to various issues in the philosophy of perception and the philosophy of memory.

A second and less central reason for not pursuing the strategy above is practical. Because more than half of the chapters have been accepted for publication, making changes to them would alter their content, which could create potential conflicts between the versions composing the thesis and the published versions. Since the published versions are more likely to have a far reach, and are thus bound to be treated as the “official” discussion of the topic by an author, I thought that having an accurate representation of them in the thesis would diminish the risks of creating conflicts or misrepresentations.

Refraining from pursuing the initial strategy does not mean, however, that no attempt has been made to provide a coherent overall picture of the thesis. Instead of rewriting the chapters, I opted for the inclusion of introductory sections at the beginning of each of them, where I discuss the relationship of the chapter in question with one or more of the other chapters. Such introductions include discussions of how common notions used in the chapters are understood and/or formulated differently, how my thinking about a certain topic developed, and the reasons for certain argumentative moves that, out of context, would seem unmotivated or odd. The introductions are followed by short outlines providing overviews of the structure of the chapters. On top of chapter introductions and outlines, I have also included this fairly long introduction to the thesis, which offers a preliminary discussion of most of the problems that will be discussed later in the chapters, along with some discussion about how I have approached those topics more recently. The thesis is finished with a Conclusion making some final remarks as to how the chapters contribute to achieving the overall goal of the thesis.

In sum, by pointing these things out, I hope to make explicit that this thesis is not meant

to be read as a book-like project, but rather as a collection of philosophical essays on topics that are broadly related. The broad topic uniting those essays is, as the title of the thesis suggests, the dispute between representationalism and relationalism in perception and in memory. In an attempt to situate most of the questions that will appear later in the chapters, the next section of this Introduction will discuss indirect and direct realism, the historical predecessors of representationalism and relationalism,¹ and some problems and advantages of each of those views. After contextualizing the dispute, I move on to consider how it unfolded in the contemporary debate about perception and about memory. I then introduce hybrid views of perception and discuss how they provide an insightful way to think about representationalism and relationalism not only in perception, but in memory too. Finally, the last part of the introduction offers a brief summary of each chapter and outlines the structure of the thesis in more detail.

1. Indirect and direct realism

Traditionally, indirect realism is the view that the objects of perception and memory are mental representations of the real world — in particular, in the case of perception, what is represented are things in the environment, and in the case of memory, what is represented are the past things or events. For traditional indirect realists, our access to the world is indirect, in the sense that it is mediated by mental representations. Among defenders of indirect realism are philosophers such as the early empiricists (e.g., Hume 2011; Locke 1975) and early twentieth century sense-data theorists (e.g., Russell 1912; Broad 1925; Ayer 1956). Traditional versions of direct realism are, in contrast, the views according to which perceiving and remembering involve being directly related to the real world — i.e., in the case of perception, being related to the objects residing in the environment, and in the case of memory, being related to the past things or events themselves. For traditional direct realists, thus, our access to the real world is unmediated. Proponents of direct realism include philosophers such as Reid (2000), Laird (1920), and John Cook Wilson (1926).

¹I want to acknowledge here that this is a controversial claim, especially for philosophers of perception. For a more detailed discussion of the motivations behind it, see the discussion about perception in this Introduction.

The traditional dispute between direct and indirect realists revolves around two related, but importantly distinct, questions about perception and memory. The first is the question of the *objects of perception*: When one perceives something, what is the object of one's perceptual experience? The second is the question of the *objects of memory*: When one remembers something, what is the object of one's memory?² Depending on how one answers these questions, different, and often incompatible, pictures of the epistemology and the metaphysics of perception and memory will emerge.

Consider, for example, the questions of how we know something about the environment based on our perceptual experiences and how we know something about the past based on our memories. Call this the *knowledge problem*. Direct realism offers a neat explanation of how this is possible. In the case of perception, it is argued that perceptual experiences can ground our knowledge of the world because they make us *directly* related to what is out there in the world. Similarly, in the case of memory, direct realists say that we know things about the past because, when we remember, we stand in a direct relationship to the past events themselves. Because the relationship is direct, the objects in the world and the past events are simply presented to us; there are no intermediaries that can deviate, or contaminate, our access to those things. Indirect realists, in contrast, face difficulties to provide simple answers to these questions. The reason is that, both in perception and in memory, our access to the world is dependent on our relation to mental representations of, respectively, the objects in the environment and the past events. Because the relationship between representations, on the one hand, and the worldly objects and the past events, on the other hand, is often contingent, there is no guarantee that one's mental representation of the environment or of the past has anything in common with how the environment actually is, or with how the past was. Therefore, the indirect realist needs to provide an additional explanation of how, in being related to mental representations, we can form knowledge about reality.

While direct realism provides a neat account of the relationship between perception,

²Because the focus of my discussion will be memory of events, or what psychologists call *episodic memory* (see Tulving 1972, 2002), whenever I use the term "memory" without any further qualification, I will be referring to episodic memory.

memory, and the knowledge formed on their basis, it faces important difficulties to explain the possibility, and indeed the actual occurrence, of perceptual and memory errors. Call this the *error problem*. Perception and memory are about the real world, but it is not always the case that what we perceive, or that what we remember, corresponds to how the world actually is or was. It is well-known, both from common sense and from empirical research, that it is not uncommon for us to perceive and remember things erroneously — for example, when a straight stick that is half-submerged in water appears bent to me, or when I remember there being a clown at my last birthday party when it is false that there was one — or even for us to perceive and remember things that are not there, or that did not occur — e.g., seeing a spider on the wall when there is no such spider, or remembering having a birthday party ten years ago where there was no such party. On the face of such occurrences, it is not clear how they can be explained in a direct realist framework. For direct realists, perception and memory consist of direct relations to the things that are perceived and remembered, but if those things do not exist, it is hard to see how we could possibly have perceptual experiences or memories of them. Indirect realism deals with this question in a simple and coherent way. It explains the occurrence of error in terms of the presence of mental representations that, despite appearing to be about the real world, fail to be satisfied by, or to be related to, things that exist. In this sense, the problematic aspect of indirect realism discussed above, namely, that representations do not necessarily correspond to how the world is, becomes a strength of the view when it comes to explaining the occurrence of error.

A third issue, which has received a lot of attention in the contemporary literature, concerns the phenomenal character of, or to how things appear to us in, perception and memory. Call this the *phenomenology problem*. When we perceive something, or when we remember an event, it seems to us that we are perceiving the thing itself, or remembering the event itself. In other words, from a first-person point of view, perception and memory seem to put us in direct contact with the things or events that are perceived and remembered. This seems to lend support to direct realist views, as for indirect realists, the objects of perception and memory are mental representations, and it certainly does not seem to us that we are

perceiving or remembering representations, but instead the things themselves. The advantage of the direct realist is only apparent, though. When we take the occurrence of errors into account, direct realism simply fails to provide a satisfactory account of the phenomenal character of those mental states. Perceptual and memory errors are not perceived or remembered as such; usually, when we perceive or remember erroneously, it still seems to us that we are perceiving and remembering real things. Thus, finding out that we are not perceiving or remembering correctly involves an additional step, which often has to do with the consideration of those mental states in the context of some background knowledge of the world, or even with interacting with the relevant people who are or were also related to the relevant objects and events. The problem that arises for the direct realist is, therefore, that one can have the relevant experiences of objects and events even when there are no such objects and events. This undermines the initial suggestion that being directly related to objects and events explains the appearance of directness of perception and of memory.

The discussion above describes some epistemological, metaphysical, and phenomenological problems that arise in the context of the dispute between direct and indirect realism. Despite the initial questions motivating these approaches to perception and memory being very similar, the way the literature on perception and memory developed in recent decades has been different. On the one hand, philosophers of perception have moved away from discussions focused solely on the question of the objects of perception, to more general discussions about issues pertaining to the metaphysical and phenomenological aspects of perception. On the other hand, philosophers of memory have maintained the traditional framework of direct and indirect realism, only starting to discuss the metaphysics and the phenomenology of memory more systematically in the past couple of decades.

The goal of this thesis is to explore the disagreement between the contemporary descendants of indirect and direct realism, the views which have been called, respectively, *representationalism* and *relationalism*, and to propose a reconciliation of those views with the purpose of advancing the debates in perception and memory.³ For the rest of this in-

³I should note that my discussion will be focused on a set of specific views that have been developed in the context of debates regarding perception in the past few years in the subfield of philosophy of perception. As such, I will follow the standard way to characterize the main views being discussed in the debate. By clarifying this point, I want to make it explicit that I do not take this way of “mapping” the terrain to be

roduction, I will explore how the traditional dispute between direct and indirect realism relates to ongoing disputes in the philosophy of perception and in the philosophy of memory.⁴ Afterwards, I will move on to consider hybrid views that attempt to reconcile the contemporary versions of those views. In doing so, I will set the stage for the development of the main argument of the thesis, which is that hybrid accounts of both perception and memory should be preferred.

2. Perception

In the contemporary perception literature, it has been a commonplace to assume that debates about perception have moved away from discussions about the (in)directness of perception to broader discussions about the nature of perception. The questions moving these discussions have to do not only with the epistemology of perception, but also with the metaphysics and the phenomenology of perception. The main alternatives currently on the table to answer these questions are *representationalism* and *relationalism*.

2.1 Representationalism

Representationalism is the view that, at the most fundamental level, perception is a form of intentionality — which is why the view is also sometimes called intentionalism (e.g., Searle 1983; Byrne 2001; see also Crane and French 2017) — for perception represents the world as being a certain way similarly to other mental states that are described in intentional terms, such as beliefs and desires. The characterization of perception in terms of intentionality has led philosophers to claim that representationalism can hold on to a form of direct realism. This is because, according to its proponents, representationalism eliminates the “veil of perception” standing in between the mind and the world, which has been the source of the problems raised for indirect realism. And this is possible because, following Brentano

exhaustive, in the sense that some theories of perception, especially those that have not been developed in the context of this debate, might not be appropriately represented.

⁴For a more detailed discussion of this dispute and its implications for contemporary debates about perception, see Fish (2010) and Crane and French (2017). While there are no contemporary systematic reviews of the discussion in memory, see Bernecker (2008, chs. 5 and 6), Senor (2014, Sect. 1), and Michaelian and Sutton (2017, Sect. 5) for some useful discussion.

(1973), representationalists say that representing an object *O* as being *F* does not imply that there is an object *O* that is *F* which is the object of awareness of perception. The objects of perception are, instead, the real objects themselves, and we become aware of those objects when the content of our perceptual experiences is satisfied by the relevant state of affairs in the world. However, when this does not happen, it is not the case that we are aware of non-material or non-existent things, as sense-data theorists suppose (e.g., Russell 1912; Broad 1925; Ayer 1956), but rather that we fail to be aware of anything. For example, if my perceptual experience represents a red book sitting on the table and there is in fact a red book sitting on the table, the experience is successful in making me aware of the red book itself. However, if I have an experience that is phenomenologically indistinguishable from the previous one, but there is not a red book sitting on the table, representationalists deny that I am perceiving a non-material or non-existent red book, saying, instead, that my representation failed to make me aware of an object in my environment.

Despite the attempts of contemporary philosophers to formulate representationalism as being a form of direct realism, it is not entirely clear whether they have been successful in doing so. While some, such as J. L. Austin (1962), have argued that it is not theoretically useful to distinguish between ‘direct’ and ‘indirect’ perception, others have maintained that there is an important sense in which representationalism amounts to an indirect realist view of perception (Snowdon 1992; Martin 2017; Travis 2017). Snowdon (1992), for example, argues that because content is supposed to be abstract, perception is not capable of putting us in a position to make demonstrative judgments about particulars in the surrounding environment (see also Campbell 2002; Schellenberg 2016). In this sense, Snowdon says that perception only puts us in an indirect relation to what is perceived. More recently, Travis (2017) has made a similar point when he says that, despite not postulating the existence of non-physical objects standing between perception and the world, requiring the presence of content in perception still implies a form of mediated relationship between mind and world.

Besides, there are other reasons why representationalism should not be viewed as a form of direct realism. These have to do with the three problems introduced above. Consider, first, the phenomenology problem. Because characterizing content in abstract terms is cen-

tral to the idea that not all perceptual experiences need to involve a relation of awareness between a subject and an object, representationalism faces trouble to explain why it seems to us that we perceive particular, as opposed to abstract objects. Thus, it is not obvious how this phenomenological particularity of perception (Schellenberg 2010) can be explained by saying that perception is characterized in terms of content.

Consider, second, the fact that the advantages often ascribed to representationalism are similar to the advantages traditionally associated with indirect realism. As Martin (2017, 270) notes “the typical concerns that intentional theorists appeal to in order to make us recognize the intentionality of sense perception are the kinds of concerns that moved sense-datum theorists at the beginning of the twentieth century”. More specifically, the strategy adopted by representationalists to deal with the error problem is very similar to to strategy adopted by sense-data theorists. While the former talk about content and the latter talk about non-material objects, the appeal to introduce such entities is similar, that is, the assumption that there is something that is shared, at the most fundamental level, by successful and unsuccessful forms of perception. For sense-data theorists, what is shared is the fact that both involve an act of awareness to objects, which is why sense-data are introduced. For representationalists, in contrast, what is shared is the fact that both successful and unsuccessful perception possess content. Thus, the motivation to postulate a common kind factor shared by successful and unsuccessful perception is at the bottom of the efforts of both contemporary representationalism and traditional indirect realism.

I shall emphasize that I am not suggesting that contemporary representationalism and traditional indirect realism are, in a straightforward sense, the same views. The motivation to think of perception in terms of intentionality is, indeed, one important difference between those views. I am also not trying to provide a definitive argument for the view that representationalism is a form of indirect realism. What I want to highlight is, instead, that the proximity, or the continuity, as Martin (2017) puts it, between these two views is more significant than what is acknowledged by most contemporary philosophers of perception. While acknowledging this proximity provides in itself an important contribution to the philosophy of perception, as is shown by the works of Snowdon, Martin, and Travis, the reason

why I bring it up here is primarily instrumental. As will become clear later, an important part of the work of this thesis will be based on the assumption that current debates about representationalism and relationalism in philosophy of perception can inform similar but underdeveloped debates in the philosophy of memory. Since, in the memory literature, the prevalent framework is still the traditional one of indirect and direct realism, the parallels that I trace between representationalism in perception and representationalism in memory will be more meaningful if the former can be viewed, at least for the sake of the overall framework of the thesis, as a form of indirect realism.

2.2 Relationalism

Relationalism is the view that, at the most fundamental level, perception is a form of relation to the world (Campbell 2002; Martin 2004; Brewer 2007; Fish 2009). Relationalists inherit the idea that perception is a form of awareness of objects from sense-data theorists, but unlike them, they believe that relations of awareness only obtain in what are usually described as successful or good cases of perception. Because in unsuccessful or bad cases of perception — i.e., hallucinations — there is not a relation of awareness obtaining between a subject and an object, such cases are said to be occurrences of different kinds from successful or good cases.

Relationalism provides a straightforward direct realist view of perception. However, as Genone (2016) rightly notes, it goes beyond traditional direct realist views in that it is not only concerned with the problem of the directness of perception, but also with providing an account of how the objects of perception — i.e., the particulars in the world — play a fundamental role in any complete account of the epistemology, the metaphysics, and the phenomenology of perception.

To see this, consider the knowledge problem. Campbell (2002) has famously argued for relationalism on the grounds that it provides a neat account of how perception grounds singular thoughts, or thoughts about particulars in the world, because perceiving puts us in direct contact with those particulars. In terms of Snowdon's (1992) characterization, relationalism explains how perception puts us in a position to make demonstrative judgments

about the environment, and, as such, to form knowledge about it. The same has been said of relationalism in relation to the phenomenology problem. Because relationalists see particulars as constitutive parts of perception, the external objects “shape the contours of the subject’s conscious experience by actually being the contours of the subject’s conscious experience” (Fish 2009, 6; see also Martin 2004; Brewer 2007). This provides a simple account of why it seems to us that we perceive particulars and, moreover, why we perceive them as being directly available to us.

Despite these advantages, the biggest challenge faced by relationalists is, just as with traditional direct realism, the error problem. Recall that the main motivation driving sense-data theorists to postulate the existence of sense-data was that the objects that one seems to be aware of in unsuccessful perception are not real objects. Since relationalists hold on to the idea that successful perception involves a relationship of awareness, they need an account of error that avoids the postulation of sense-data if the view is to be successful. One common strategy adopted by relationalists to deal with this problem has been to deny that unsuccessful perception involves a relationship of awareness to objects. By adopting a disjunctivist theory of perception, where successful and unsuccessful perception are viewed as mental occurrences of different kinds, relationalists explain away occurrences of hallucinations by claiming that they are not genuine cases of perception. Despite being introduced by Hinton (1967a; 1967b) in the 1960s, disjunctivism only got traction more recently with the works of Snowdon (1980) and Martin (2004) (see also Byrne and Logue 2008; Soteriou 2016). While relationalists have been consistent in their defenses of disjunctivism (see, e.g., Fish 2009), this solution has been met with skepticism by non-relationalists. One general worry has been that disjunctivism is counterintuitive, for it underplays the metaphysical importance of the phenomenological similarities between successful and unsuccessful perception. Another worry is that disjunctivism is not appealing in itself because it goes against both our common sense and our empirical understandings of perception.

I shall emphasize, again, that I am not suggesting that relationalism and traditional direct realism are the same views. The goal of this discussion is, rather, to show that similar motivations and problems that were present in the traditional debate are also present in

the contemporary debate. As I mentioned above, acknowledging this theoretical proximity will be important to properly understanding the motivations behind my proposal of using discussions in contemporary philosophy of perception as starting points to promote similar discussions in the philosophy of memory.

3. Memory

While the perception literature attempts to depart in important senses from the traditional framework of direct and indirect realism, the memory literature still mostly relies on it. In particular, the question of whether the objects of memory are mental representations of events or the events themselves is still very important, although this has not been properly acknowledged by contemporary philosophers of memory. The problems faced by direct and indirect realism about memory are, however, similar to the problems faced by representationalism and relationalism in the contemporary debate about perception.

3.1 Representationalism

In the philosophy of memory, representationalism or indirect realism is the view that the objects of memory are mental representations of past events. Hume (2011) is, perhaps, the most well-known defender of representational or indirect realist views. According to Hume's empiricist theory of memory, memories are less vivid and less intense forms of perception, thus making mental representations the objects of memory. In the early twentieth century, Russell (1912) and Broad (1925) have also proposed similar accounts of memory.⁵

Another form of representationalism, which has been very influential in the contemporary literature, is implied by the causal theory of memory proposed by Martin and Deutscher (1966). The causal theory's central claim is that a subject counts as remembering if and only if the subject's current mental state is causally related, in an appropriate manner, to the past events themselves. But not only this, the actual mental state must be a representation that is sufficiently similar to the past representation of the event in perception. Thus, the theory

⁵See Holland (1954) and Michaelian and Sutton (2017) for more detailed discussion.

assumes a form of representationalism or indirect realism about memory, for subjects become aware of the past events only by being indirectly aware of actual representations of them.⁶

Representationalism about memory has been very popular and it is often taken for granted in philosophical debates about memory. In part, this is because it provides a simple and intuitive answer to the error problem and because it is taken to be the most adequate theory in the context of empirical research. For example, it has been suggested by a number of different researchers that memory is a result of a more general cognitive capacity to think about events, whether those events occurred or not (see, e.g., Suddendorf and Corballis 2007; Schacter et al. 2007, 2012; De Brigard 2014a; Michaelian 2016c). Thus, because relationalism seems to require the existence of the objects of memory — and hence of the objects of other forms of episodic thinking — the view simply appears unappealing to most people when considered in relation to these results. Another issue has to do with the fact that it is now widely known that, rather than simply retrieving information from the past, memory is constructive, in the sense that not all the information that one remembers needs to be originated in one's past experience of the relevant event (see, e.g., Schacter et al. 2007; Michaelian 2011, 2016c; Robins 2016b). This suggests that memory need not be able to reach back to past experiences or events, but rather construct plausible and actual representations of what happened. For this reason, a form of representationalism seems to be the most simple and intuitive option, as representations can be constructed by the brain in the absence of the represented objects. A third issue is that representationalism seems to accommodate memory errors better. By memory errors, I mean occurrences of *misremembering* (see, e.g., Robins 2016a; Michaelian 2016b), such as when one remembers some details of an event incorrectly — e.g., remembering having strawberry cake at your birthday party, when you had chocolate cake — and occurrences of *confabulation* (see, e.g., Hirstein 2005; Michaelian 2016b; Bernecker 2017; Robins 2017a), such as when one re-

⁶It is important to note here that not all causal theories of memory are representationalist, for the requirement for causal connections between memories and events does not by itself entail a form of representationalism (see, e.g., Debus 2008). However, since Martin and Deutscher's (1966) version of the causal theory, which happens to be the most influential one in the contemporary debate, clearly assumes representationalism, most philosophers have been tempted to accept the latter because they are committed to the former.

members events that did not happen — e.g., remembering having a birthday party when there was no such party.

Despite its advantage in accounting for the error problem, because representationalism is a form of indirect realism, it will face similar problems to the ones discussed initially. With regard to the knowledge problem, the challenge is to provide an account of how episodic memories allow us to know something about the past. That is, because representations mediate our contact with the past events, we cannot be sure that what we are aware of in remembering actually corresponds to how the past was. With regard to the phenomenology problem, representationalism faces trouble to explain the directness of memory. When we remember past events, we seem to be remembering the events themselves, and not representations. However, if the objects of memory are mental representations, it is hard to see how that can be the case. So, in the case of memory too, the success of contemporary versions of representationalism depends greatly on how they will deal with the problems that traditional indirect realist views have faced.

3.2 Relationalism

In contrast to representationalism, relationalism or direct realism in philosophy of memory is the view that the objects of memory are the past events events themselves and, therefore, that our relationship to them is unmediated. Despite the attempts of philosophers such as Thomas Reid (2000) and, later, of Russell (1921) and John Laird (1920), relationalism about memory has not been very popular in the past few decades. Recent proponents of the view include Dorothea Debus (2008) and Sven Bernecker (2008). Drawing inspiration from relationalists in perception, especially from Martin (2004), Debus argues that, when one successfully remembers, one stands in a recollective relation to the past events that is absent in unsuccessful forms of remembering. As she puts it, successfully remembering puts one in a direct relation to the past events themselves. Bernecker (2008), in contrast, pursues a different strategy. He argues that, while memory involves representing the past, the representation itself is not the object of mnemonic awareness. That is, remembering makes us aware of the past events themselves by means of being related to representations

of those events.

Although, on the one hand, Bernecker takes himself to be a direct realist, and thus arguably a relationalist, his view seems to be more accurately described as a form of representationalism along the lines of representationalism in perception. In other words, Bernecker seems to accept that, in remembering, we are related to representations, but denies that this makes our relationship to the past events mediated. However, it is not clear whether this consists in a genuine form of direct realism or relationalism, for representations still seem to stand in between one's memory and the past events themselves. Debus, on the other hand, seems to provide a view that is clearly relationalist. However, the kind of analysis of remembering suggested by her view has been met with a lot of skepticism. This might be due to two aspects of the current debate in the philosophy of memory. The first is the popularity of the causal theory, which relies on a form of representationalism, in philosophy of memory. The second is the idea that representationalism is better supported by the empirical research discussed above. In addition to that, the prospects of relational views in relation to empirical research on memory have not been addressed in enough detail. This is, in part, because there are not many relational views out there, and also because those who have proposed relational views have not explored this topic in detail. In this respect, the contemporary debate on memory has not developed in the same way that its perception counterpart has.

A third problem is the fact that relationalism seems unable to handle memory errors. As I argue in Chapter 4, relational theories in their current form cannot distinguish properly between successful and unsuccessful cases of remembering. Moreover, I show that, even if disjunctivism is in place, relationalists still face trouble to explain a particular form of memory error, that is, misremembering (Robins 2016a). Despite all these problems, relationalism is still appealing for some philosophers for similar reasons that traditional direct realism about memory is appealing. In being a form of direct realism, relationalism provides a straightforward answer to the knowledge and phenomenology problems. That is, because the objects of memory are the past events themselves, there is no mystery in how we can know something about the past and in why it seems to us that we are directly

remembering those events.

4. Towards a hybrid approach

Until recently, many philosophers have taken representationalism and relationalism, both in perception and in memory, to be incompatible with each other. However, it is not entirely clear why they take this view. For example, in the perception literature, hybrid views started to emerge recently, where elements of representationalism and relationalism are put together in order to form unified and novel views (see, e.g., Siegel 2010; Schellenberg 2010, 2014; Logue 2013, 2014; Hanna 2015). This suggests that at least some elements of representationalism and relationalism are not incompatible with each other. While there is still much work to be done to identify which of those elements can be integrated, the prospects for hybrid views of perception seem to be promising at this stage. The literature on memory, in contrast, has not paralleled the developments in perception. As I said before, with the exception of a handful of philosophers, the standard assumption seems to be that some form of representationalism is correct. Again, this is not so much due to the fact that relationalism has been shown to be problematic, as to the fact that this question has not received the same amount of attention that its counterpart in the perception literature has. Thus, the possibility of hybrid views of memory remains largely unexplored to date.

One particular hybrid proposal, which is defended by Susanna Schellenberg (2010; 2011; 2016), provides a potential useful framework to think of hybrid views in perception, and, as I will argue, in memory too. Although Schellenberg's hybrid view shares important elements with traditional representational views, she argues that some forms of relationalism are not incompatible with it. Her claim is that perceptual representations can be object-involving, in the sense that, if the particular objects that we perceive fail to exist, the content of our perceptual representations will vary in significant ways. While defending this claim requires some bold moves, especially with respect to the relationship between the representational content and the phenomenal content of perception, if right, her view provides an overall account of how perceptual representations can be inherently relational.

This account, I will argue, can be extended to the case of memory, which provides a helpful starting point to think about hybrid views of memory. The idea is that, like perceptual representations, mnemonic representations can be object-involving, or, to be more precise, event-involving, in the sense that, if the particular events that we remember fail to exist, the content of our mnemonic representations will vary in significant ways. This will require making similar moves to the ones that Schellenberg does in the case of perception, but again, if successful, this adapted view will provide an overall account of how mnemonic representations can be inherently relational. One particular advantage of this hybrid approach, as I discuss in Chapter 5, is that it is compatible with a broader theoretical framework in psychology, according to which memory is just a form of mental time travel (see Suddendorf and Corballis 1997; Tulving 2002). That is, it respects the theoretical intuitions from other domains, while preserving important features of representationalism and relationalism. In this sense, the hybrid view provides an important resource for those sympathetic to relationalism to reconcile it with empirical research on memory.

This brief discussion makes explicit that there are still important questions to be resolved when it comes to the dispute between representationalism and relationalism both in perception and in memory. In the perception literature, the question of what hybrid views of perception will look like is only now starting to be addressed, which opens up different possibilities for future research. In the memory literature, there is the question of whether the general assumption that a form of representationalism is correct is warranted, and if not, whether its alternative, relationalism, can provide a coherent view in light of empirical research. In addition, there is a more general question, which has not been asked so far in the literature, of whether hybrid views of memory are possible, as perhaps even plausible. Given the prospects of this question in the perception literature, it seems to make sense to ask it in the case of memory too.

Thus, given the current context of the dispute between representationalism and relationalism is situated today, this thesis consists in an attempt to move the debate forward. By addressing the questions briefly discussed here, the main idea defended is that a form of hybridism about perception and memory is the most promising view. Rather than provid-

ing a final hybrid view of perception and memory, the general suggestion of the thesis is that, instead of seeing representationalism and relationalism as two opposite and irreconcilable views, philosophers will benefit from exploring ways in which these views can be reconciled.

More precisely, the thesis has one major goal, which is to enhance our understanding of the dispute between representationalism and relationalism, both in perception and in memory, by defending a hybrid approach. Two sub-goals are, accordingly, subsumed to the main goal.

1. The first is to contribute to the developing trend of hybrid views in the philosophy of perception. In particular, I will propose a new hybrid view of perception, which draws inspiration from a philosophical tradition that has been underexplored in contemporary analytic philosophy, namely, classical pragmatism. The option for pragmatism is due to the fact that one of its founders, Charles Peirce, was one of the first philosophers to propose a hybrid view of perception. Thus, I provide an interpretation of Peirce's theory of perception that engages with the contemporary literature in philosophy of perception.
2. The second is to advance the debate about representationalism and relationalism in philosophy of memory. The work done in philosophy of perception will serve as a starting point to conceive of the problems in the memory literature. The suggestion will be that a form of hybridism is also the best alternative for philosophers of memory. However, since no one has proposed hybrid views in the literature, my efforts will be directed at proposing a framework, rather than a definitive view, for the future development of hybrid views of memory.

The focus on Charles Peirce's pragmatism in my approach to perception is due to two reasons. The first reason is that Peirce himself attempted to bridge traditional direct and indirect realism about perception (Haack 1994; Bergman 2007; Legg 2014b; Wilson 2016), which naturally provides a potentially useful starting point to think about hybrid views. The second is that, as some philosophers have recently noted (Genone 2016; Locatelli and

Wilson 2017), contemporary hybrid views of perception have been biased toward representationalism, and one important question open in this literature is precisely whether alternative views are possible. As it will become clear in Chapters 1 and 2, I take Peirce's pragmatic theory of perception to offer such an alternative, that is, a hybrid view that is not biased toward representationalism and consequently one that can be potentially appealing to philosophers who are sympathetic to relationalism. Peirce's pragmatism thus provides the foundation to understand the novel contributions made by my hybrid view of perception.

5. The thesis

The thesis is organized as a collection of five original and self-standing articles dealing with different issues pertaining to representationalism and relationalism in perception and in memory. Chapters 1 and 2 compose Part I of the thesis, which is dedicated to perception. Chapters 3, 4, and 5 compose Part II, which is dedicated to memory.

Part I – Perception

Chapter 1 The first chapter, *Contextualizing perception*, is an early attempt to deal with some of the issues discussed above in perception. It employs some general classical pragmatist ideas to provide a relationalist theory of perception. Although the explicit suggestion is that a form of relationalism should be preferred over representationalism, the relational view defended there departs radically from most relational views in the contemporary literature. Because one central aspect of my proposal is that disjunctivism, which is assumed by virtually all contemporary relationalist views, should be abandoned, the relationalist view that I offer already indicates an inclination to incorporate insights from representationalism.

Chapter 2 The second chapter, *Perception pragmatized*, explicitly proposes a hybrid view of perception based on the ideas of the pragmatist philosopher Charles Peirce. While the influence of Peirce's pragmatism was only implicit in Chapter 1, his ideas are developed here in more detail and in relation to contemporary hybrid views of perception. More

specifically, I argue that my hybrid view, perceptual pragmatism, offers a reconciliation of representationalism and relationalism that does not underplay the importance of relationalism, as is the case with some competing hybrid views. This chapter offers, therefore, an updated version of the relational view offered in Chapter 1, but now explicitly recognizing the importance of representationalism.

Part II – Memory

Chapter 3 The third chapter, *Thinking about events*, is more speculative in nature. It investigates a topic that has not been explored so far. Moreover, it represents an important shift in the thesis: namely, the ideas developed in Chapters 1 and 2, which pertain to perception, are now put to work in relation to memory. This chapter, which was co-authored with Kourken Michaelian, proposes a novel question in the philosophy of memory. This is the question of what the objects of episodic hypothetical thinking are. Because some have suggested that episodic memory is just another form of thinking about hypothetical scenarios, the chapter explores the prospects of representationalist and relationalist accounts of episodic hypothetical thinking. The argument offered there is that neither representationalism nor relationalism alone can provide a satisfactory account of episodic hypothetical thinking. The alternative proposal that we put forward is that a form of pragmatism about episodic hypothetical thought, which is influenced by the account of perception developed in Chapters 1 and 2, provides a more promising account of those mental states and their objects. This allows us, finally, to extract a pragmatist account of memory from the more general pragmatist account of episodic hypothetical thinking.

Chapter 4 The fourth chapter, *Failing to remember*, explores an important topic in contemporary philosophy of memory. This topic is closely related to one prominent discussion in the philosophy of perception, which refers to the possibility of a relational account of non-veridical perception, such as illusions and hallucinations. The chapter thus raises a parallel issue for relationalists about memory. More specifically, the argument proposed is that, given the occurrence of unsuccessful remembering, such as misremembering and

confabulations, relationalism cannot distinguish properly between these cases and cases of successful remembering. I focus on one prominent relational view, developed by Dorothea Debus, to show that, if a relational account of memory is to succeed, it needs to deal successfully with unsuccessful remembering.

Chapter 5 The fifth chapter, *The hybrid contents of memory*, adopts a more focused and positive strategy, dealing explicitly with the question of whether a hybrid view of memory is possible. This chapter attempts to use the insights from hybridism in philosophy of perception to begin to develop hybridism in philosophy of memory. I discuss one prominent hybrid account of perception, developed by Susanna Schellenberg, and expand its scope to the case of memory. In particular, I offer an account, based on Schellenberg's work, of why representationalists and relationalists disagree in the case of memory. I then propose an adaptation of her account of perceptual content to address the relevant issues in the case of memory. Besides dealing with the dispute between representationalism and relationalism, the hybrid view developed here also provides insightful prospects for contemporary discussions in philosophy of memory. That is, building on some of the work developed in Chapter 3, I propose that a hybrid account of memory can give us a coherent picture of the relationship between memory and other forms of episodic thinking.

Appendix Finally, in the short discussion paper *Mental time travel and the philosophy of memory*, which is included as an Appendix, I offer an overview of the problems raised by the idea that episodic memory is a form of mental time travel in philosophy of memory.

Part I

Perception

Chapter 1

Contextualizing perception

Introduction

This chapter is an early attempt to discuss, in the domain of perception, some of the problems outlined in the Introduction. I develop and defend a novel relationalist theory of perception — perceptual contextualism — grounded in pragmatist ideas. The central claim of perceptual contextualism is that the distinction between veridical and non-veridical perceptual experiences should be understood in terms of the capacity of the former to generate pragmatic agreement among perceptual subjects in relevant contexts. Perceptual contextualism endorses the core relationalist claim that the objects of perception are constitutive parts of perceptual experiences but denies the assumption, made by many relationalists, that those objects are necessarily ordinary material objects. I argue that the assumption needlessly restricts the explanatory reach of relationalism and show how the notions of context and pragmatic agreement provide a means of distinguishing between veridical and non-veridical experiences that allows us to expand the explanatory reach of relationalism by avoiding commitment to any particular view about the nature of the objects of perception.

Chapter 1 is not only the first in terms of structure, but also the first in chronological order. It is my first attempt to bring two bodies of literature together: contemporary philosophy of perception and classical pragmatism. At this early stage, the connection between them is only implicit, though. While the influence of Charles Peirce's pragmatism is clear

in various parts of the chapter, I did not explore the actual connections between his writings and current discussions about perception in any further detail. This was done in Chapter 2. Thus, the reader should keep in mind that most of the ideas developed in the Chapter 1 are still embryonic. Some of them have, in fact, been abandoned along the way, but some others have been developed in more detail in Chapter 2, where I present my final take on perception in the thesis. For example, in Chapter 1, I take the enterprise of providing a pragmatist-inspired account of perception to be fundamentally a relationalist enterprise. However, this idea is dropped in Chapter 2, for a closer study of the works of Charles Peirce made me realize that his pragmatist theory of perception is better understood as a hybrid representational-relational view. Despite this important shift, various important ideas that were introduced in this chapter re-appear in my discussion of perception in Chapter 2. In particular, the notions of *pragmatic agreement* and *pragmatic context* are developed in more detail there. Moreover, the denial of disjunctivism, which is first discussed in Chapter 1 and constitutes one of its main motivations, is also a central motivation for my discussion of perception in Chapter 2. Thus, despite not being my final discussion of perception, Chapter 1 plays an important instrumental role in the overall scheme of the thesis, for it sets the stage for most discussions developed in Chapter 2. As such, it provides potentially useful insights into the motivations of the work done Chapter 2.

Outline

This chapter develops and defends a novel relationalist theory of perception — perceptual contextualism — that is grounded on pragmatist ideas. Section 1.1 reviews relationalist and representationalist approaches to the problem of perceptual error. Section 1.2 develops perceptual contextualism as an alternative relationalist view. The central idea of the view is that veridical experiences are those that would generate pragmatic agreement among relevant perceptual subjects. Perceptual contextualism is also discussed in relation to alternative relationalist theories. Section 1.3 concludes by replying objections to perceptual contextualism raised in relation to illusions and hallucinations.

1.1 Perceptual error and the nature of perceptual experiences

The question of whether perception is fundamentally relational has been widely debated in contemporary philosophy of perception. Among those who give a negative answer are representationalists, according to whom perceptual experiences are fundamentally characterized by their representational content (see Searle 1983; Harman 1990; Dretske 1997, 2003; Byrne 2001, 2009). In opposition, relationalists give a positive answer, claiming that perceptual experiences are fundamentally characterized by their relations to external objects (see Snowdon 1980; Travis 2004; Fish 2008, 2009; Brewer 2007, 2011, 2013).¹ In historical terms, this debate relates to but does not reduce to the older dispute between direct and indirect realist theories of perception. On the one hand, relationalists explicitly maintain that at least the fundamental aspects of perception are determined by its relation to external objects. On the other hand, while representationalists are not indirect realists in the strong sense of the term, i.e., they do not accept that in perception we are related to mediating entities, they hold that at least the fundamental aspects of perceptual experiences are determined by their representational content. The question of how to account for perceptual error thus plays a central role in the debate between representationalism and relationalism, just as it did in the dispute between indirect and direct realism.

1.1.1 Perceptual error

Representationalists and relationalists have given opposing accounts of perceptual error. The view that I defend here — perceptual contextualism — is a relational theory of perception that offers an alternative to current relationalist theories. For this reason, it will be

¹This way of setting the stage is not uncontroversial at all. Some representationalist views understand perceptual content as being Russellian propositions or Fregean senses, meaning that the objects and properties represented are viewed as constitutive parts of perceptual content (e.g., Speaks 2009; Chalmers 2004; Schellenberg 2010, 2011). Other philosophers have argued that representationalism and relationalism are not ultimately incompatible, defending “hybrid” accounts of perception (e.g. Schellenberg 2010, 2011; Siegel 2010; Logue 2014; Hanna 2015). I will not be concerned with these views here, so when I use the terms “representationalism” and “relationalism”, I will be referring to what Schellenberg (2010; 2011) calls “austere representationalism” and “austere relationalism”, i.e., views of perception that accept either that objects or representational content fundamentally and exclusively constitute our perception of the world.

useful to review those accounts briefly.

Perceptual error can be understood here as referring to cases in which one has a perceptual experience that fails to capture, completely or partially, the way the world is. For representationalists, perceptual experiences are fundamentally characterized by what they represent to be the case (see Searle 1983; Byrne 2001, 2009). When John sees a white cat, John's visual experience represents a cat as having the property of being white. On this view, perceptual errors happen when an experience misrepresents the world, meaning that it is inaccurate with respect to how the world is. Thus, if John has a visual experience of a white cat, but no such object exists in the world, or if the cat is of another color, then his experience represents the world inaccurately. Alternatively, we can say that the satisfaction conditions of John's experience (i.e., the conditions under which his experience would be accurate) are not met.

Relationalists hold that it is not adequate to think of perception as being representational. For example, Charles Travis (2004) argues that perception does not by itself establish satisfaction conditions. When I see an apple sitting on the kitchen table, this experience can be only accurate or inaccurate relative to something else, such as my prior beliefs about how the world is. If I know that my partner bought some wax fruit to decorate the house the day before, then I would be inclined to think that although it seems to me that I see a real apple, I might be mistaken. If no such information is available, then I will probably just infer that what I see is a real apple. No intrinsic feature of my experience can establish its satisfaction conditions, but only its relation to other mental states, such as my beliefs. Therefore, perception is not representational.

Although Travis (2004) does not explicitly endorse the view that satisfaction conditions are determined by other mental states, it has been defended by other relationalists (see e.g. Antony 2011 and Genone 2014). On these accounts, perceptual errors would be due to those mental states and not to perception itself. If the (non-)veridicality of a mental state depends on whether its satisfaction conditions are met in the world, and if perception cannot establish those conditions alone, then perception cannot be said to be veridical or non-veridical. Consequently, perceptual error is due to those mental states and not to perception.

The above conception of perceptual error is congenial to relationalism because it implies that perceptual error is not really “perceptual”, which implies, in turn, that even non-veridical perceptual experiences can be relational now (see Martin 2004; Brewer 2007, 2011; Fish 2008, 2009). A further development of this approach is provided by James Genone’s (2014) doxastic account of illusions. Genone accepts Travis’s claim that perceptual experiences do not establish satisfaction conditions and then proceeds to explain cases of illusion, such as when a straight stick appears bent when one sees it half-submerged in water, by claiming that they are caused by doxastic states accompanying the relevant perceptual experiences.

To make the doxastic account fully compatible with a relational view of perception, Genone (2014) introduces the notion of an “appearance property”. Appearance properties are manifestations of the intrinsic properties of objects that are displayed or not depending on the contexts in which the objects are perceived. It is important to note that, despite the fact that they are context-dependent, Genone argues that appearance properties should not be viewed as mind-dependent;² they are dependent on physical contexts and not on minds.

Appearance properties play a central explanatory role in Genone’s account because they allow us to explain error without having to appeal to representations. For example, when a straight stick appears bent when one sees it half-submerged in water, the stick displays the appearance property of being bent in that particular physical context. The doxastic account says that this is a “perceptual illusion” because the experience is accompanied by a mistaken belief, namely that the stick should continue to be bent.³ Therefore, if perceptual error is due to doxastic states, we need not invoke representations in order to explain how it occurs.

One problem with the doxastic account is that it is not clear why relationalists should

²One important worry here is that it is not clear that intrinsic properties appearing in different contexts implies that they are context-dependent. I am taking this for granted here, since it is Genone’s view on the subject. Later on, I will argue that we cannot think of context-dependent properties as being intrinsic properties unless we make unwanted changes to our understanding of the latter.

³One might argue that this misreads the relationalist core account of perceptual error, for the relationalist does not understand error as a cognitive process that happens *after* or separately from perception. I am not saying here that error is a cognitive process, but rather that perceptual error happens because of doxastic states that are somehow involved in perception. The distinction I am trying to convey here is analytic, that is, it identifies different logical elements of perception, but it does not say that they are separated in reality.

adopt it. The answer can be found in a commitment taken by most relationalists, namely, that the objects of perception are ordinary material objects (Brewer 2007). This view, often called *metaphysical naïve realism* (hereinafter *naïve realism*), states that the objects of perception are ordinary mind-independent objects, such as chairs and tables. Once relationalism and naïve realism are conjoined, it becomes clear why the doxastic account appeals to relationalists; one can explain how perceptual error is possible without appealing to representations. In perceptual error, we are related to ordinary material objects, but error is due to doxastic states, and not to perception (or its objects). In the next section, I shall argue that, despite the fact that the doxastic account fits well with relationalism, relationalists should avoid adopting it, for it requires an overly restrictive view of the objects of perception.

1.1.2 Relationalism and naïve realism

I have suggested that the doxastic account becomes attractive to most relationalists because they are committed to a particular view of the objects of perception, namely, naïve realism. But these are independent views. Relationalism says that perception is fundamentally shaped by its relation to objects. Naïve realism says that those objects are ordinary mind-independent objects. While naïve realism implies relationalism, the contrary does not hold, for relationalism, strictly understood, makes no claim about the nature of the objects of perception. In this section, I argue that relationalists should reject both the doxastic account and naïve realism. I provide two different arguments: the first shows that, instead of saving naïve realism, the doxastic account makes it implausible. The second says that naïve realism should be abandoned as a view of the objects of perception because it provides an overly restrictive account of those.

Let me start by clarifying what I understand by *naïve realism*. Naïve realism, as stated above, is the view that the objects of perception are ordinary material objects. It can be described by the two following claims

- (i) The objects of perception have *intrinsic properties* (i.e., they have mind-independent and non-relational properties);

- (ii) The objects of perception are *self-subsistent* or *self-contained* (i.e., they have a non-relational existence).

This definition requires further clarification. Intrinsic properties are understood as properties that objects have independent of being related to other individuals, such as the rectangularity of my laptop. This property is intrinsic because my laptop is rectangular independent of being related to something else. Moreover, its intrinsic properties allow it to have a self-contained or non-relational existence, meaning that the laptop is individuated as an object independent of its relation to other things. Thus, intrinsic properties are sufficient to individuate objects, but they do not necessarily provide a full description of them, for some properties, such as colors, can only be individuated in particular contexts. In sum, I will understand intrinsic properties as the properties without which an object could not be what it is. I shall call this the *Naïve Individuation View* (NIV).

There are two problems for relationalists who accept NIV and the doxastic account. First, it is not obvious whether these two are compatible at all. Genone's introduction of appearance properties can be seen as an attempt to address this worry. He takes appearance properties to be intrinsic *and* context-dependent simultaneously, because "part of what is involved in an object possessing certain intrinsic properties is being disposed to appear various ways in various conditions" (2014, 357). The problem with this claim is that it contrasts explicitly with the notion of intrinsic properties as being non-relational. If intrinsic properties involve dispositions to appear in different ways in various conditions, then (i) should be abandoned. But this would make NIV and the doxastic account incompatible. The problem here is not that it is implausible to think of objects as having relational properties, but why we should take relational properties to be intrinsic. Doing so seems to require an alternative notion of what an intrinsic property is, which does not help to solve the problem of reconciling NIV and the doxastic account.

The second problem is that relationalists who accept the doxastic account have to deal with difficulties pertaining to NIV. One of the difficulties is that NIV restricts the explanatory reach of relationalism because it offers an overly restrictive metaphysical view of the objects of perception. Casey O'Callaghan (2011; 2016) has recently noted that philosophers

of perception have focused excessively on vision, making them think of other perceptual modalities in analogy to it. But this is problematic, because it is not possible to individuate the objects of those modalities in the way that the objects of vision are individuated (see O’Callaghan 2011). For example, in audition we seem to be related to sounds instead of the objects causing them. When you hear a car crashing against a tree a few hundred meters away, the object of your auditory experience is a sound caused by the ordinary material objects and not the car and/or the tree. Unlike in vision, in which you are acquainted with ordinary material objects in your visual field, in audition you relate to sounds. But sounds are extended in time, meaning that they can only be fully experienced in more than one time slice. To see this, consider the differences in the sirens of police and ambulance cars. In some countries, those cars use the same type of sirens, but the way we experience them differs because of the variations in pitch and progress over time.⁴

The above shows us that an essential feature of the objects of auditory perception, i.e., their temporal extension, requires the presence of a perceiving subject in order to be properly individuated. This makes NIV inadequate to individuate auditory objects, for sounds are essentially relational entities, i.e., they only exist in relation to perceptual subjects. But this is incompatible with (i) and (ii), which suggests that NIV poses important constraints on relationalism.

One possible way to avoid this conclusion is to either claim that the objects of different perceptual modalities can be adequately individuated by NIV (or some reformulated version of it), or to explain the differences highlighted above in terms of appearance properties. The first alternative will only succeed if one can show us that (1) vision is metaphysically prior to other sensory modalities, and (2) that the same criteria used to individuate the objects of vision also apply to other modalities. While it seems uncontroversial that humans rely mostly on vision to guide themselves in the environment, it is not obvious why this should

⁴One might argue that this objection against NIV is misplaced, for nothing is said to motivate the view that NIV is visuocentric. Relationalists can consistently claim that besides relating to objects, perception also relates to events, e.g., in auditory perception (O’Callaghan 2011). However, this actually supports the argument for abandoning NIV. NIV is a commitment that most (but not all) relationalists make that is incompatible with other modalities of perception. Thus, it is hard to see how things such as events could be consistently individuated by NIV. Therefore, since we have good reasons derived from other perceptual modalities to abandon NIV, and since the doxastic account is an attempt to reconcile relationalism with NIV, the relationalist should feel suspicious about the doxastic account.

provide any reason to think that vision is metaphysically prior in relation to other perceptual modalities. For example, in modalities such as smell or taste, it is not obvious that NIV would be the best way (or even the most intuitive way) to individuate the objects of those modalities. Like in audition, when you smell a cheese cracker, it is more intuitive to think of the object of your experience as being a smell — something like a cheesy smell — rather than the cheese cracker itself. Were you deprived of vision, it would not be obvious to you that this smell is caused by an outside object, instead of merely being a quale inside your head. Moreover, other species rely primarily on other sensory modalities, such as the case of bats famously discussed by Nagel (1973). Therefore, even if vision is more fundamental in practical terms for humans, it does not follow that it is metaphysically prior to other sensory modalities. As a consequence, (1) requires further argument, but it seems hard to see a non-arbitrary way to settle this question. As for (2), if it is right that the criteria used to individuate visual objects cannot be extended to other sensory modalities, then unified strategies along these lines are very unlikely to succeed. Moreover, if we cannot make a substantial case for (1), then (2) seems to lose most of its appeal.

But what about the second alternative, that is, appealing to appearance properties? Does it shed any light on these difficulties? One could say that one is aware of the car crash because the car has an appearance property that is manifest only when the car crashes into a tree. While this view is plausible in principle, it complicates things in an undesired way for the naïve realist. This account would have to allow that physical objects have appearance properties that are spatially extended beyond the proximal spatial region that they inhabit — e.g., a property that, despite being a property of a particular object residing in a restricted spatial region, only manifests itself beyond the limits of that spatial region. Again, while this view might be plausible in principle, it is hard to see how to reconcile it with naïve realism. The difficulty with this strategy is that, in extending NIV to audition, important features of auditory perception are left unexplained. As O’Callaghan rightly points out, this strategy “fails to recognize the diversity of types of individuals that are objects of perceptual awareness” (2011, 151).

1.2 Perceptual contextualism

I have argued that the doxastic account is inconsistent with NIV. I have also argued that NIV limits the prospects of relationalism to individuate the objects of perception in modalities other than vision. This does not mean that relationalism is wrong, but only that it becomes problematic when associated with NIV. But since these two theses are independent of one another, one can still be a relationalist and abandon NIV. In this section, I argue for a relationalist theory that avoids ontological commitments, such as NIV, about the nature of the objects of perception. Moreover, to provide a relational account of perceptual error and the plurality of the objects of perception, this view introduces pragmatist elements to distinguish between veridical and non-veridical experiences and to individuate the objects of perception. I call it *perceptual contextualism*.

A central feature of perceptual contextualism is that it is not committed to any particular ontological view of the objects of perception. Instead, it characterizes the objects of perception in terms of how they would prompt subjects to act. Hence, perceptual experiences should be analyzed relative to what I will refer to as “pragmatic contexts”, which are understood in terms of the biological, evolutionary, and practical features of the settings in which experiences occur. Perceptual contextualism makes two important claims: first, that perceptual experiences are veridical if they would generate pragmatic agreement in pragmatic contexts, and second, that the objects of perception can only be individuated in pragmatic contexts. This section is dedicated to motivate and develop these claims.

1.2.1 The resistance view of experiences

Perceptual contextualism starts by proposing a new way to look at perception, one that considers the actions subjects perform on the basis of their experiences.⁵ In this respect, it opposes what I will refer to as the *Revelation View* of experiences. The Revelation View, which is implicitly accepted by most representationalist and relationalist views, says that perception is a *passive* relation to self-contained and non-relational objects residing in the

⁵For views that share similar motivations, see Nöe (2004) and Nanay (2013)

world (see Chirimuuta 2017 for a related discussion). On this view, perception is essentially a matter of receiving some external stimulation through our sensory organs, which is then processed internally to produce perceptual experiences of the world. The subject is essentially passive in this process, for perceiving only requires being stimulated by the world and processing the stimuli.

Perceptual contextualism introduces an alternative view that I call the *Resistance View* of experiences. The Resistance View opposes the Revelation View by ascribing the subject an essentially active role in perceiving the world. Besides how internal stimulation is processed, perception is also determined by how subjects would act on the presence of the stimuli. In this respect, the Resistance View denies that there is a sharp separation between perceiving and acting. It thus becomes attractive for two reasons. First, as I will argue, it provides a framework to distinguish between the objective and the subjective elements of perception in pragmatic terms. This is important because, in doing so, we can avoid appealing to ontological notions to individuate the objects of perception. Second, it provides an account of this distinction that is internal to the theory, meaning that its plausibility will not depend on the plausibility of other philosophical theses. Thus, the Resistance View does not depend on other views, such as NIV, to explain how the objects of perception are individuated.

The core claim of the Resistance View is that perception is essentially impositional. By “impositional” I mean that it offers a sense of external resistance to perceptual subjects. This resistance prompts subjects to act in the presence of experiences, thus revealing elements in perception that resist any conscious efforts to be changed. For example, in seeing an apple, I see an object with certain shape and color. If I wish that the apple becomes blue, the experience will not simply change because of my wish. This is what is meant by imposition or resistance here, namely, that some elements of our experiences do not depend on our minds to be presented to us as they are. They are simply imposed on our stream of consciousness.⁶

⁶In terms of its historical roots, the Resistance View is motivated by Charles Peirce’s view that perception is impositional, in the sense that subjects cannot control whether or how they will experience the world. This idea is developed in more detail in Chapter 2, Sect. 2.2, where I talk about the *percept* in Peirce’s theory of perception.

Now, perception is essential for subjects to act in the environment. These actions will be motivated by how subjects interpret the impositional elements. If I see a red apple, I will likely eat it, for it is a sign of a ripe fruit. Moreover, actions will have impact on other perceptual subjects. For instance, if a mother wants to teach her child that red apples are good to eat, she can put them next to spoiled apples and point out the relevant features to identify the good apples, such as their color. The child can then either identify the good apples correctly or fail to do so. In the former case, the child coordinates her actions to comply with the mother's instructions, meaning that her experience of the apple was similar to the mother's to the extent that it allowed for coordinated actions in that context. This case shows that, in having different experiences of the world, different perceptual subjects — the mother and the child — can interact based on what they perceive, which allows for the coordination of their actions to achieve a certain goal. Perception is therefore active, for our interactions with the world and other perceptual subjects determine how we perceive.⁷

The Resistance View says that the impositional elements are the objects of perception. More importantly, it says that the objects of perception are always objects *for* action. They are not, however, mere creations of the mind because they resist the attempts of subjects to change them. On the example above, when the mother and the child are successful in coordinating their actions, the Resistance View says that it is because their experiences have the same relevant impositional elements, and hence the same object. But these elements are not dependent on individual minds, for they are neither a creation of the mother's mind nor of the child's. In this sense, they are *external* to individual minds. Being external, however, does not imply that the object exists concretely, e.g., as ordinary material objects. If a child sees a unicorn, her experience does have impositional and external elements, but it does not follow that these elements exist concretely, i.e., that there is a unicorn. What will determine whether the object exists concretely is how the actions prompted by perceiving it would be coordinated with the actions of other subjects in the same context. For example, if the child tells her parents about the unicorn, but they fail to see it, then the object does not exist. The

⁷For a clearer and more detailed discussion of this topic, see Chapter 2, Sect. 2.3.2, where I talk about the phenomenology-first approach to perception.

general point here is that the existence of an object of perception depends on whether it could serve as basis for coordinating the actions of multiple subjects.⁸

To make this point clear, let us see what the Resistance View says of illusions and hallucinations. Consider a straightforward visual illusion, such as seeing an apple as being gray. When you see this object, it is presented to you with certain qualities, such as being round and gray. This experience has the impositional element of gray, for even if you try to change it, you will not be able to. Now, suppose that you tell your friends that you see a gray apple on the table. When they look at the object, they see a red apple instead of a gray one. In this case, your experiences have different impositional elements, i.e., gray and red, and when you act on the basis of this experience, you find out that other subjects perceive the object differently. Despite being impositional and external to your mind, your experience failed to be the source of actions that would be coordinated with other subjects, which makes it illusory. The same analysis applies to hallucinations. Suppose, again, that you see a gray apple on the table, and you tell your friends about it. Now, instead of seeing a red apple, your friends do not see anything on the table. In this case, assuming that you have no control over your hallucination, your experience has impositional elements because you cannot simply stop seeing the apple. However, despite being external to your mind in this particular sense, the object of your experience does not exist, for the impositional elements that individuate it as an object (its shape and color) failed to be the source of actions that would be coordinated with other subjects, thus making it a hallucination.

To sum up, the Resistance View says that the objects of perception are the impositional elements. It also says that those objects exist only if they could serve as basis for coordinated action in the environment. On this account, no reference to the nature of the objects of perception is made to explain how they are individuated. The distinction between the impositional elements that exist and those that do not is explained by the actions they would prompt in subjects. Similarly, the objective and subjective elements of perception are distinguished in terms of the actions that they would promote. Therefore, we do not need to say that existing objects are physical, or that non-existing ones are mental, or even that the

⁸These claims might strike some as wildly implausible. I address some worries related to this point in more detail in section 1.2.2.

objective elements of experiences are caused by the world and that the subjective ones are caused by the mind. Again, all we need is reference to the actions promoted by those elements, which eliminates the need to make ontological assumptions about the nature of the objects of perception. I will develop these ideas in more detail below.

1.2.2 Pragmatic contexts

The discussion above introduced the Resistance View. I will now explain how it helps relationalism with perceptual errors and the plurality of the objects of perception. Perceptual contextualism says that the distinction between veridical and non-veridical experiences and the individuation of the objects of perception are only possible in pragmatic contexts. Bluntly put, an experience is veridical if it would generate pragmatic agreement in pragmatic contexts, and an object of perception is properly individuated if perceptual subjects would pragmatically agree on the nature of those objects in pragmatic contexts. To make this clear, I will now discuss what I mean by “pragmatic context” and “pragmatic agreement”.

A pragmatic context is a situation in which a subject has a perceptual experience p , such that a full description of p is given by:

- (i) how the perceptual subject S of which p is a part is wired;
- (ii) the selective history that describes how S came to be wired in that particular way;
- (iii) the set D of dispositions that S may have in the presence of p by acting in its environment; and
- (iv) how the actions and dispositions to act that S has relates to the actions and dispositions to act that other perceptual subjects of the same kind as S have.

So, for perceptual contextualism, the distinction between veridical and non-veridical perceptual experiences, and the individuation of the objects of perception, are only possible when the relevant pragmatic context characterized by (i)–(iv) is identified.

To clarify, let us consider this in more detail. Perceptual contextualism says that an experience is correctly described by making explicit that you are a human being S who belongs to an evolutionary species E that has been the subject of a natural selective process, which has wired members of that species to be a certain way (i and ii above). Moreover, for the large majority of members of E , the way they are wired makes them disposed to act in a set D of ways when p occurs (iii above), such that they can interact with other members of E , and *pragmatic agreement* among them may be reached (iv above), in the sense that members of E would be able to coordinate their actions when p is the case. “Pragmatic agreement” refers to the situation in which perceptual subjects are successful in coordinating their actions with other perceptual subjects because they have the same or similar experiences of the world. For example, human individuals with normal functioning visual systems who are asked to separate small red balloons from the big white ones would be in pragmatic agreement as long as they could coordinate their actions to realize this task. In terms of the Resistance View, each individual will have experiences with impositional elements and they will engage in actions because of those experiences. These actions allow for the identification of common impositional elements in their experiences (e.g., that everyone sees size in the same way), thus making the coordination of further actions possible to achieve the initial goal.

Similarly, on the apple case discussed above, perceptual contextualism says that the object of the mother’s and child’s experiences is the apple, because it generated pragmatic agreement. In other words, the mother and the child had experiences with the same relevant impositional elements in a pragmatic context, which allowed for the coordination of their actions in relation to the object of their experiences. This is possible because humans share similar perceptual mechanisms (condition i), which is explained by the fact that members of the human species were selected to have those mechanisms (condition ii). Consequently, humans are wired in such a way that allows for a certain set of reactions when they have p , such as being able to identify the apple as a round and red object (condition iii). Finally, by engaging in actions on the basis of p , other individuals will be able to evaluate one’s actions and coordinate their own actions accordingly, thus reaching pragmatic agreement or not.

Note that this account does not require that humans have identical experiences of the world, but only experiences with the same relevant impositional elements in pragmatic contexts. The mother and the child can have different experiences of the apple, as long as those experiences have the relevant impositional elements in the relevant pragmatic contexts. The same holds for other species. In fact, one worry with the present account is that it looks “speciesist”, for pragmatic contexts only provide an account of pragmatic agreement among individuals of the same species. This worry can be resolved by noting that the notion of pragmatic contexts is consistent with the idea that different species may have evolved in a similar way, such that how they respond to the environment can be similar. For example, both humans and pigeons will avoid being hit by a ball that is rapidly coming in their direction. For perceptual contextualism, there is a pragmatic context in which the human and the pigeon could be at least in partial pragmatic agreement, for their experiences had impositional elements that prompted reaction in a similar way. However, humans and pigeons have different perceptual systems — e.g. humans are trichromatic while pigeons are tetrachromatic — which means that there are some pragmatic contexts in which there would not be pragmatic agreement. If, for instance, we could devise an experiment in which pigeons had to somehow identify the color of objects, there would be pragmatic disagreement between humans and pigeons. This highlights an important feature of the notion of pragmatic agreement. Although I have focused on cases of human interaction, pragmatic agreement does not require linguistic capacity or sophisticated cognitive processes, as with the pigeon case. Pragmatic agreement is, therefore, ultimately about the actions subjects would perform on the basis of their experiences.

Perceptual error

How does appealing to pragmatic contexts help to explain cases of perceptual error? Given the importance ascribed to the active role of perceptual subjects in perception, perceptual contextualism claims that the distinction between veridical and non-veridical experiences can only be made in pragmatic contexts. More precisely, an experience is not veridical or non-veridical in itself, but only in relation to pragmatic contexts. In this sense, percep-

tual contextualism resembles Travis's (2004) view, according to which perception cannot establish satisfaction conditions by itself.

To see this, imagine that you are walking in the park and you see an exotic bird sitting on a tree. You try to show it to your partner, but she says that there is no bird on the tree. To solve the disagreement, you ask another person to look at the tree, and this person reports that there is no bird there. Here there is pragmatic disagreement between you and other individuals of your species. According to the Resistance View, your experience has impositional elements that prompted you to act, i.e., to show your partner the exotic bird. Therefore, it has an object. However, if we consider the pragmatic context of your experience, we will notice that it failed to prompt actions that could be coordinated with other individuals of your species. So, the object of your experience failed to be the source of pragmatic agreement in that particular context. Since you are the only person who can see the bird, perceptual contextualism says that your experience is non-veridical because it fails and would fail (if further interaction took place) to be a source of coordinated actions in that particular context. In terms of (i)–(iv) above, you had an experience p in a pragmatic context determined by how you are wired (condition i) and the species to which you belong (condition ii). Furthermore, you acted on the basis of p (condition iii), and those actions were evaluated by other subjects of your species in that context (condition iv). Since those subjects could not coordinate their actions with your actions, your experience failed to generate pragmatic agreement in that context, and is consequently non-veridical.

There are two important things to note about this account. The first is of crucial importance. It refers to the fact that pragmatic agreement need not always be an actual process, in the sense that subjects always need to engage with their experiences to determine their veridicality. Suppose that, instead of trying to show the bird to your partner, you did not say anything. In this case, there is no interaction based on the experience. For perceptual contextualism, your experience is still non-veridical (or hallucinatory) because, if you had acted on the basis of that experience, it would have generated pragmatic disagreement. Thus, what is important is whether there would be pragmatic agreement if enough interaction happens in a pragmatic context, and not so much the actual occurrence of agreement.

Hence, a more informed definition is that an experience is veridical if it *would* generate pragmatic agreement in relevant pragmatic contexts.

Again, this is not a trivial point. Perceptual contextualism purports to be a realist theory of perception, and this is only possible if we analyze veridicality in terms of whether there *would* be pragmatic agreement.⁹ So, the proposal here is not to provide an infallible set of necessary and sufficient conditions to distinguish veridical and non-veridical experiences. Perceptual contextualism allows for the possibility of there being veridical experiences that fail to generate pragmatic agreement and non-veridical experiences that generate pragmatic agreement in some contexts. These cases happen when there is not enough interaction in pragmatic contexts. A more qualified claim would be, therefore, that an experience is veridical if it would generate pragmatic agreement given that enough interaction with the object had taken place.

The main point here is that understanding what it means for an experience to be veridical requires understanding the notion of pragmatic agreement. In this sense, perceptual contextualism adopts a broad Peircean pragmatist stance towards perception. Charles Peirce believed that pragmatism was essentially a method to logically analyze our ideas (see Hookway 2012 and Legg 2014a) and not some metaphysical view of the universe. For example, Peirce said that to fully understand the notion of truth, we need to understand what it means to reach agreement at the end of inquiry. This does not mean, however, that truth is determined by agreement. For Peirce, a proposition would be true if, given enough inquiry about its nature, it would not generate disagreement among inquirers. So, not all situations of agreement are situations in which inquirers reached truth. As Hookway (2012) points out, Peirce is not concerned so much with offering a theory of truth, but rather an account of what it means to say that something is true (see also Legg 2014a, 210). I am making a similar point for veridicality: i.e., understanding this notion requires understanding the notion of pragmatic agreement. But this does not mean that veridicality is determined by pragmatic agreement. I talk about some motivations for this view in more detail in the objections, but for now, we can simply consider the fact that, if we did not interact based

⁹I'm indebted to Cathy Legg for initially pointing this out and for many helpful discussions that helped clarify this idea.

on what we perceive, and if we did not take for granted that some experiences lead to coordinated action and some do not, we could not possibly have formulated the notion of a veridical experience. This suggests that the relationship between perception and action is much more important than most views in contemporary philosophy of perception acknowledge¹⁰, which is why perceptual contextualism and the Resistance View are important.

The second important thing is that, despite being non-veridical, your experience has an object to which it is related. This does not mean, of course, that the object exists concretely (see section 1.2.1). As discussed above, the Resistance View avoids using ontological notions to talk of the objects of perception, so whether an object exists concretely is not relevant to identify the object of an experience. This raises an important issue about the nature of those objects. If non-veridical experiences have objects, and if they are not physical, then they should be mental objects. While this holds for most theories of perception, perceptual contextualism says that this is a false disjunction when it comes to identifying the objects of perception. The idea is that, in studying perception, we start by looking at how the world appears to us, and only later we investigate the nature of the objects of our experiences. This investigation takes place when subjects interact with those objects and observe the actions that can be employed by being related to them. Thus, instead of starting with the assumption that the objects of perception possess properties such as physical and mental, existence and non-existence, perceptual contextualism says that we ascribe those properties to the objects of perception based on the actions that they would prompt. On this view, then, the veridicality of an experience does not have to do with the nature of its objects, but with the actions that would be generated by the experience. This is an important distinctive feature of perceptual contextualism: it is a theory of perception that avoids initial commitment to the nature of the objects of perception. Instead of assuming that perception is about the interaction of two fundamentally distinct domains, i.e., the mind and the world, it starts by looking at perception alone, and only then provides an account of the distinction between mind and world in terms of our interactions with what we perceive.

¹⁰There are substantial works in philosophy of mind and cognitive science that take this idea seriously, e.g., the vast literature on enactivist views of the mind such as Varela et al. (1991); Thompson (2007); Rowlands (2010); Chemero (2011); Hutto and Myin (2012), but these have not received much attention from most philosophers concerned with the questions I addressed so far.

One objection here would be that despite being neutral about the objects of perception, perceptual contextualism entails an extreme metaphysical view, for it will only be plausible if one is already inclined to accept a pragmatist metaphysics. This objection is vague and uninformative. For one thing, there is no such thing as a pragmatist metaphysics, in the sense that different formulations of pragmatism will imply different and sometimes incompatible metaphysical views. So it is not clear what the “extreme” elements of a pragmatist metaphysics are. Moreover, there are no a priori reasons to dismiss pragmatist views of metaphysics as incoherent, so it is not clear why being committed to a pragmatist metaphysics is inherently problematic. I am not denying that the burden of proof is with the perceptual contextualist here, for it is up to him to motivate such views, but only that being associated with more general pragmatist views is not an inherent problem.

But more importantly, while it is true that perceptual contextualism would be initially more attractive to those who accept broad pragmatist views, this is not the central motivation of the view. So it is simply misleading to dismiss it on these grounds. As I argued in section 1.2.1, perceptual contextualism relies on a new view of experience that considers actions as being fundamental to understand the nature of perception. By proposing to define the nature of the objects of perception in terms of the effects generated by our interaction with them, the Resistance View starts with a theory of perception and builds a metaphysical view out of it (e.g., by claiming that the objects of perception are physical or mental in relation to the actions they generate).

This might strike some as an implausible move, but the motivation here is similar to classical pragmatist views. Perceptual contextualism shares with these views the idea that subjects are not mere spectators of the world, but are instead constantly interacting with it. Thus, how we think about the nature of the world cannot be separated from how we interact (and have previously interacted) with it.¹¹ And understanding how we interact with this world is one thing that any theory of perception should be able to explain. The Resistance View denies, therefore, the assumption that a theory of perception should be constrained by a prior metaphysics.

¹¹In this respect, perceptual contextualism share some of the motivations of enactivist views of perception. See e.g. Varela et. al. (1991) and Nöe (2004).

This is a quick response and I do not expect it to convince anyone that pragmatism is right. However, I hope it shows that instead of being a wrong-headed view, perceptual contextualism shares important motivations with an important philosophical tradition, namely, classical pragmatism. And since discussing the plausibility of pragmatist views is beyond my scope in this paper, I take it to be a mistake to dismiss perceptual contextualism merely on these grounds.

To conclude, doing justice to the second point would require discussing how perceptual contextualism, and the pragmatist view that it adopts more generally, understands the nature of physical objects. For my current argument, all that matters is that, if we accept the Resistance View, then it makes sense to say that non-veridical experiences have objects. As I will discuss in section 1.2.3, this is one distinctive feature of perceptual contextualism in comparison to other relationalist views. For now, it suffices to note that perceptual contextualism provides a picture of how a relational view can explain perceptual error in pragmatic terms.

Plurality of the objects of perception

Now I will consider how perceptual contextualism explains the plurality of the objects of perception. Like with perceptual error, the objects of different perceptual modalities are individuated in specific pragmatic contexts. The idea is that different modalities will individuate different pragmatic contexts depending on how pragmatic agreement could be reached in each modality. I will further discuss the case of auditory perception to clarify this point.

Recently, Casey O’Callaghan (2011; 2016) has pointed out that traditional theories of perception focus too narrowly on vision, leading philosophers to think of the objects of other modalities in analogy to the objects of vision. However, O’Callaghan argues that this strategy is misleading, particularly in relation to auditory perception. He defends the view that auditory objects have a temporally extended structure, which makes them structurally different from visual objects, which are not temporally extended. For him, auditory objects should be conceived in light of their mereological structure,

What it [the mereological conception of objects] must allow is that perceptual objects have proper parts that are treated collectively as a single perceptual unity. It must also allow that perceptual objects of different varieties differ in internal structure. For instance, visible material objects have a complex spatially-extended internal structure. Audible sounds or auditory objects have a complex temporally-extended internal structure. (O'Callaghan 2011, 152)

Leaving aside the details of his account, what is important for my purposes is O'Callaghan's suggestion that a proper way to individuate the objects of perception should take into account the structural differences of different modalities. In light of this, my claim is that, by individuating the objects of perception in pragmatic contexts, perceptual contextualism provides a view that follows this suggestion. This becomes clear by looking at the neutrality of the Resistance View in adopting any ontological theory to individuate the objects of perception. Given that having impositional elements is the only requirement for an experience to have objects, perceptual contextualism does not discriminate between objects with different structures outside of pragmatic contexts.

To see this, consider that the way the impositional features of experiences present themselves in different modalities varies, and the way pragmatic agreement would be reached in those modalities varies accordingly. One may need different standards to evaluate the acts and judgments made about a certain object, say a piece of chalk, that could be related to different sensory modalities. If one points to a piece of chalk and says to a colleague "There is some chalk", then it might suffice for one's colleague to only look at the chalk and nod in agreement. However, if one claims that the chalk makes a squeaky noise when one uses it on a chalkboard, then one's colleague will have to use the chalk on the blackboard to see whether this is true.

In the case above, different standards are required to individuate the objects of each experience. This means that different pragmatic contexts would be identified, and that pragmatic agreement would be reached in different ways in different sensory modalities. In particular, time is an essential feature to reach pragmatic agreement in the latter case, for how an experience feels in different temporal moments is essential to determine whether it

has the quality of being squeaky. In the visual case, however, time does not play a fundamental role, for one can usually describe the impositional elements of a visual experience in a single temporal moment. Thus, depending on the modality that we are considering, perceptual experiences will have different impositional elements, which is compatible with the idea that objects might differ structurally across modalities.

By introducing the Resistance View, and the notion of pragmatic contexts, perceptual contextualism makes room for the pluralist view of the objects of perception that is required to deal with cases of non-visual sensory modalities. This is possible because, contrary to most relationalist views, perceptual contextualism adopts a view of experiences that is neutral about the ontological status of the objects of perception. The objects of perception are considered in the context of the actions that they would generate in subjects, i.e., in their pragmatic contexts. Therefore, the neutrality of the objects of perception and the pragmatic account of veridicality allow perceptual contextualism to solve the difficulties raised in section 1.1 in a relationalist framework.

1.2.3 Perceptual contextualism and disjunctivism

One striking feature of perceptual contextualism is that it seems to be inconsistent with disjunctivism, which is prevalent in most contemporary relationalist views. Although there are different forms of disjunctivism (see Byrne and Logue 2008), I will focus on what is often called *metaphysical disjunctivism*, and from now on simply *disjunctivism*. Disjunctivism says that veridical and non-veridical experiences belong to two exclusive metaphysical kinds. Consider the case of Macbeth hallucinating a dagger. For the disjunctivist, either Macbeth sees a dagger *or* it merely seems to Macbeth that he is seeing a dagger. The usual philosophical interpretation of the disjunctive operator has it that either Macbeth has a veridical experience of seeing a dagger or Macbeth has an experience of another kind entirely (see Martin 2004; Snowdon 2008; Fish 2009).

Disjunctivists thus deny that there is a “common kind” factor that veridical and non-veridical experiences share. This means that, even if both veridical and non-veridical experiences seem exactly alike to Macbeth when he reflects on them, it does not follow that they

belong to the same metaphysical kind. For the disjunctivist, phenomenological indistinguishability does not imply metaphysical identity. Now, the question is whether perceptual contextualism can accept this kind of disjunctivist reasoning, and if not, how it purports to be a relationalist theory that denies disjunctivism.

Perceptual contextualism accepts that veridical and non-veridical experiences *have* a common kind factor, and, consequently, it denies disjunctivism. But even without being a disjunctivist theory of perception, perceptual contextualism is a genuine relationalist theory. For it recognizes that we are related to objects both in veridical and non-veridical experiences. This is one important feature that distinguishes perceptual contextualism from current relationalist views in philosophy of perception. In opposition to other relationalist views, perceptual contextualism denies disjunctivism because one of its core claims is that the objects of perception — or the impositional elements of experiences — are fundamental in shaping experiences. But given that even non-veridical experiences have genuine objects, and that veridicality is determined in the pragmatic domain of perception, all experiences are essentially of the same kind. Being veridical (or non-veridical) is a feature that perceptual experiences acquire in pragmatic contexts and not something intrinsic to them. As a consequence, for the perceptual contextualist, veridical and non-veridical experiences belong to the same metaphysical kind.

However, denying disjunctivism and accepting relationalism might have some counter-intuitive consequences. Does perceptual contextualism say that there are no differences between perceiving a physical object, such as my laptop, and an object that does not exist, such as a unicorn? If no, then how should we distinguish between these? Perceptual contextualism says that the objects of our experiences are the impositional elements, so both can be objects of our experiences. However, impositional elements are not physical or mental in themselves, but only in relation to the pragmatic contexts in which they are perceived. Perceptual contextualism assumes that a theory of the objects of perception should be a corollary of a theory of perception more generally, meaning that it does not initially take a standpoint on which objects exist and which do not, or which objects are physical and which are mental. This is not to say, however, that I deny that these distinctions are

meaningful. I do believe that some objects exist and some do not, but instead of taking this distinction for granted, I think that such distinctions are only possible when we have a clear understanding of the nature of perception and how it allows us to interact with the world. For example, we know that most physical concrete objects, as opposed to non-existing or abstract objects (assuming that there are such things), can be manipulated with our bodies because we interacted with those objects before based on our perceptual experiences. In fact, taking these things for granted is part of the reason why the relationship between acting and perceiving is often underemphasized. Perceptual contextualism is an attempt to show why it is important to think about them.

But, as I said in section 1.2.2, my aim here is not to convince one that a theory of perception requires a radical and novel metaphysics, but only to show that some important problems can be reconceived from a pragmatist perspective. Whether the resulting metaphysical revisions would be plausible is the topic of another discussion. For now, I will restrict myself to pointing out that this way of conceiving of the distinction between physical and mental is another idea very familiar to classical pragmatist views, which should suffice to show that it is not simply an arbitrary or wrong-headed way to see the matter.

These considerations might lead one to think that perceptual contextualism is a sort of disguised sense-data theory of perception, for it says that we are related to entities that are not ordinary material objects in perception. This is incorrect for two reasons. The first reason is that perceptual contextualism does not claim that perception is related to some intermediary entity standing between mind and world. Perception is a *direct* relation to objects, but the nature of this relation, and the nature of the objects involved, are ultimately determined in pragmatic contexts. Perception is, therefore, always relational and direct, for nothing stands between mind and world. This leads us to the second reason, i.e., that the veridicality/non-veridicality distinction is not cashed out in terms of whether some intermediary mental entity relates to the external world, but rather in terms of whether perception would generate pragmatic agreement. Again, instead of explaining this distinction by referring to the nature of the objects of perception, perceptual contextualism does so by adopting a pragmatist account.

To conclude, perceptual contextualism does not accept disjunctivism but is not a sense-data theory of perception either. It offers a new way to think of perception, one that takes into account its active and pragmatic dimensions, to develop a new relationalist view that deals with problems relating to perceptual error and the nature of the objects of perception. Moreover, it provides a new perspective to think about relationalism by denying one assumption widely accepted by current relationalist views, namely disjunctivism. Although more needs to be said in order to provide a full-blown pragmatist view of perception, the current paper takes one step forward in providing the basis for such enterprise.

1.3 Objections

I have presented and developed the main aspects of perceptual contextualism. In this section, I will consider three important objections to the view. The first objection says that perceptual contextualism cannot make sense of cases of veridical hallucinations, such as when one perceives one's environment veridically, but one's experience is being caused by something other than the environment. The second objection maintains that perceptual contextualism does not offer a reasonable account of illusions because it fails to acknowledge that well-known cases of illusions, such as the Müller-Lyer illusion, are really illusory. The third objection says that perceptual contextualism tries to introduce representationalism by the back door, since it allegedly fails to acknowledge important insights of relationalism.

1.3.1 Objection 1: Veridical hallucinations

Perceptual contextualism claims that a perceptual experience is veridical if it could generate pragmatic agreement in pragmatic contexts. But if this is the case, one can object that it cannot explain how two subjects can have phenomenologically indistinguishable experiences of some given object that nonetheless have entirely different causes. For example, when you look at a newspaper sitting on the table and have a veridical experience of that object, your partner, who is also looking at the same newspaper, has an experience that is phenomenologically indistinguishable from yours, but that is being caused by some evil

scientists manipulating her brain. In this case, you reach pragmatic agreement about the newspapers, but, the objection goes, there is an obvious sense in which your experience is veridical, while hers is not. Call this a veridical hallucination.

This objection only follows if we adopt a narrow interpretation of what it means to reach pragmatic agreement. If pragmatic agreement refers to actual situations where subjects actually reach agreement about the nature of the objects of their experiences, then it would be true that perceptual contextualism cannot distinguish veridical perceptual experiences from veridical hallucinations. However, as I have argued in section 1.2.2, what determines whether an experience is veridical is if it *would* generate pragmatic agreement in the relevant pragmatic contexts.

What this means in more concrete terms is that, if enough interaction happened in the pragmatic context of a veridical hallucination, we would find out that despite looking exactly like a veridical experience, it is not veridical. It might well be the case that my partner and I never find out that her experience was a hallucination, but this is because there are usually no good reasons to inquire about the nature of this kind of ordinary experience. Suppose that for some reason we decided to further interact with the newspapers to perform various tasks. My partner and I can describe the newspaper in a similar way (e.g., it has such-and-such properties) but if we were asked to perform an action, such as opening the newspapers, or using it to kill a fly, there would be clearly some form of pragmatic disagreement between us. The object causing her experience is different from the object causing mine, which will make our reactions to the environment differ in some important sense. The possibility of such disagreement is what distinguishes a veridical hallucination from a non-veridical one.

One natural reaction to this reply is to say that a veridical hallucination would allow her to perform the same kinds of action that I would do in that pragmatic context. However, this will only hold if we make the further assumption that not only her visual experiences, but all her perceptual experiences, are being caused by something other than the newspapers. Thus, when she moves her arm to reach for the newspapers, something other than the newspapers (e.g. an evil scientist) is providing her with stimuli that will allow her to perform the

right kind of actions in that pragmatic context. We can accept skeptical scenarios like these as legitimate to think about perception, and many philosophers have done so, but perceptual contextualism advises against it. If we accept the idea that veridical experiences and veridical hallucinations could not possibly result in any practical difference in the world, which is required by these cases, then the motivation to conceive of them as two different kinds of mental states seems to lose its appeal.

To see this, consider the related claim that there are two classes of horses, real horses and virtual horses, but our interactions with them could not possibly result in any practical differences. Virtual horses do the same things as real horses: we can ride them, they make the same noises, and if we cut them open, we would find the same organs inside. The only difference between them is that virtual horses are not physical objects; they only behave like physical objects, and the reason is that careful scientists are constantly providing the right causal stimuli to the things interacting with them. Now, if there is no possible way in which our experiences and interactions with virtual horses could ever differ from our interactions with real horses, then it seems misleading to suppose that there are such things as virtual horses. Again, this would only make sense in highly speculative contexts, such as the skeptical scenario above.

Moreover, accepting the possibility of such skeptical scenarios raises problems to all theories of perception, for they cast doubt on the coherence of the distinction between veridical and non-veridical experiences. In this sense, the problem here is not with conceiving of veridicality in terms of pragmatic agreement, but with the very possibility of knowing whether our perceptual experiences ever relate to anything in the world. But this touches on further epistemological problems that are not restricted to perceptual contextualism, which does not make it a special problem for the latter. To sum up, as long as we acknowledge that there could be some practical difference in cases of veridical hallucinations, these offer no problem to perceptual contextualism.

1.3.2 Objection 2: Illusions

This objection says that there are genuine cases of illusions that generate pragmatic agreement. Consider, for example, the well-known case of the Müller-Lyer illusion. The Müller-Lyer illusion consists of two parallel lines of the same length, one of which has an inward “arrow’s head” (two short segments spreading out of each extreme), while the other has an outward “arrow’s head” that consists of two segments converging to each extreme point. Now, when normal human perceivers look at this figure, they see one of the lines (the one with two outward-pointing arrow heads) as being greater in length than the other. This case is problematic for perceptual contextualism because despite knowing that the lines are of different lengths, humans continue to see them as being of the same length. Thus, the objection goes, there are illusions that generate pragmatic agreement.¹²

The same response applies to this case. In illusions, where things appear to be differently than they really are, individuals who do not know about the illusion will eventually be in pragmatic disagreement with those who know it. For example, a person who does not know the Müller-Lyer illusion will say that the lines are of equal length, but this will result in pragmatic disagreement with the report of those who know about the illusion. The latter will say that it appears to them that the lines are of the same length, but they know that they are perceiving lines of different length.

This response might worry some because it seems to be an arbitrary account of illusions. To see this, suppose that almost all human beings were wiped out of the planet, such that only a few who do not know about the Müller-Lyer illusion were left. When these individuals look at the figure, they will pragmatically agree that the lines are of the same length. In this case, the objection goes, perceptual contextualism would have to say that there is no illusion because pragmatic agreement is reached. But this is absurd, for it cannot be the case that an experience is veridical at some time and non-veridical at another time in the same pragmatic context.

The problem with this reaction is that it assumes that pragmatic agreement provides a set of necessary and sufficient conditions to distinguish between veridical and non-veridical

¹²I’m indebted to Kirk Michaelian for pressing me on this point.

experiences. But, as I argued before, this is an incorrect way to interpret my proposal. By saying that veridical experiences are those that generate pragmatic agreement, I am not saying that whenever there is pragmatic agreement a given experience is veridical. The whole point of perceptual contextualism is to emphasize the importance of pragmatic agreement in understanding what it means to say that an experience is veridical. To see why this is important, consider the fact that the possibility of there being pragmatic disagreement is what allows us to distinguish between veridical and non-veridical experiences. In other words, if a given experience, such as the visual experience of the Müller-Lyer lines, had never generated pragmatic disagreement, there would be no reason to call them illusions. On the case above, if those subjects started to investigate how perception works, they would eventually realize that the Müller-Lyer lines are illusory because there would be pragmatic disagreement at some point in their investigation. I am not claiming that the metaphysical status of the experience changed (it was veridical, now it is non-veridical), but only that we were wrong about it in the first case. We cannot rule out the possibility, although unlikely, that in the future we might find out that this is not an illusion for some unknown reasons.

While it is tempting to see this “fallibility” as a weakness of perceptual contextualism, and therefore a reason to reject it, a more positive interpretation is that, instead of relying merely on fixed intuitions about what would count as a veridical experience and what would not, perceptual contextualism offers a theory of perception that makes sense of the possibility of us being wrong about how we classify such cases. This is particularly important for scientific practice, for as Susan Haack (2009) correctly notes, the meaning of scientific terms, such as “veridical” and “illusory”, often change, and it is a theoretical virtue to have a theory that can accommodate these changes.

1.3.3 Objection 3: Representationalism by the back door?

A third objection says that rather than providing a relationalist theory, perceptual contextualism is introducing representationalism by the back door. This objection is motivated on the grounds that, first, the impositional elements of perceptual experiences are elements of experiences and not of the objects themselves. Thus, perceptual contextualism fails to pre-

serve the relationalist intuition that perception is a relation to external objects. Second, one might claim that perceptual contextualism is committed to a view that relationalists deny, namely, that the perceptual relation is determined only after the identification of pragmatic contexts. This claim, the objection goes, is more suited to representationalism, for the way a perceptual system represents the world depends on the context where the representation takes place.¹³

Let us start by considering the first part of the objection. Perceptual contextualism says that the impositional elements are the objects of experience. Thus, the impositional elements are constitutive parts of our experiences. This is not different from the original relationalist claim that worldly objects are constitutive parts of perceptual experiences. In this respect, the only difference between perceptual contextualism and standard relationalist views lies in how to classify the objects of our experiences. While the standard relationalist says that they are ordinary material objects, the perceptual contextualist says that they are impositional elements that are forced on our consciousness. But this disagreement is not a disagreement about the nature of perception (i.e., both standard relationalists and perceptual contextualists believe that it is essentially relational), but instead a disagreement about the nature of the objects of perception. As long as both views are in agreement that perception involves a relation to things external to itself, there is no reason to see the perceptual contextualist commitment to impositional elements as being a form of representationalism.

Since this only answers to the first worry, the second worry still remains open. Does the commitment to pragmatic contexts prevent perceptual contextualism from being a relationalist theory? This objection only holds if we accept that relationalism is necessarily committed to disjunctivism. That is, the reason why standard relationalist views are committed to the idea that the nature of the perceptual relation is determined before the identification of pragmatic contexts is that they assume, as a matter of definition, that a mental state is either a perceptual state (when there is a relation to an external object) or something else entirely. Conversely, perceptual contextualism denies disjunctivism and says that both veridical and non-veridical experiences are genuine cases of perception, and that the distinction between

¹³I'm grateful to one anonymous referee for raising this objection.

veridical and non-veridical experiences is to be given in pragmatic contexts. As a result, perceptual contextualism ends up offering a stronger version of relationalism, that is, one that recognizes that all perception is relational, and that veridicality is to be assessed in pragmatic contexts. This shows us that it is mistaken to think of perceptual contextualism as a form of representationalism, for the reason why standard relationalist views reject that the perceptual relation can be identified only in pragmatic contexts results primarily from their commitment to disjunctivism, and not to relationalism about perception.

Chapter 2

Perception pragmatized

Introduction

This chapter presents my final discussion of perception in the thesis. I develop a theory of perception that reconciles representationalism and relationalism by relying on pragmatist ideas. I call it the *pragmatic view* of perception. I argue that fully reconciling representationalism and relationalism requires, first, providing a theory in which how we perceive the world involves representations; second, preserving the idea that perception is constitutively shaped by its objects; and third, offering a direct realist account of perception. This constitutes what I call the *Hybrid Triad*. I discuss how Charles Peirce's theory of perception can provide a framework for such a view and I devote the rest of the chapter to developing my own pragmatic and Peircean theory of perception. In particular, I argue that considering perception as a continuous temporal process, which essentially involves interaction with the environment, allows us to do justice to the Hybrid Triad. I motivate this view by discussing how a pragmatic theory of perception would deal with issues such as the distinction between veridical and non-veridical experiences and the nature of perceptual objects.

Chapter 2 expands various topics discussed in Chapter 1, but now offering a different take on them. As I pointed out in the introduction to Chapter 1, my study of the works of Charles Peirce made me realize that the best way to characterize his pragmatist theory of perception in relation to contemporary philosophy of perception is to frame it as a hybrid

representational-relational theory. More importantly, though, the hybrid approach offered by Peirce's theory differs in important ways from current hybrid approaches. Current hybrid views of perception have relied too much on 'austere' representationalism, which has made them unappealing to relationalists. Chapter 2 can be viewed as an attempt to counterbalance this tendency. By introducing what I call a "Hybrid Triad", I try to develop a hybrid view that does justice to 'austere' relationalist insights. Such a view requires, however, taking on some commitments that are unfamiliar to, and are often viewed as highly implausible by, analytic philosophers. While this is a worry that is expected to arise when an unfamiliar view is put on the table, the discussion of perception in Chapter 2 can still be insightful for those who are not willing to get on board with Peirce's pragmatism for at least two reasons. The first reason is that the current status of hybrid theories in the contemporary literature is still unclear, and, to my knowledge, no effort to develop a relationalist-inspired hybrid view has been made. In this respect, Chapter 2 offers the first systematic discussion of this topic. The second reason is that, while Peirce's theory of perception might require committing to some unfamiliar views, Chapter 2 sets the stage for further pragmatist approaches to the philosophy of perception. In particular, given the unbalanced focus on austere representationalism in hybrid approaches, the suggestion that pragmatism provides a hybrid route for those who are more sympathetic to austere relationalism can be viewed as a initial motivation for them to develop alternative pragmatist theories of perception that will require less unfamiliar commitments.

Besides this shift to a hybrid approach, Chapter 2 also highlights important developments in my thinking about the notion of pragmatic agreement. Partly due to a more nuanced understanding of Peirce's theory of perception in relation to his semiotics, and partly due to a better grasp of the overall place of Peirce's pragmatism in relation to other pragmatist views, Chapter 2 employs a notion of pragmatic agreement where contexts of "agreement" — the pragmatic contexts — are understood as being *ideal* or *counterfactual* situations, as opposed to *actual* situations, which was the underlying interpretation employed in Chapter 1.¹ This provides an important improvement over the proposal in

¹While I hint sometimes at a counterfactual characterization in Chapter 1, this characterization is not yet wholly integrated into the overall framework of my proposal there.

Chapter 1, as the worries about anti-realism that arose there can now be approached in a more effective way.

In terms of its place in the thesis, Chapter 2 closes Part I, which is dedicated to perception. However, because some of the work done in Chapter 3 still relies on it, it will be important to clarify the connection between these two chapters. One important notion that is introduced in Chapter 2 is the notion of a *pragmatic object*. This notion is introduced and developed in the context of Peirce's *semeiotic*, or his theory of signs. Besides offering a more clear and detailed account of what the objects of perception are in comparison to Chapter 1, the notion of pragmatic objects will play an essential role in Chapter 3, where the questions of the objects of episodic memory and of the objects of episodic counterfactual thought are approached. The overall proposal of Chapter 3 is to provide an account of the objects of different forms of episodic thinking as pragmatic objects. The way this notion is used in Chapter 2 is not always consistent with the way it is used in Chapter 3. This is because they are being applied to different domains of investigation. Despite any possible inconsistencies, the overall idea of pragmatic objects and the motivations behind its introduction are consistent over both chapters.

Outline

This chapter develops a pragmatist theory of perception that reconciles representationalism and relationalism. Section 2.1 surveys representationalism and relationalism and raises some problems for a recent attempt to reconcile them. I also discuss how a reconciliatory theory should look. Section 2.2 discusses Charles Peirce's views on perception and introduces pragmatist ideas to theorize about perception. I argue that acknowledging what I call the *pragmatic dimension of perception* is required to fully reconcile representationalism and relationalism. Finally, Section 2.3 lays the groundwork for a pragmatist theory of perception that successfully reconciles representationalism and relationalism.

2.1 The nature of perception

What is it that determines how we perceive the world? This question has been central in contemporary philosophy of perception, motivating three different influential views. The first is *representationalism*, which takes perception to be essentially a matter of *representing* the external world (see, e.g., Searle 1983; Harman 1990; Dretske 2003). Representationalists say that when I look at a book sitting on a desk, my visual experience consists in a representation of that object and its properties. If the book is red and rectangular, my experience represents the object as being a book, as having the properties of being “red”, “rectangular”, and so on. What explains how I perceive the world is not how the object is, but how my mind represents it to be. In other words, “[t]he properties and situations one is aware of in having an experience... [are the] properties things are represented as having” (Dretske 2003, 71). This means that “[t]he world needn’t contain [the properties] in order to be represented as containing them.” (Dretske 2003, 71), for experiences can represent the world differently than it really is.

The second view is *relationalism*, which holds that an experience is ultimately shaped by its relation to an external object (see, e.g., Martin 2004; Travis 2004; Brewer 2007; Fish 2008, 2009). For the relationalist, the way I experience the world is *directly* shaped by the objects and their properties. The reason why I see a red and rectangular book is that I stand in a perceptual relation with it, meaning that I am directly *acquainted* with the book. The objects of perception “shape the contours of the subject’s conscious experience by actually being the contours of the subject’s conscious experience.” (Fish 2009, 6). That is, the object shapes the experience by being a constitutive part of it (Martin 2004, 64). As Bill Fish correctly notes, representationalists and relationalists can agree that the way experiences present the world to us is the same; what they disagree about is whether the phenomenology of our experiences is ultimately shaped by the subject representing features of the world or by the features themselves (Fish 2009, 13–14).

The third view, which I will call *hybridism*, says that representationalism and relationalism need not be seen as exclusive (see Schellenberg 2011, 2014; Logue 2013; McDowell

2013; Hanna 2015). Susanna Schellenberg (2011; 2014) argues, for example, that perception is partly determined by some internal structures and partly determined by external objects. According to Schellenberg, external objects can be parts of the content of experiences, but they do not determine, or at least not directly, the phenomenology of experiences. What makes my experience of the red rectangular book have the properties “red” and “rectangularity” are the concepts recruited by the capacities I have to identify particular features of the environment (Schellenberg 2014, 210). For example, a red rectangular book recruits my capacity to discriminate things that are books from things that are not books, my capacity to discriminate red and rectangular things, and so on. These capacities and the concepts employed thus determine the phenomenology of my experience. Therefore, perception is representational because it is shaped by the capacities recruited and also relational because objects recruit the capacities.

The motivations for these views vary, but two of them will be important for my purposes. The first has to do with the phenomenological directness of perception. When I see a red rectangular book, it seems to me that I stand in a direct contact with an external object that is red and rectangular. Perception is *direct*, in the sense that I do not perceive redness as being a property of my experience, but as a property of the book in front of me (see, e.g., McDowell 1996; Crane 2006; Hellie 2007; Fish 2009; see Millar 2014 for a critical assessment). The second motivation relates to the nature of illusions and hallucinations. Illusions, on the one hand, consist in cases in which we perceive an object as having a certain property when it does not, e.g., seeing the red book as being green under different lighting conditions. Hallucinations, on the other hand, are experiences about things that do not exist; for instance, seeing a unicorn in my room. For my purposes, I will use the term *non-veridical experiences* to cover both illusions and hallucinations.

Relationalists take the idea that objects shape perception seriously, which allows for an intuitive account of the phenomenology of perception, for it preserves the idea that perception is a direct relation to the world. *Directness* is understood here in relation to the idea that what fundamentally shapes the way we perceive the world are worldly objects, and not in relation to the idea that perceptual experiences are directly or indirectly related

to worldly objects by means of intermediary entities, such as sense-data. So, while most representationalists would indeed say that perception is direct in the sense that there are no intermediary entities connecting our perceptual experiences to the world, they would deny that perception is direct in the sense that it is fundamentally shaped by worldly objects. In this sense, since it seems to me that I have phenomenologically identical experiences in either the presence or in the absence of an object — e.g., when I hallucinate a red and rectangular book — relationalists have to deny that veridical and non-veridical experiences are both instances of perception, for the latter do not relate to objects.² This results in a *disjunctivist* account of this distinction, i.e., either a mental state is a perception, in which case it relates to objects, or it is something else entirely.

Representationalists offer an intuitive account of why non-veridical experiences can appear exactly like veridical experiences. Since my experiences are shaped by how I represent the world to be, I can have an experience that represents a red rectangular book sitting on the table when there is no such book — or when its properties are markedly distinct, such as when I see the book as being green. However, this account is at odds with the idea that our experiences put us in direct contact with the world, for how the subject perceives it is not ultimately shaped by how the world is.

As an alternative, hybridists try to reconcile these two views. Susanna Schellenberg (2011; 2014) argues that veridical and non-veridical experiences can be phenomenologically indistinguishable because, in both cases, the same capacities to identify a given object can be triggered. When I have a veridical experience of the red rectangular book, the book recruits my capacities to identify an object as being a book, as having the color red, and a rectangular shape. These capacities explain, in turn, why I see a red rectangular book. However, I can have phenomenologically identical experiences in the absence of objects, since “[w]e can employ a perceptual capacity even if a relevant particular is not present — where a relevant particular is a particular of the type that the capacity functions to single out”, for in such cases, “one could be prompted to employ the capacities due to nonstandard circumstances, such as unusual brain stimulation or misleading distal input.” (Schellenberg

²Although some relationalists, such as Fish (2009), deny that veridical and non-veridical experiences have the same phenomenology.

2014, 211). This account allegedly reconciles relationalism and representationalism because it preserves the relationship to external objects in veridical experiences and explains how veridical and non-veridical experiences can be phenomenologically alike.

There are two important problems with Schellenberg's (2011; 2014) hybrid view. First, while it provides an account of why veridical experiences involve a relation to their objects, it is not the kind of relation that matters for the relationalist. As Bill Fish points out, when the relationalist says that objects shape our experiences, "the metaphor of 'shaping' is read in a constitutive rather than a merely causal sense." (2009, 6). Therefore, the fact that objects trigger our capacities to identify their properties consists merely in a causal connection holding between objects and experiences. Schellenberg's account fails to incorporate this, for the subject can have that same experience when the object is not present, that is, as long as the relevant capacities are triggered.

The second problem is that if Schellenberg's hybrid account wants to do full justice to relationalism, then it should adopt some form of disjunctivism, for if her hybrid view is both fundamentally relational and representational, then there must be a fundamental difference between veridical and some cases of non-veridical experiences. Since a relation to an external object is at least part of what fundamentally makes an experience what it is, if such a relation is absent at least in cases of hallucinations, some non-veridical experiences will be fundamentally different from veridical ones. But if this is true, then it cannot be the case that perception is fundamentally relational.

One worry here would be that it is not clear why hybrid views should preserve these relationalist intuitions. It is plausible to think of hybrid views as attempts to reconcile *some* important features of representationalism with *some* important features of relationalism, while consistently leaving out some other important features from each theory. In fact, Schellenberg's hybrid view is often taken to be a variant of representationalism that incorporates relational elements by introducing perceptual contents that are object-involving.³ So, it might be argued that not incorporating all the important relationalist features into a

³Representational views of this sort usually define perceptual content as being structured propositions, which are partly determined by objects. See, for example, Speaks (2009); Chalmers (2004); Schellenberg (2010); Siegel (2010) for variants of such views.

hybrid view is not an inherent problem for Schellenberg's view.⁴

In response, I should say that my aim here is not to provide an argument against Schellenberg's view, but rather to point out aspects of her view that might be seen as problematic from a relationalist perspective. In other words, in claiming that her view leaves out important elements of relationalism, I want to suggest that a hybrid view that incorporates these elements could be potentially insightful and also more appealing to those with relationalist inclinations. For this reason, I will simply take it for granted from now on that the underlying motivation behind hybridism is right, namely, that to explain perception, we need to take into account both representational and relational elements. As a result, the following discussion will be an attempt to develop a hybrid view that is more sensitive to the relational aspects of perception.

Before moving on, it will be instructive to discuss what a hybrid view should look like in this context. Preserving the relationalist idea that objects stand in a constitutive relation to perceptual experiences, and the representationalist idea that perception involves representing the world, requires acknowledging that perceptual objects can be at least partly individuated by how the subject represents the world. If the act of representing and how an object is are completely independent of one another, then it seems impossible to have a hybrid view that genuinely reconciles the two views above. However, in acknowledging this, I am not saying that hybridism should accept that perception requires the postulation of intermediate entities, such as sense-data, between the subject and the world (see Russell 1912; Ayer 1956). This would mean giving up on another idea dear to relationalism, namely, that perception puts us in direct contact with the world. I am suggesting, instead, that hybridism should rethink the idea that how the world is is completely independent of how we perceive it.⁵

Therefore, I suggest that a genuine hybrid view should, first, recognize the representationalist insight that perception depends on how subjects represent the world to be; second,

⁴I am grateful to an anonymous referee for raising this point.

⁵Some might worry that this view flirts with subjective idealism. However, the kind of dependence I have in mind is intersubjective dependence. In other words, we should recognize that the world is partly shaped by how multiple perceptual subjects interact with it, as opposed to one single subject (see Zahidi 2014 for related discussion). For more on this, see Section 2.3.2.

preserve the relationalist insight that the objects of perception shape our experiences in a constitutive sense; and third, offer a direct realist view of perception. I call this the *Hybrid Triad*. This paper develops a view that satisfies the Hybrid Triad. I rely on Charles Peirce's theory of perception in order to show how representationalism and relationalism can be combined so as to provide a fully consistent hybrid view. I will not be defending Peirce's view here, but only take it as a starting point to develop my own pragmatic view. In particular, I will argue that Peirce's view that perception is a continuous process of interpretation of what we sensorily apprehend in our interactions with the world is the key to satisfy the Hybrid Triad. I then develop this theory in more detail by discussing how it conceives of the nature of perceptual objects and how it conceives of the distinction between veridical and non-veridical experiences.

2.2 Peirce's theory of perception

In this section, I will discuss the main aspects of Peirce's theory of perception, which introduces an alternative hybrid theory. Peirce's theory of perception follows a general trend in classical pragmatism to overcome dichotomies created by opposing philosophical systems, such as rationalism and empiricism, materialism and idealism, and so on. Peirce wanted his theory of perception to reconcile both direct and indirect realist theories of perception. Although the debate between representationalism and relationalism is not strictly about the directness or indirectness of perception, it presents itself as an updated version of this old quarrel: i.e., instead of asking whether perception is direct, we now ask whether the objects themselves, or representations of them, ultimately shape how we perceive the world. In this respect, Peirce's view can be instructive.

Peirce says that perception is composed of two important and irreducible elements, the *percept* and the *perceptual judgment*. The percept is that which immediately imposes itself on us and contributes something positive to our current stream of consciousness. Suppose, for instance, that you are sitting next to a river thinking about the future. Then, all of a sudden, lightning strikes the river, and you experience a very intense flash invading your

visual field followed by loud thunder. Your stream of conscious experiences is abruptly interrupted by an experience possessing certain features that positively contribute to the contents of your mental states (i.e., the experience has certain qualitative features such as a brief but intense brightness in your visual field). The percept, according to Peirce, “is a forceful thing. Yet it offers no reason, defence, nor excuse for its presence. It does not pretend to any right to be there. It silently forces itself upon me.” (CP 7.621).⁶ This means that, in perception, we seem to be related to something that is external to our own consciousness and that presents itself as being insistent, in the sense that it is forced on our consciousness, but makes no claims of any kind (see Haack 1994; Bergman 2007; Legg 2014b; Wilson 2016). The percept “is absolutely dumb. It acts upon us, it forces itself on us, but it does not address the reason, nor *appeal* to anything for support” (CP 7.622, his emphasis), meaning that it cannot establish truth conditions alone.

The perceptual judgment, on the other hand, consists in the interpretation of the percept in propositional terms, such as when we say, “The lightning was very intense”. In opposition to the percept, the perceptual judgment is abstract, for it prescind from the immediate and particular elements of perception in order to establish a general, coarse-grained connection in propositional form. To use Peirce’s own example, when we see a yellow chair, the percept conveys the exact qualitative features of that object — e.g., that the object’s color is a certain tone of yellow. The perceptual judgment, on the other hand, abstracts from these particularities, focusing instead on a general description of the object, e.g., “x is yellow”.

Like many relationalists in contemporary philosophy of perception (see, e.g., Travis 2004 and Genone 2014), Peirce acknowledges that there is an important sense in which perception is “silent”. The percept is not the kind of thing that is either “true” or “false”; it is just imposed on us and does not depend on anything else to be what it is. The “discursive” or representational aspect of perception that allows me to say that I see a yellow chair is introduced by the perceptual judgment. The perceptual judgment announces its representation of something, and hence it can be said to be either true or false (CP 7.630). While the percept delivers us particular objects which have a particular color, shape, etc.,

⁶For convenience, I refer to Peirce’s work from the *Collected Papers*.

the perceptual judgment abstracts from these particular properties to form propositional judgments.

One important problem is how to understand the relationship between the percept and the perceptual judgment. Two points are of interest here. The first relates to Peirce's claim that the distinction between the percept and the perceptual judgment is analytical (CP 7.626–36), that is, it is intended to help us understand two conceptually distinct elements of perception, rather than to establish a clear-cut division in how perception unfolds in time. This idea is reflected in his introduction of the notion of the *percipuum*, which is the analytical moment in which the percept is interpreted by the perceptual judgment (CP 7.643). The *percipuum* is required because, for Peirce, thought is always a mediation, and placing the percept as an object of analysis implies standing in a mediated relation to it. So, the *percipuum* is not the percept per se, but the first stage at which the perceptual judgment interprets the percept.

The second point is about the nature of the *percipuum*. One natural interpretation would be to think of it as being a single mental state combining both the percept and the perceptual judgment. But, for Peirce, this is not right, for the *percipuum* is an analytical rather than a temporal moment of perception. In the same way that it would be misleading to pinpoint an exact temporal moment when a given species originated in the course of evolution, it is misleading to see the *percipuum* as a single mental state combining the percept and the perceptual judgment in the course of our perceptions (see also Legg 2014b).

To make these points clear, I will appeal to Sandra Rosenthal's (2001) interpretation of the *percipuum* and discuss how these elements work in perception. Rosenthal distinguishes between two uses that Peirce makes of the notion of the *percipuum*. The first use, which she calls *narrow*, is when the perceptual judgment captures the percept in relation to a synthesis of previous percepts. This synthesis is responsible for recruiting certain habits developed by the subject based on his previous experiences, such that he will perceive the world influenced by those habits. Consider the case in which I see a red apple sitting on the table. The perception of this object will be influenced by habits I developed in my previous interactions with apples; for example, I will perceive that object as something that I can eat.

It is important to note that this is not a two-step process: I do not see the apple first and only then form a judgment that it is an object that I can eat it. Rosenthal (2001) notes that, according to Peirce, perception is not a matter of passively receiving some “given” content from the world, but rather a “taken”. Thus, when considered in this narrow sense, the percipuum should not be understood as delivering the subject an “appearance” of the world that is distinct from it, upon which we further form judgments, but rather as presenting the world itself from the perspective of the subject’s past experiences.

The second use that Peirce makes of the notion of the percipuum, which Rosenthal calls *wide*, is when we form a judgment, in propositional form, that interprets the percipuum in its narrow sense. These judgments come in the subject-predicate form, such as in the proposition “The chair is yellow”. The difference between the narrow and the wide percipuum is that the former does not establish any truth conditions, while the latter can be said to be true or false. However, the truth-value of the perceptual judgment is not solely determined by its correspondence to some external reality. Peirce notes that merely thinking about a judgment takes time, however short, which implies that, by the time we are finished thinking the perceptual judgment, the initial percept is no longer present (CP 5.544). The initial percept, as it is interpreted by the perceptual judgment in the wide sense, can only be remembered as interpreted, thus allowing for its truth-value to be evaluated only with respect to future experiences. This means that whether my initial perception of the object will be said to be “true” or “false” will depend on how my future interactions with the world will take place.

To see how this works, take the apple case again. I see an apple sitting on the table and I judge (in the wide sense) that it looks red. By the time I finish making the perceptual judgment, I no longer have access to the initial percept, which makes me unable to determine whether the perceptual judgment is true of the percept. To determine this, I have to engage in future experiences with the apple. For example, I could take it outside my house and look at it in sun light. This would present the apple as being red to me. Or I could ask another person to look at the apple and tell me its color. This would show whether my actual interactions with the object conform with what I remember to be the initial percept. If it does, I might take my actual experience to confirm the past experience, thus making

it true (or veridical); if it does not, then I might take my initial experience to be false (or non-veridical) in relation to my actual experience.

One worry is that this view is simply false in relation to how perception actually works. When one sees a red apple, one does not need to perform any experiment in order to know that one is seeing a red apple. This is right with respect to how we ordinarily deal with objects in our environment, and we do not perform experiments with them because we assume, based on previous experiences, that certain regularities will hold for them. I do not need to experiment with every apple that I see to know that my judgments about apples are usually true. In this respect, it is important to note that the truth-value of a perceptual judgment depends on possible future experiences too, meaning that if we were to perform actions based on that experience, we would learn that it is true or false.

This point can be further clarified by thinking of perception as analogous to the way scientific investigation develops. When a scientist is investigating a certain phenomenon, he interprets what he perceives based on prior knowledge and generates a hypothesis to explain what is being perceived. In order to determine whether his hypothesis is true, he has to interact with the phenomenon (e.g., by conducting experiments), such that he will reformulate certain aspects of his hypothesis, or drop it altogether. Additionally, he has to see whether other scientists can interact with that phenomenon, and arrive at the same (or very similar) interpretation or hypothesis. If a certain agreement is reached about the nature of the phenomenon being investigated, then, and only then, the scientific community will agree on the nature of the phenomenon, and on the most adequate theory of it. Ian Hacking (1983) notes, for example, that interaction with the world often comes before we have a theory of it. He mentions the case of J. J. Thomson and Robert Millikan, who did not know about the existence of electrons — as we understand the term today — when they studied electric charge. They could perceptually interact with what we now call “electrons” through instruments, but it was only when the scientific community learned enough about electrons, that they could be used to intervene in other parts of reality, thus making them “real” (Hacking 1983, 262–263). Now, whenever scientists are concerned with electrons, they do not need to perform experiments to determine their nature, for they take as granted

that the regularities observed in prior interactions with them would hold in the present moment.

The discussion above illustrates aspects of Peirce's theory of perception that will be central to my proposal in the next section. The first is that Peirce sees perception as being a *continuous* and ongoing process in which *time* plays an essential role. One implication of this is that Peirce denies that perception is a succession of distinct mental states, meaning that we cannot clearly distinguish between the role of what is merely given in terms of sensory inputs (the percept), and how we interpret these inputs (the perceptual judgment). For Peirce, these are different features of an ongoing process of interacting with the world. I shall call this continuity of perception over time its *temporal aspect*.

The second aspect, which is a consequence of the temporal aspect, is that perception is not a passive process, meaning that fully understanding how we are perceptually related to the world requires considering, first, the nature of our previous interactions with the world; second, how we are prompted to act based on what we perceive; and third, how our interactions relate to the experiences of other individuals. I call this the *active aspect* of perception.

These two aspects also highlight an important feature of perception that is often under-emphasized by most contemporary theories. This is the idea that perception is determined by our interactions with other individuals in the environment. This is not a new aspect of perception in itself, but rather an important implication of the temporal and active aspects. As with the process of scientific investigation, perceptually interacting with the world also involves evaluating our experiences in terms of how other individuals act, and also in relation to what we are taught by those individuals. For example, instead of reaching for the apple to verify if it is real, I can ask other people whether they see the same apple. Their answer will then influence how I evaluate the wide percipuum formed on the basis of my narrow percipuum. Hence, there is an important sense in which social interactions influence how we perceive the world. In the next section, I develop the temporal and active aspects in more detail and show how they can help us theorize about perception; in particular, how they can satisfy the Hybrid Triad.

2.3 The pragmatic view

I will now discuss how a theory of perception inspired by Peirce's theory can satisfy the Hybrid Triad. I will call it the *pragmatic view* of perception. The pragmatic view takes as a starting point the temporal and active aspects of perception. I proceed by discussing the representational and relational aspects of the pragmatic view, and apply it to the issue of the veridicality of perception. This discussion will allow us to see how the pragmatic view can satisfy the Hybrid Triad.

2.3.1 The representational aspect of perception

I have argued that Peirce's theory of perception takes the temporal and active aspects of perception seriously. I do not mean to suggest that other views in philosophy of perception completely ignore this. My aim in this section is rather to show how we can deal with important contemporary problems when we think about perception explicitly in these terms. That said, when I claim that perception is continuous in time, I mean two things: first, that it is not possible to draw a clear separation between perception and other mental states, such as beliefs, and second, that perceiving is an ongoing active process and not merely a succession of discrete mental states.

Note that the first claim is *not* that perception and beliefs are the same thing. The impossibility of drawing a clear separation between them relates to how perception unfolds in our continuous interaction with the environment. So, for example, when I see a yellow chair in my office, it is plausible to say that I have a visual experience of a yellow chair and that I formed the belief "There is a yellow chair in my office". For the pragmatic view, the perceptual experience and the belief are continuous in the sense that it is misleading to see this case as involving a two-step process where we first perceive and only then form beliefs on the basis of perception. In other words, seeing a yellow chair in my office partly consists in believing that there is a yellow chair in my office. These two things (perceiving and believing) are indeed different, but this difference is only clearly visualized in analytical contexts, meaning that there is no clear way to separate them we we look at our actual interactions

with the world. Moreover, with respect to the second claim, when I say that perception is active, I mean that perception is not simply a matter of a subject being sensorily stimulated by the world. Being perceptually related to the world also requires interacting with it based on sensory stimuli.

When we consider these two aspects together, it is possible to see how the pragmatic view makes sense of the “representational aspect” of perception. The fact that perception develops in time and requires the interaction of subjects with their environment allows subjects to change, based on this ongoing interaction, their interpretation of what is being perceived over time. Imagine that I see an apple on the kitchen table and decide to have a bite. However, after the first bite, I notice that it is a fake wax apple. I then reconsider my initial experience and realize that it was misleading. That is, I had a visual experience of an object (the percept) and my initial perception was influenced by my previous interactions with apples (the percept plus the perceptual judgment yielding a percipuum in the narrow sense). I thought that the apple looked tasty and that I should have a bite (the percept plus the perceptual judgment yielding the percipuum in the wide sense). By the time I formed this perceptual judgment, the initial percept was gone, so what determines whether my experience is true is how my future engagement with the object would line up with the how other subjects would engage with the object in that situation.⁷ Finally, I decided to have a bite and discovered that my initial perception was misleading.⁸

Although this is one way in which I can revise my representations of the world over

⁷Here we stumble on an important notion in Peirce’s philosophy, that of *truth*. For Peirce, a proposition would be true if, given enough inquiry about its nature (what he usually called “the end of inquiry”), it would generate agreement among inquirers. However, Peirce is not concerned with offering a theory of truth, in the sense of providing necessary and sufficient conditions for truth, but rather an analysis of what it means to say that something is true (see Misak 2004; Hookway 2012; Legg 2014a). As Cathy Legg (2014a) points out: “[the] ‘end of inquiry’ [...] is not ‘end’ in the sense of ‘finish’. It is ‘end’ in the teleological sense of ‘aim’ or ‘goal’. Rather than a description of some future time where all questions are settled, Peirce’s explication of truth is an idealised continuation of what scientists are doing now, namely settling questions about which they genuinely doubt” (206). What I have in mind here is something similar, but now applied to perception: that is, a perceptual experience is said to be true or veridical if it generates agreement among relevant inquirers in relevant situations when their perspectives are taken into account. I develop this claim in more detail in section 2.3.3.

⁸One worry with this example is that it assumes that perception can represent *high-level properties* (see Peacocke 1992; Siegel 2006; Nanay 2011), such as kind properties, as opposed to *low-level properties* (see Tye 1997; Dretske 1997; Clark 2000), such as color and shape. While this is a controversial issue, it is only tangential to the idea being conveyed here, as the same points could be made in relation to low-level properties, e.g., in cases where we see objects as instantiating different colors because of different illumination settings.

time, it is not the only one. In some cases, these revisions require interacting not only with inanimate objects, such as apples, but with other perceptual subjects in the environment. Suppose that I arrive one day at my office and see a pink chair in the corner. I do not own a pink chair, so I simply assume that someone left it there for some reason. However, since I feel curious, I ask my colleagues whether they know why the chair is there. To my surprise, I find out that they cannot see the pink chair and they say that the only chair in the room is my usual black chair. In this case, I have a visual experience of a pink chair and I perceive that chair as being there based on my prior experiences — i.e., chairs are objects that I usually see in offices. But when I asked my colleagues if they knew anything about it, I found out that I was hallucinating the chair, for none of my colleagues could see it.

One worry here is that this will only work if we assume that perception is *cognitively penetrable*, which is a highly controversial issue (see Siegel 2006; Stokes 2012, 2013). Roughly, as Stokes puts it, the idea is that “[...] perception is, sometimes, penetrated by cognition.” (2013, 647). Note, however, that the example above does not require assuming that perception is cognitively penetrable. The “revisions” are not revisions in the phenomenology of my experience, but rather in the way I evaluate my experience. After talking to my colleagues and learning that there is no pink chair, I might still continue to see the chair for some reason or another. The only difference now is that I know that the experience is not veridical, which I could not tell before by merely introspecting.

These examples show that, while it is true that we form representations of the world when we perceive, those representations are constantly revised as we interact with the world and other individuals. For the pragmatic view, perception is representational in the sense that each individual perceives the world from the perspective of his previous experiences. These representations prompt subjects to act, thus creating opportunities to reformulate them. If we understand the representational aspect of perception along these lines, then the first part of the Hybrid Triad is satisfied.

2.3.2 The relational aspect of perception

The pragmatic view is also fundamentally relational, which is similarly possible because of its temporal and active aspects. Relationalists say that the objects of perception shape experiences in a constitutive sense, meaning that the properties of objects determine how we see them. Yet, it is hard to understand how this constitutivity can be reconciled with the idea that perception is representational.

It is important to distinguish here between relationalism, which is a view about perception, and metaphysical naïve realism, which is a view about the objects of perception. Relationalism says that perception is constitutively shaped by its objects, but does not say anything about the nature of those objects. Most relationalists, however, accept metaphysical naïve realism, which says that the objects of perception are ordinary material objects, such as tables and stones, which are completely mind-independent (Brewer 2007, 88). Once we associate these two views, relationalism is taken to imply that perception is constitutively shaped by ordinary material objects. However, these two views are independent: it is possible to be a relationalist without being a metaphysical naïve realist about perceptual objects.

That being said, by saying that the pragmatic view has a fundamentally relational aspect, I endorse the first view above, but reject the second. In other words, I believe that perception is ultimately a relation to what I will call *pragmatic objects*, which include what we ordinarily call material objects, but is not restricted to them. Broadly understood, pragmatic objects are the objects of the mind that promote action. I understand action in two senses here: first, there is physical action, such as moving my arm to reach for an apple; and second, there is mental action, such as thinking that an apple looks tasty. To motivate this view of perceptual objects, I shall introduce what I call the phenomenology-first approach to perception and then connect it to the temporal and active aspects.

The phenomenology-first approach consists in looking at the qualitative elements of our experiences without assuming any theory about their nature. When I see a red apple sitting on a table, the phenomenology-first approach says that I see a thing which has certain

qualities grouped together, such as color and shape. Moreover, I see that object as being distinct from myself, in that its existence does not seem to depend on its relation to my mind. More specifically, what I mean here is that this object seems to resist any of my attempts to change it — e.g., if I wish that a banana could be red, this wish cannot change the qualities of the perceptual object. This is all the phenomenology-first approach extracts from perception: it recognizes that, when I perceive something, it seems to me that I see some qualities united together, which I call a “thing”, and that thing seems to be something external to my own consciousness. At this level, I refrain from judging whether the thing is physical or mental. The objects that we perceive in this minimalist phenomenological sense and that make subjects to act in the environment, irrespective of whether they are real, are what I call *pragmatic objects*.

To further clarify this point, pragmatic objects can be compared to a more familiar notion, the idea of *intentional objects*. Broadly understood, intentional objects are the objects that our mental states are directed at (Brentano 1973; Crane 2001, 2013). Pragmatic objects are similar to intentional objects in this respect. However, while Brentano (1973) conceived of intentional objects as being distinctively mental, pragmatic objects are defined in neutral ontological terms. In other words, they are objects in the sense that they are the kind of things that minds can be directed at, but facts related to their existence (or inexistence) are to be determined in terms of the actions that they promote in subjects. That is why the phenomenology-first approach is important, namely, it requires us to put aside our knowledge of the world acquired through previous interaction with it to properly identify pragmatic objects. In this sense, the phenomenology-first approach resembles the phenomenological reduction proposed by Husserl (1982; 1988), in which we are required to suspend our “natural attitude” to investigate what is presented to the mind.⁹

Before we proceed, it is important to highlight an assumption being made here in relation to the neutral ontological aspect of pragmatic objects. That is the view that how we ultimately understand metaphysical notions, such as existence and non-existence, physical

⁹While such methodological assumptions might seem controversial for some analytic philosophers, there has been a growing interest in how the phenomenological reduction can be a useful tool for philosophy of mind and cognitive science more generally (see Gallagher and Zahavi 2013, ch. 2 for a more detailed discussion).

and mental, and so on, are best characterized in pragmatist terms. The basic idea here refers to what Peirce called the pragmatic maxim in his well-known essay *How to Make Our Ideas Clear* (CP 5.388). Peirce conceived of the pragmatic maxim as a logical principle that analyzes the meaning of concepts in terms of the conceivable practical effects the concepts have or would have have in our experience.¹⁰ For example, Peirce says that a thing is hard if it would not be scratched by many other things. The concept HARD is therefore defined in terms of the effects it would have in our experience if we were to interact with hard things. The same is assumed here of more general metaphysical notions, which is what allows us to say that a pragmatic object exists only in relation to the actions that it promotes in subjects. Instead of viewing opposing notions as *discontinuous*, the pragmatist opts for a *continuous* characterization by appealing to actions (see Menary 2009; Solymosi 2013 for related discussions).

This is a bold metaphysical assumption to make and a full pragmatic account of perception has to motivate it properly if it is to succeed. However, my scope here is more modest. I do not want to provide a full-blown argument for a pragmatist view of perception, but simply show that understanding perception in pragmatist terms can be useful to deal with a particular issue in philosophy of perception. I will not, consequently, attempt to defend any pragmatist metaphysics here. By making this assumption explicit, however, I hope that it becomes clear that, instead of being wrong-headed, the attempt to bring pragmatist elements to theorizing about perception draws at least part of its motivation from an independent and well-established philosophical tradition, namely, classical pragmatism.

With these clarifications in mind, let us return to the discussion of pragmatic objects. The phenomenology-first approach, as discussed above, does not allow us to know much about pragmatic objects and therefore to take any position on whether they are physical or mental, or even real (or existent). We only know that we can possibly interact with them. Throughout the course of our lives, we learn that some pragmatic objects can be

¹⁰This is not, however, Peirce's initial formulation in *How to Make Our Ideas Clear*. As he later recognized, his first formulation of the maxim was too nominalist, in that it did not take into account how things would effect our experiences in possible but not actual experiences. Later in his works, Peirce provided a new formulation of the maxim in the *Harvard Lectures on Pragmatism* (CP 5.14), which is the one described here. Since I do not intend to provide a complete picture of Peirce's philosophy, I will leave these exegetical subtleties aside.

used in a certain way, and that certain regularities hold in our interactions with them. More importantly, I know that the regularities hold not only for me, but for other people who interact with those objects. I can talk to my partner about the apple on the table and we can discuss its properties. If my partner is hungry and decides to eat the apple, her hunger will be satisfied. The point is that we do not need to make metaphysical assumptions about the nature of those objects in order to make sense of their role in perception. All we need to do is to pay attention to their qualities, the effects generated by our interactions with those objects, and how they allow us to get along with other individuals who can also interact with them.

This discussion highlights that, on the phenomenology-first approach, we perceive different pragmatic objects, and our interactions with them reveal different regularities. Knowing these regularities allows us to interact with individuals who can also perceive those pragmatic objects. Now, some of those objects will not be useful to engage in further actions in my environment. For example, if I see a car parked outside but my friends do not, and I tell them that we can use it to go the beach, that will obviously not be possible. In this case, my perception has a pragmatic object, but fails to allow me to engage in further interactions with the world and other individuals. Thus, a pragmatic object is said to be “real” or “not real” depending on the paths of action that I undertake based on my perception of it, and on whether those actions line up with the actions that other individuals undertake in the same context.

One might worry that this fails to be a realist theory of perception, in the sense that it does not say anything about whether the things we perceive are really out there, or mere figments of our minds. I will address part of this worry later when I talk about the veridicality of perception, but the pragmatic view is certainly not an anti-realist theory of perception. The pragmatic view takes a different starting point than most contemporary theories of perception, i.e., it does not assume beforehand that perception is a relation between two fundamentally distinct entities, namely, a subjective mind and an objective world. The phenomenology-first approach urges us to start considering perception blinded from ontological theories. Thus, it says that our notion of reality, understood as something that

exists independent of any perceivers, is a late product of our perceptual interactions with the world. It is only after we interact with the world and learn which pragmatic objects can serve as successful guides for action in the environment for different individuals that we form this notion. An object is said to be “real” when more than one individual can interact with it, such that they can undertake successful paths of action in the environment. This view can deliver us a notion of “reality” understood as mind-independence without having to talk about the nature of the objects in question.

I do not expect to make a complete case for this notion of “reality” here, but only show that the pragmatic view of perception does not ignore these worries. My claim is, instead, that *all* perceptual experiences involve a relation to pragmatic objects. Moreover, I want to say that pragmatic objects constitutively shape the way we perceive the world. This preserves the relationalist insight that objects shape experiences and avoids disjunctivism, because all experiences relate to objects. And once disjunctivism is discarded, we can understand the relational side of non-veridical experiences without having to say that they are not instances of perception.

To illustrate this, consider the following example. Imagine that you are at the beach with some friends, and you spot something floating on the sea.¹¹ At first glance, it looks like a log to you. Your friends can also see this object, but it looks different to them. John thinks that it is a broken piece of wood from a boat, and Peter sees it as seaweed. Given that the object is far away, it is hard to tell with precision what it really is. However, in this particular context, that black dot floating on the sea stimulated your visual systems, and each of you “created” an interpretation of what is being perceived based on your previous experiences.

This case illustrates a situation where different individuals are sensorily stimulated by the same external thing, but they end up seeing different things. It is natural for us to ask what these individuals are perceiving. Since we want a hybrid theory that preserves constitutivity, we cannot simply say that they perceive the same thing, for otherwise they would not have different experiences. How would the pragmatic view conceive of this case?

¹¹This example has been adapted from Hausman’s (2006) discussion of the relationship between perception and semeiotics in Peirce.

The first thing to note is that you, John, and Peter have developed different habits during your past interactions with the world. So, when you see that thing on the beach, each of you will interpret it according to your previous experiences. This is Peirce's percipuum in the narrow sense. At the initial moment, you are related to different pragmatic objects: a log, a piece of wood, and seaweed, respectively. You then make judgments about the nature of this object, e.g., "The black dot on the sea is a log", which is Peirce's percipuum in the wide sense. Note that, despite these interpretive elements influencing how you perceive the black dot, they influence perception only to the extent that they determine that different pragmatic objects will be picked up because of them. On the phenomenology-first approach, the piece of wood, the seaweed and the log are all perceived as objects out there in the environment.

The pragmatic view says that perception is temporal and active, so we should take into account the future interpretations we make based on what we sensorily perceive. Suppose that, as the object gets closer, you learn that it is a log. It turns out that the pragmatic object of your initial perception is the object that persisted, for it allowed you, John, and Peter to make reliable inferences about the world. Throughout this process, however, you have been related to different pragmatic objects which have shaped the way you see the world. I am not suggesting that the initial pragmatic objects are intermediary entities separating you from the final pragmatic object, but that during the perceptual process, you relate to different pragmatic objects, and your interactions with the world eventually lead you to relate to the "right" pragmatic object. This is possible, initially, because your habits make you pick different initial pragmatic objects, but as you interact with the world, you can correct your interpretations, such that you end up picking the "right" pragmatic object.

When I say that the initial pragmatic objects are different from the final one, I mean that they are qualitatively different, but not numerically different. In the same way that the same object can appear with different colors to a daltonic individual and an individual with a normally functioning visual system, the same external object can appear as a different object for me and someone else — i.e., a log as being seaweed. Instead of saying that individuals represent the same object differently, the pragmatic view says that the same object appears differently to individuals because of their prior habits, which can be corrected by their

future interactions with it. Therefore, the final pragmatic object is said to be the “real” thing residing in the world, but this is not because it has a physical or material nature, but rather because it is the object that allows subjects to interact successfully in the environment.

This discussion allows us to see how the constitutivity of relationalism can be incorporated by the pragmatic view. By relying on the idea that subjects perceive different pragmatic objects because of their past experiences (Peirce’s percipuum in the narrow sense), we can make sense of how our representations can influence the way we perceive the world without giving up on the idea that objects constitutively shape our experiences. However, since perception also involves making judgments (Peirce’s percipuum in the wide sense), it can only be fully understood in relation to future actions undertaken because of the initial experience. Thus, when subjects interact with the environment, they revise their initial interpretations, such that they can eventually pick up the “right” pragmatic object. The temporal and active aspects of perception are essential in this process, for it is only by acknowledging them that we can understand how subjects can relate to different objects in the course of perception without saying that those objects stand between mind and world. This satisfies the second part of the Hybrid Triad, that is, that a genuine hybrid view should preserve the constitutivity idea.

2.3.3 The veridicality of perception

Can the pragmatic view provide a direct realist view of perception? This is the third part of the Hybrid Triad and bears on the distinction between veridical and non-veridical experiences. For my purposes here, I take it that veridical experiences are those that deliver the subject correct information about the external world. To recapitulate the example from the beginning, when I have a visual experience of a red rectangular book and there is a red rectangular book in front of me, my experience is veridical. In contrast, non-veridical experiences provide partially or completely misleading information about the external world to a subject. These are usually classified as either illusions or hallucinations respectively. I have an illusory experience when I perceive some feature or property of the external world misleadingly, and I have a hallucinatory experience when I have experiences of things that

do not exist in the external world.

The dispute between representationalism and relationalism depends, in part, on how one conceives of the veridicality of perception. One advantage of representationalism is that it provides a unified account of veridical and non-veridical experiences, which differ only with respect to the accuracy of a given perceptual experience. Relationalists take a difference position on this issue, stating that perception is constitutively shaped by its objects. However, non-veridical experiences can be indistinguishable from veridical experiences, in the sense that the subject cannot tell, from his own point of view, whether it really involves an object or not. So, the problem is to explain the possibility of there being similar experiences in the absence of objects. To deal with these cases, some relationalists have argued that the subject not being able to tell veridical from non-veridical experiences does not imply that there is no difference between them. In fact, some relationalists, such as Martin (2004), go so far as to claim that the only similarity between veridical and non-veridical experience is the fact that they are indistinguishable. This results in a bold metaphysical claim, often called *metaphysical disjunctivism*, which claims that either an experience is a perception, in which case it is constitutively shaped by an object, or it is a mental state of a different kind.

How does the pragmatic view conceive of the distinction between veridical and non-veridical experiences? I have suggested that a satisfying hybrid view should preserve the constitutivity of relationalism for non-veridical experiences, which casts doubt on the plausibility of disjunctivism. My proposal is that we can avoid disjunctivism and do justice to the constitutivity idea by claiming that perceptions, veridical and non-veridical alike, are related to pragmatic objects. Moreover, the distinction between veridical and non-veridical experiences depends on whether subjects can take *successful* paths of action based on such experiences. I shall say that an experience is successful when the subject can pursue paths of action that allows him to coordinate his actions with both the environment and the actions of other subjects.

To motivate this approach, let us reconsider the case of the black dot on the sea. Each of you see different pragmatic objects: a log, a piece of wood, and seaweed respectively. You

also learn, later on, that the object is a log. At this initial stage, who is perceiving veridically? One natural answer is that you are the one who has a veridical experience, while John and Peter have non-veridical experiences. The reason is that your experience captures how the world is, while their experiences do not. However, if perception is relational, this cannot be the case, for you, John, and Peter are perceiving the same thing.

The pragmatic view holds on to the initial intuition that you are perceiving veridically, while John and Peter are not. However, it says that, considered only in relation to their objects, there is no absolute question of whether an experience is veridical. An experience is only said to be veridical or non-veridical in relation to what I call *pragmatic contexts*, which are broadly understood as the practical situations in which the actions we undertake on the basis of our experiences relate to the actions that other subjects undertake on the basis of their experiences. In more detail, I understand pragmatic contexts as being *actual* or *possible*¹² situations in which perception could occur, in which (i) the physical constitution, (ii) the biological selective history, (iii) the way a perceptual subject interacts with the world, and (iv) how it interacts with other perceptual subjects, mutually co-determine whether an experience is veridical, or not.

Let us explore this in more detail. On the example above, your experience is said to be veridical because it is the only one that leads to successful inferences about how the world is in the relevant pragmatic context. In other words, your experience is the only one that allows you to successfully interact with other subjects in the environment. To see this, take the case of Peter. Peter's initial experience is of seaweed. However, as you discuss the nature of that thing on the sea, you formulate new interpretations of what you see, such that Peter eventually agrees with you when the object is close enough for him to identify it as a log. This is only possible because you and Peter are human beings, and therefore you have a similar physical constitution and you share the same evolutionary history (i and ii above). Moreover, due to these similarities, you can interact with the world and evaluate your actions on the basis of your interactions with each other, such that you can

¹²This is where realism enters into the theory. Veridicality does not depend only on actual contexts of interaction, but also on possible but not necessarily actual contexts. Therefore, we do not need to say that it depends on the existence of individual subjects.

coordinate your actions in the environment accordingly (iii and iv above). Thus, if subjects can agree on what the pragmatic object of their experiences is after engaging in actions in the world, then the experience is veridical. Otherwise put, an experience is veridical if it would generate *pragmatic agreement* in relevant pragmatic contexts.

Two important clarifications are required. First, although the notion of agreement seems to imply some kind of verbal communication, I understand it more broadly here. I can reach pragmatic agreement with other subjects even when there is no verbal communication of any sort. When I pass the basketball to another player in my team, we are in pragmatic agreement about what we perceive, for we can coordinate our actions based on how we perceive the ball. Second, it is important to note that pragmatic agreement depends on the *actuality* and *possibility* of subjects interacting in pragmatic contexts. The example above seems to suggest that an experience will be veridical or non-veridical only if subjects actually interact on the basis of it. However, this would imply that if no interaction takes place, there is no answer to the question of which experience is veridical. I want to say that the veridicality of an experience depends not only on subjects actually engaging in actions on the basis of the experience, but also on possible and non-actual paths of action that would have been the case *if* subjects had engaged with the experience. In other words, an experience is veridical not only if it generates actual pragmatic agreement in actual pragmatic contexts, but also if it could have generated pragmatic agreement if the subjects had further engaged with the experience. Thus, even if after looking at the object on the sea you, Peter, John had left the beach, your experience would still be veridical and theirs would not be, for if you had engaged with them, your experience would be the one to have generated pragmatic agreement.

To further motivate this approach, let us now consider a hallucination, the pink chair that I see in my office. The pragmatic view says that the pink chair is the pragmatic object of my initial experience and it constitutively shapes it. However, when I talk about that chair to my colleagues, I learn that they cannot see it, and I convince myself that despite appearances, there is no pink chair. In this case, I have a hallucinatory experience because my experience failed to be the source of useful inferences about my environment in relevant

pragmatic contexts, such that I could not coordinate my actions with the actions of other subjects. That is, individuals who have the same physical constitution and who share the same evolutionary history with me (i and ii) could not engage in the same actions with the environment in that context (iii and iv). Therefore, my experience of the pink chair failed to generate pragmatic agreement, and thus is non-veridical.

It might be argued that this account does not explain cases in which I am perceiving veridically, but my friends pragmatically disagree with me. For example, it might be the case that there is a pink chair in the office, but for some reason on another, my colleagues fail to see it. The pragmatic view seems to imply that my experience is still hallucinatory in this case. Here it is important to remember that a full account of veridicality depends not only on the actual, but on the possible interactions that could happen in a pragmatic context. Thus, while in most situations it is better to rely on actual pragmatic agreement to assess the veridicality of an experience, this does not mean that this agreement is absolute or infallible. The idea here is that my experience counts as veridical in this case because, if other people were to enter the office, they would see the chair and pragmatically agree with me.

One worry likely to arise in relation to this account is that it does not distinguish between veridical and illusory experiences. A number of illusions do generate pragmatic agreement, but we still take them to be non-veridical. For example, in the Müller-Lyer case, where we see two lines as being of different lengths when they are not, due to the arrows attached to their extremities, we might pragmatically agree that the lines are of different lengths. Therefore, the pragmatic view has to say that this experience is veridical. In response, it is worth noting that we know this experience to be illusory because we further interacted with it, for instance by finding out that if we remove the arrows at the extremities, we can see that the lines are of the same length. Thus, the interaction with the lines allows us to see that the initial state of pragmatic agreement was wrong. It is only then that we pragmatically agree that the lines are of the same length and that the initial experience is illusory. This means that the pragmatic view is ultimately fallibilist, in the sense that absolute pragmatic agreement about the status of a given experience is impossible. Although this might seem

problematic at first glance, this is what actually makes the idea of illusions and hallucinations possible. If we want to be direct realists about perception, then we have to take our interaction with the world into account, for this is the only way to know that what we sensorily perceive is actually real. This does not give a complete answer to the problem of illusion, but I hope it is enough to show that the pragmatic view does not ignore it.

I do not expect this to be a full account of the veridical and non-veridical distinction, but rather an outline of how we can think of it in pragmatic terms. However, one important worry relates to perception in the animal realm. The definition of pragmatic context seems to be restricted to individuals belonging to the same species, thus ignoring the fact that other animals are perceptually related to the same world as humans. The pragmatic view does not preclude humans from pragmatically agreeing with non-human animals. The difference in these cases is that the relevant pragmatic contexts will differ across different species. Suppose that I am standing next to a pigeon and a crazy person throws a brick at me. Both the pigeon and I see the brick approaching and dodge it, for we see it as a potential source of harm. According to the definition above, I am in pragmatic agreement with the pigeon, for we perceive the object in similar relevant ways, which allows us to engage in further actions to achieve an end, i.e., avoiding being hit. However, pigeons are tetrachromatic animals, as opposed to humans, who are trichromatic, which makes us see colors differently. If we could devise an experiment in which pigeons were required to identify the color of that object, there would be likely pragmatic disagreement between me and the pigeon. So, instead of ignoring animal perception, the pragmatic view provides a useful framework to understand how different species perceive the world differently. On this view, we can avoid potential problems relating to questions such as which species perceives the world correctly. We can opt, instead, for a notion of veridicality that is context-sensitive, thus making room for the idea that different species can veridically perceive the same world (see Chirimuuta 2017).

2.4 Conclusion

To conclude, if this view of veridicality is plausible, the pragmatic view can finally offer a theory of perception that is direct, i.e., perception is directly shaped or constituted by its objects (the pragmatic objects), that moves beyond disjunctivism. Since perception is temporal and active, we can say that even non-veridical experiences have objects, and are constitutively shaped by them, without saying that they are fundamentally different from veridical experiences. This reconciles representationalism and relationalism on the issue of veridicality, for the constitutivity idea and the view that veridical and non-veridical experiences belong to the same metaphysical kind are preserved. We are thus left with a theory that satisfies the Hybrid Triad, that is, it provides a satisfying reconciliation of representationalism and relationalism about perception. Although this paper has not developed a full pragmatic account of perception, it has laid the groundwork for the development of such theory. In particular, it has established the basic framework to think of important problems in contemporary philosophy of perception, such as the nature of veridical and non-veridical experiences, that moves beyond current theories in the field.

Part II

Memory

Chapter 3

Thinking about events

Introduction

This chapter is my first attempt (in collaboration with Kourken Michaelian) to deal with the questions mentioned at the Introduction in philosophy of memory. It starts by considering the debate over the objects of episodic memory, which has for some time been stalled, with few alternatives to familiar forms of direct and indirect realism being advanced. The chapter thus moves the debate forward by building on insights from the recent psychological literature on memory as a form of episodic hypothetical thought (or mental time travel) and the recent philosophical literature on relationalist and representationalist approaches to perception. The former suggests that an adequate account of the objects of episodic memory will have to be a special case of an account of the objects of episodic hypothetical thought more generally. The latter suggests that an adequate account of the objects of episodic hypothetical thought will have to combine features of direct realism and representationalism. We develop a novel pragmatist-inspired account of the objects of episodic hypothetical thought that has the requisite features.

Chapter 3 inaugurates Part II of the thesis, which is dedicated to memory. In chronological terms, Chapter 3 was written concomitantly to parts of Chapter 1 and Chapter 2. Because the chapter is a result of a collaboration with Kourken Michaelian, it brings two bodies of literature together to raise a new question in the philosophical landscape. The

question of the objects of episodic hypothetical thought, as we call it, is an updated version of the question of the objects of memory in the context of recent empirical work done in psychology. The proposal of the chapter is to motivate this question and provide a positive solution to it. The positive solution that we offer is what we call the pragmatic approach to, or the pragmatic view of, the objects of episodic hypothetical thought.

The pragmatic approach to episodic hypothetical thought is based on the pragmatic view of perception developed in Chapter 2. The notion of pragmatic objects reappears here, but now applied to single out the events that serve as objects for episodic hypothetical thoughts. The phenomenology first approach, which is an important element in Chapter 2, also appears here when I discuss how to individuate pragmatic objects, but I do not use this specific terminology here. One important conceptual contribution of this chapter, particularly in relation to the pragmatist-inspired work that I have done in the thesis, is that the formulation of the notion of *habits of action* is first given here. While the idea behind the notion was already present in Chapter 2, it now receives a clearer and more detailed treatment.

One important discontinuity between Chapters 1 and 2, on the one hand, and Chapter 3, on the other hand, is the fact that considerations about pragmatic agreement do not appear in Chapter 3. This is a question that is likely to come up for the reader, so I shall say something about it here. When I first started thinking about memory and episodic hypothetical thought, I was tempted to adopt the framework that I developed in Chapter 2 in this new topic. In other words, I wondered if an account of the objects of memory, and likewise an account of the veridicality of memory, could be given in terms of pragmatic agreement. Despite there being close similarities between these questions and their perception correlates, an important difference is that, unlike the objects of perception, the objects of memory are not contemporary with states of remembering, which makes it hard to see how subjects can interact with them and reach pragmatic agreement. Moreover, because the distinction between *actual* and *ideal* or *counterfactual* pragmatic contexts of agreement was not yet clear to me when I first started working on this chapter, I thought that it would not be possible to simply transfer the framework of Chapter 2 to Chapter 3. For this reason, the

strategy I adopted was to apply *some* elements of the framework proposed in Chapter 2 and adapt them to conceive of a new framework that could be applied to episodic hypothetical thought. So, while there are obvious connections between the perception framework and the one that will be developed here, they are not the same. The closest that I get to providing an analysis of memory in terms of the framework provided in Chapter 2 is section 3.4.3, and while the notion of “context of thought” is used there, it is not meant to be the same notion as the notion of “pragmatic contexts”. The “pragmatic” bit was left out on purpose because of the difficulties that I just discussed.

Finally, the reader might find it helpful to read the discussion article in the Appendix A as a complement to this chapter. In that short discussion piece, I provide an accessible discussion of the main motivations to take episodic memory to be a form of mental time travel (or, alternatively, episodic hypothetical thought). Moreover, I offer a concise discussion of the main philosophical issues that arise in the context of this idea, which are explored in more detail in this chapter.

Outline

This chapter introduces a novel question, the question of the objects of episodic hypothetical thought, and provides an answer to it based on classical pragmatism. Section 3.1 introduces the question of the objects of episodic hypothetical thought and discusses the motivations to consider it. Sections 3.2 and 3.3 consider traditional accounts of the objects of memory as possible answers to the question of the objects of episodic hypothetical thought. Both direct realism or relationalism (Section 3.2) and indirect realism or representationalism (Section 3.3) are shown to be problematic. Section 3.4 explores our positive account of the objects of episodic hypothetical thought influenced by the classical pragmatism of Charles Peirce. Finally, Section 3.5 concludes by responding to an objection raised to the pragmatist account.

3.1 From the objects of memory to the objects of episodic hypothetical thought

When a subject remembers an event, to what is he related, in the first instance? What, in other words, are the direct objects of episodic memory? This is the traditional question of the *objects of episodic memory*.¹ Following Reid, direct realists have argued that the direct objects of memory are past events themselves. Following Locke and Hume, indirect realists have argued that the direct objects of memory are not events but rather representations of events, claiming that it is only in virtue of his (direct) relation to a representation that a remembering subject is (indirectly) related to an event. Contemporary discussions of the objects of memory continue to unfold largely within this early modern framework, but the framework predates the empirical psychology of memory, and we suggest that it may be time to move beyond it. We suggest, in particular, that it may be time to replace the traditional question of the objects of episodic memory with a question that takes contemporary psychological thinking about memory into account: the question of the *objects of episodic hypothetical thought*.

Adopting De Brigard's (2014a) terminology, we use "episodic hypothetical thought" as a blanket term referring not only to episodic memory (roughly, the capacity to remember events that occurred in the personal past) but also to episodic future thought (the capacity to imagine events that might occur in the personal future) and episodic counterfactual thought (the capacity to imagine events that might have but did not occur in the personal past).² In philosophy, it has generally been assumed that there is no need for an adequate positive account of episodic memory to refer to these other forms of episodic hypothetical thought — that, if they figure in an account of memory at all, they will do so merely negatively, as forms of imagination from which memory must be distinguished. In psychology, in contrast, it is increasingly taken for granted that episodic memory can be adequately

¹The term "episodic memory" originates in psychology (Tulving 1985a). Philosophers have historically referred to episodic memory using a variety of different terms but are increasingly adopting the psychological terminology (Perrin and Rousset 2014).

²In section 4, we introduce a fourth category of episodic hypothetical thought: future-oriented episodic counterfactual thought.

understood only if it is seen as one instance among others of a more general capacity for simulating possible episodes. On this view, episodic memory overlaps heavily with other forms of episodic hypothetical thought — also known as “mental time travel” — at every level from the neural to the phenomenological. This is not the place to summarize the evidence in favour of the psychological approach (see Michaelian 2016c for a review), and we will simply take for granted the consensus view that the more general category of episodic hypothetical thought is prior to the more specific category of episodic memory. This view suggests that the traditional question of the objects of episodic memory should be replaced with a new question: What are the direct objects of episodic hypothetical thought? When a subject thinks about a possible event, to what is he related, in the first instance?

In what follows, we defend a specific answer to this question. Our argument unfolds in two stages, with the first stage being primarily negative in character. We build in this stage on two distinct literatures. The psychological literature on episodic hypothetical thought, on the one hand, provides new insights into the relationship between episodic memory and other forms of episodic hypothetical thought. The philosophical literature on the objects of perception, on the other hand, which is considerably more advanced than the literature on the objects of memory, points to the need for and possibility of alternatives to traditional forms of direct and indirect realism. Building on these two literatures, we arrive at two conclusions. In section 3.2, we show that, when the traditional question is replaced with the new question, as research on episodic hypothetical thought suggests it ought to be, a previously overlooked problem for direct realism comes to light: direct realists have usually been disjunctivists, but it turns out that direct realism leads to a form of disjunctivism far more extreme than has hitherto been appreciated. Indeed, the form of disjunctivism in question is unacceptably extreme, and our first conclusion is thus that direct realism should no longer be seen as a viable option. In section 3.3, however, we show that considerations originally adduced by philosophers of perception demonstrate, when applied to the domain of episodic hypothetical thought, that indirect realism is equally problematic: episodic hypothetical thoughts turn out to be “silent” in the sense that they do not establish their own satisfaction conditions. This, in turn, undermines indirect realism, and our second conclu-

sion is thus that indirect realism should likewise no longer be seen as a viable option.

Overall, the first stage of the argument suggests that an adequate account of the objects of episodic hypothetical thought will preserve the desirable features of direct and indirect realism while eliminating those that render them unviable. The second stage of the argument is positive in character, aiming to develop an account of this sort. We build in this stage on a third, older literature, finding, in section 3.4, the starting points for an alternative to direct and indirect realism in Peircean pragmatism and developing these into a detailed pragmatist account of the objects of episodic hypothetical thought. Given its unfamiliar character, the pragmatist account is bound to be controversial, and, while we cannot hope to respond here to all objections that the account is likely to encounter, we do, in section 3.5, respond to what we take to be the most pressing objection to the account, arguing that it can, initial appearances to the contrary notwithstanding, successfully accommodate the role of auto-noetic consciousness in episodic hypothetical thought. Thus, while we may not succeed in persuading readers antecedently sceptical of pragmatism that the pragmatist account is the way to go, our third and final conclusion is nevertheless that it represents a promising avenue for future research in this area.

We note in advance that the argument of the paper is somewhat complex, as it both takes on an unfamiliar question (the question of the objects of episodic hypothetical thought) and approaches that question via an unorthodox theoretical framework (pragmatism). We ask our readers to bear with us. The pragmatist account may be unorthodox, but its ability to avoid the problems that undermine orthodox direct and indirect realist accounts speaks in its favour. Moreover, even readers who are ultimately unpersuaded by the pragmatist account (the second stage of our argument) may be persuaded by the claim that direct and indirect realism fail to provide adequate accounts of the objects of episodic hypothetical thought (the first stage of our argument). As this claim does not itself rest on pragmatist assumptions, it is, we believe, incumbent upon opponents of pragmatism to show how the objects of episodic hypothetical thought can be accounted for within a nonpragmatist theoretical framework.

3.2 Direct realism and extreme disjunctivism

This section sets out the first part of the negative stage of our argument. After motivating direct realism (section 3.2.1), we describe two problems for the view (section 3.2.2). The first, the co-reality problem, is a generalization of the co-temporality problem familiar from discussions elsewhere in the philosophy of memory. The second, the problem of confabulation and misremembering, is a counterpart to the problem of hallucination and illusion familiar from discussions in the philosophy of perception. Our discussion of these problems will show that, once episodic memory has been situated as a form of episodic hypothetical thought, it becomes clear that direct realism implies an unacceptably extreme form of disjunctivism.

3.2.1 Motivating direct realism

Just as direct and indirect realism about the objects of perception can be seen as particular ways of spelling out basic relationalist and representationalist accounts of the objects of perception, direct and indirect realism about the objects of memory can be seen as ways of spelling out basic relationalist and representationalist accounts of the objects of memory. According to relationalism, memory is fundamentally relational, in the sense that it necessarily involves a relation between the remembering subject and the remembered event. For the relationalist, states of remembering involve remembered events as constituents, which implies that, if one is not appropriately related to an event, one is not literally remembering. According to representationalism, memory is fundamentally representational, in the sense that it does not necessarily involve a relation between the remembering subject and the remembered object. For the representationalist, states of remembering do not involve remembered events as constituents; instead, they involve internal representations, which implies that one might literally remember despite not being appropriately related to an event. Relationalism comes in a variety of forms, and it will be crucial to our argument here that forms of relationalism other than direct realism are available, as the pragmatist account can be seen as an unusual form of relationalism. Direct realism — which can be

defined as the result of combining generic relationalism with the naive realist assumption that the event to which the remembering subject is related is an ordinary, concrete, particular event — is, however, the only form of relationalism that has so far received sustained attention, and it therefore provides the starting point for our discussion.

One reason often offered in favour of direct realism is that, unless the indirect realist can identify criteria that enable the subject to distinguish genuine from merely apparent memory representations, indirect realism may lead to scepticism about memory knowledge (Bernecker 2008). Some indirect realists have argued that such criteria can in fact be identified (Michaelian 2016c), but these arguments do not respond to what what is often considered to be the most intuitively powerful reason in favour of direct realism, namely, that it aligns with the phenomenological directness of remembering. Its phenomenological plausibility has meant that, while indirect realism has also found many adherents, direct realism has remained a popular view.

There are significant disagreements among different varieties of direct realism, but all direct realists are in agreement on the naive realist assumption identified above. In an oft-quoted formulation, Laird (1920), for example, maintains that “memory does not mean the existence of present representatives of past things. It is the mind’s awareness of past things themselves”. More recently, Bernecker emphasizes that, “on the [direct] realist view, what one is directly aware of in memory is the past event *in propria persona*, and not some representation of it” (Bernecker 2008, 68). Debus is similarly explicit, arguing that “[t]he [recollectively] remembered object or event itself is a part of the [recollective] memory” (Debus 2008, 405). There is no hint in any of these formulations that the objects of memory are anything other than ordinary, concrete, particular past events.

3.2.2 Extreme disjunctivism

Though it remains popular, direct realism faces serious problems. Here, we argue that two such problems ultimately undermine the appeal of direct realism.

From the co-temporality problem to the co-reality problem

The first problem arises because the mental state of remembering and the event that is remembered are located at different points in time, making it difficult to see how the latter might be the direct object of the former. Direct realism about perception may face a version of this “co-temporality problem” (as Bernecker 2008 refers to it) as well, since the process of perceiving takes time to unfold (Russell 1992). But, whereas the object perceived and the state of perceiving often at least overlap in time, the event remembered and the state of remembering may be separated by periods of many years, and the problem is arguably more urgent for direct realism about memory.

Responding to the co-temporality problem, Debus (2008) argues that direct realists can allow that the past event can serve as the direct object of the present mental state as long as they assume that the latter stands in the right sort of causal relationship to the former — in other words, as long as direct realism is combined with the causal theory of memory (Martin and Deutscher 1966). Given, however, that — as Debus herself emphasizes — what it is for an event to be the direct object of a mental state is for it to be a constituent of that state, it remains difficult to see how an event located in the past might be the direct object of a state of remembering located in the present. Bernecker (2008), who, like Debus, combines direct realism with the causal theory, argues that direct realists can overcome this difficulty if they reject presentism, the view that events exist only at the times at which they occur, in favour of eternalism, the view that events do not cease to exist once they have occurred (Price 1936). If past events continue to exist once they have occurred, they are in principle available to stand in constitutive relations with present mental states, and adopting eternalism might thus in principle allow the direct realist to avoid the co-temporality problem.

Eternalism is an intuitively implausible view. Some may already be convinced of eternalism for independent reasons. To them, the eternalist solution to the co-temporality problem comes at no extra metaphysical cost. Others, of course, are not so convinced. To some of these, the price of the solution may seem to be too high. To others, the price may seem to be worth paying, but we argue in the remainder of this section that, if direct realism is

understood as an answer not just to the question of the objects of episodic memory but also to the question of the objects of episodic hypothetical thought, the cost of the solution goes up dramatically. It eventually turns out, moreover, to be inapplicable to forms of episodic hypothetical thought other than episodic memory, rendering it unavailable even to those convinced of eternalism for independent reasons.

Direct realism about the objects of episodic hypothetical thought as a whole would be an unusual view, but it is not difficult to motivate. The most intuitively powerful reason in favour of direct realism about the objects of episodic memory, as we have seen, is the phenomenological directness of remembering: when one remembers an episode, one (in most cases) feels as if one is attending to the episode itself, not to an internal representation of the episode. Similarly, when one imagines a future or counterfactual episode, one (in most cases) feels as if one is attending to the episode itself, not to an internal representation of the episode. Thus forms of episodic hypothetical thought other than episodic memory share the phenomenological directness of episodic memory, and we might appeal to this shared phenomenological directness to motivate a direct realist account of the objects of episodic hypothetical thought as a whole. According to direct realism about episodic hypothetical thought, the direct objects of episodic hypothetical thought would be possible events — actual past events, in the case of episodic memory, possible future events, in the case of episodic future thought, and merely possible past events, in the case of episodic counterfactual thought.

Despite its phenomenological plausibility, this account immediately encounters serious stumbling-blocks. Because direct realism about episodic hypothetical thought holds that the direct objects of episodic future thought are future events, it faces a version of the co-temporality problem significantly more severe than that faced by direct realism about episodic memory. What we might think of as the “extended co-temporality problem” is the problem of explaining how events that were experienced by the subject in the past, as well as events that merely might be experienced by the subject in the future, can be constituents of his present mental states. Building on the eternalist solution, the direct realist might argue that events do not exist only at the times at which they occur, not only in the sense

that an event does not cease to exist once it has occurred but also in the sense that it does not begin to exist only when it occurs. If events are (so to speak) eternal in both directions, both past and future events are in principle available to stand in constitutive relations with present mental states, and adopting this extended form of eternalism might thus in principle allow the direct realist to avoid the extended co-temporality problem.

Extended eternalism does not, however, provide a full explanation of how possible events can be the direct objects of episodic hypothetical thought. Because direct realism about episodic hypothetical thought holds that the direct objects of episodic counterfactual thought are counterfactual events, it faces a problem yet more severe than the extended co-temporality problem. What we might think of as the “co-reality problem” is the problem of explaining how events that merely could have been experienced by the subject in the past might be constituents of his present mental states. In order to avoid the co-reality problem, the direct realist might argue that events exist regardless not only of when they occur but also of whether they occur at all. This solution, it should be emphasized, comes at a much higher metaphysical price — a price on a par with that of modal realism — than the eternalist solutions introduced in response to the co-temporality problem and the extended co-temporality problem.

Even if we are in principle willing to pay the price, moreover, the solution is subject to a fatal difficulty. The original co-temporality problem arose because it was unclear how an event might serve as the direct object of a mental state despite not existing at the same time as the mental state. Eternalism provides a solution to this problem, but it does so, as we have seen, only given the further assumption that episodic memory involves an appropriate causal relationship to the subject’s experience of the past event. The variants of eternalism considered here fail to provide a solution to the extended co-temporality problem or to the co-reality problem simply because, in forms of episodic hypothetical thought other than episodic memory, the subject has not experienced the relevant events and therefore stands in no suitable causal relationship to them. In principle, merely possible events might be available to stand in constitutive relationships with thoughts about them; in practice, it is entirely unclear what mechanism might enable them to do so.

The prospects for applying direct realism to episodic future thought and episodic counterfactual thought thus appear dim. Indeed, Debus (2014; cf. Perrin 2016) has explicitly argued that there is a difference in kind between episodic memory and episodic future thought, on the ground that, in episodic memory, the subject stands in an experiential (ultimately, a causal) relation to a particular event, whereas, in episodic future thought, the subject does not and could not stand in an experiential (causal) relation to a particular event. She thus in effect adopts a form of disjunctivism, with episodic memory as one disjunct and episodic future thought as the other. A subject is, of course, no more capable of standing in an experiential relation to a counterfactual event than he is of standing in an experiential relation to a future event; adopting a disjunctivist view of episodic memory and episodic future thought will thus naturally lead us to adopt disjunctivism about episodic hypothetical thought as a whole, with episodic memory as one disjunct and episodic future thought and episodic counterfactual thought as the other.

The problem of confabulation and misremembering

Disjunctivism about episodic hypothetical thought is distinct from traditional disjunctivism, according to which there is a difference in kind, within the category of episodic memory, between successful and unsuccessful remembering. Traditional disjunctivism is imposed on the direct realist about memory by the need to respond to a second problem, a problem posed by confabulation and misremembering analogous to the problem posed by hallucination and illusion the need to respond to which imposes a similar form of disjunctivism on the direct realist about perception.³

Disregarding unnecessary technical details, the argument runs as follows. (1) The mental states at issue in cases in which we remember events that did not occur, i.e., cases of confabulation (or “memory hallucination”), or in which we remember events that did not occur as we remember them occurring, i.e., cases of misremembering (or “memory illusion”), cannot have past events as their objects. (2) The mental states at issue in cases of successful remembering and those at issue in cases of unsuccessful remembering, includ-

³See Robins 2016a on the relationship between the confabulation/misremembering distinction and the hallucination/illusion distinction.

ing confabulation and misremembering, have objects of the same kind. Therefore, (3) the mental states at issue in cases of successful remembering do not have past events as their objects. While this is not directly entailed by 3, the representationalist invites us to draw the further conclusion that (4) the direct objects of memory, both in cases of successful remembering and in cases of unsuccessful remembering, are internal representations. Of the steps of this argument, 1 is the least controversial. In order to avoid having to accept 3 — and hence potentially being led to accept 4 — direct realists have therefore rejected 2, the claim that the objects of memory must be of the same kind in successful and unsuccessful remembering. Bernecker, for example, maintains that “[w]hat we are aware of in genuine memory is different from what we are aware of in ostensible memory. In genuine memory we are aware of past events. In ostensible memory, if we are aware of anything . . . we are aware of mental entities” (Bernecker 2008, 74).

Putting traditional disjunctivism together with disjunctivism about episodic hypothetical thought, it becomes clear that direct realism in fact leads to a disjunctivism considerably more extreme than has generally been recognized, with successful memory as one disjunct, and future thought, counterfactual thought, *and* confabulation and misremembering as the other. This extreme disjunctivism might seem to be acceptable as long as we take ourselves to be entitled to single successful remembering out as having a special status. But a look at, first, the ordinariness of confabulation and misremembering and, second, the centrality to our mental lives of forms of episodic hypothetical thought other than episodic memory suggests that we are not so entitled. Indeed, it suggests that direct realism in fact leads to a disjunctivism so extreme as to be clearly unacceptable, a disjunctivism according to which the size of the “successful memory” disjunct at best is dwarfed by that of the “other” disjunct and at worst may be outright empty.

Setting forms of episodic hypothetical thought other than memory aside for the moment, consider confabulation and misremembering. It is important to note at the outset that, by positing a difference in kind between successful memory, on the one hand, and confabulation and misremembering, on the other hand, disjunctivism makes what intuitively seems to be a unified mental phenomenon into something fundamentally disunified. Disjunctivism

should therefore be seen not as a position which is intrinsically attractive but rather as a position to which the direct realist is compelled to retreat in order to save his view. It is thus unsurprising that defences of disjunctivism often implicitly assume that successful memory is the norm, i.e., that confabulation and misremembering are exceptions to the rule, as this assumption makes the “successful memory” disjunct seem larger or more central than the “unsuccessful memory” disjunct, thus minimizing the intuitive implausibility of the view. The assumption is, however, unjustified. There is no way of estimating the precise frequency of confabulation and misremembering outside the laboratory, but what we have learned about the workings of memory from laboratory studies tells us that they must be very frequent indeed. Loftus and her colleagues, for example, have demonstrated that it is possible to implant memories of entire events that were not experienced by the subject (Loftus 2005). Crucially, no special technology is required to implant such memories. Simply having subjects repeatedly imagine an event is often sufficient to lead them to “remember” it, which suggests that we sometimes unknowingly implant memories in ourselves. Such cases of confabulation may be relatively rare, but cases of misremembering are likely to occur more frequently. Again, no special technology is required to induce subjects to misremember. In the Deese-Roediger-McDermott paradigm, for example, subjects who are simply presented with lists of thematically-related words often end up remembering non-presented but thematically-consistent lure words (Gallo 2010). Analogous conditions occur routinely in non-laboratory settings, suggesting that misremembering is anything but rare.

Confabulation and misremembering occur because memory has a reconstructive, rather than a reproductive, character: rather than being a matter of the preservation of a representation or a relationship established at the time of experience, remembering is, as Bartlett put it in his foundational study (1932), always a matter of creating something anew on the basis both of the subject’s past experience and of his present state of mind. The evidence for the reconstructive character of remembering that has accumulated since Bartlett’s time is overwhelming, and we will not attempt to summarize it here (see Michaelian 2011 for an overview). Suffice it to say that the evidence makes it clear that, while full-blown confabulation appears to be relatively rare in healthy subjects, misremembering is clearly a

widespread and inevitable consequence of the normal functioning of memory. This suggests that the size of the “unsuccessful memory” disjunct is far larger than traditional disjunctivists have taken it to be. Depending on how the argument from confabulation and misremembering is interpreted, in fact, it may even be virtually empty. Research on constructive memory suggests that, because all memories include details not derived from experience of the relevant events, they inevitably depart to some extent from experience and are thus to some extent false (e.g., Conway and Loveday 2015). All *memories*, in other words, are to some extent *mismemories*. This goes, for example, even for so-called “flashbulb memories”, exceptionally detailed and phenomenologically compelling memories of dramatic events (Newman and Lindsay 2009). Thus, if the argument from confabulation and misremembering is interpreted in such a way that it implies that all mismemories belong in the “unsuccessful memory” disjunct, there may be little or nothing left in the “successful memory” disjunct.

Turning to episodic hypothetical thought, while it is difficult to estimate the relative frequencies with which different forms of episodic hypothetical thought occur, there is every reason to think that episodic future thought occurs quite frequently (D’Argembeau et al. 2011): from an adaptive perspective, it is to be expected that we spend as much time thinking about and planning for the future as we do thinking about the past, and quite possibly more. Much of the time that we do spend thinking about the past, moreover, is devoted not to attempting to faithfully reproduce past events but rather to exploring what might have been had things gone differently — that is, to episodic counterfactual thought, rather than to episodic memory (De Brigard 2014a). Occurrences of episodic memory thus may account for a relatively small fraction of occurrences of episodic hypothetical thought.

In fact, research on mental time travel, like research on constructive memory, suggests that direct realism may ultimately be unable to acknowledge any instances at all of successful or genuine memory. The argument from confabulation and misremembering focuses on cases in which memory is at least to some extent inaccurate with respect to past events. But there are also cases in which memory is accurate with respect to past events but which direct realism cannot classify as instances of genuine memory. On views such as Debus’s and

Bernecker's, we have seen, a past event is able to serve as a constituent of a present state of remembering only because memory involves an appropriate causal relationship between the present state and the past event. In other words, in cases in which no causal relationship obtains, the past event cannot serve as a constituent of the present state, regardless of the accuracy of the apparent memory. In such cases, the present state will, by direct realist standards, fail to qualify as an instance of genuine memory. What mental time travel research seems to imply, however, is that, even in many cases of accurate apparent memory, no causal relationship between the present state and the past event obtains (Michaelian 2016a). The overlap among episodic memory, episodic future thought, and episodic counterfactual thought suggests that episodic memory is a form of episodic imagination. Like other forms of episodic imagination, remembering draws on past experience to simulate a target event. Forms of episodic imagination other than episodic memory, of course, cannot and do not draw on experience of the particular events they target; similarly, episodic memory can but need not draw on experience of the events it targets. In many cases, the subject's experience of a target event does play a role in shaping his memory of the event, and, in such cases, there may be an appropriate causal link between the present memory and the past event. But even in such cases, experience of the target event is typically not alone in shaping the memory — there may be equally strong or even stronger links to one or more other events. And in some cases, experience of the target event plays no role at all in shaping the memory — instead, the memory builds entirely on experience of other events. In such cases, memory involves no causal link whatsoever with the past event. The upshot, again, is that the “successful memory” disjunct may be empty or virtually empty.

This concludes the first part of the negative stage of our argument. To summarize: The form of disjunctivism traditionally thought to follow from direct realism is already unappealing. The extreme form of disjunctivism that in fact follows from direct realism is far less appealing: if extreme disjunctivism is right, genuine memory is, if it occurs at all, not the rule but rather a rare exception. We conclude that, once episodic memory has been situated as a form of episodic hypothetical thought, direct realism no longer appears to be a viable option.

3.3 Representationalism and the silence of episodic hypothetical thought

Turning to the traditional alternative to direct realism — indirect realism or representationalism — this section sets out the second part of the negative stage of our argument. Indirect realism would in principle provide a means of avoiding extreme disjunctivism. After reviewing the motivation for representationalism (section 3.3.1), however, we show that it, like direct realism, is subject to a previously-overlooked but fatal problem (section 3.3.2): just as perceptual states are “silent” in the sense that they do not establish their own satisfaction conditions, so, too, are episodic hypothetical thoughts. This undermines the representationalist approach.

3.3.1 Motivating representationalism

The most powerful reason in favour of direct realism, we saw above, is the phenomenological directness of remembering. The most powerful reason in favour of representationalism, in contrast, is its capacity to give an account of the objects of memory that avoids both the co-reality problem and the problem of confabulation and misremembering. Representationalism can provide such an account simply because it holds that, when we remember past events, we are directly related not to events themselves but rather to internal representations of events.

We take a mental state to be representational if it presents the world as being a certain way to the subject, i.e., if it has satisfaction conditions. A representational state is accurate if its satisfaction conditions are met by a state of affairs in the world; otherwise it is inaccurate. Taking episodic memories to be representational in this sense straightforwardly allows us to avoid the co-reality problem. According to representationalism, when we engage in episodic hypothetical thought, the direct objects of our mental states are internal representations of events. Because these representations are located in the present, there is no mystery about how they might serve as the direct objects of present mental states: just

as the representations involved (on a representationalist account) in perception can represent scenes that are not before the subject's eyes or do not occur at all, the representations involved in episodic hypothetical thought can represent events that occur at other times or that do not occur at all. Taking episodic memories to be representational also allows us to avoid the problem of confabulation and misremembering. Representationalism avoids this problem because it denies that events serve as the direct objects of memory. The difference between successful memory and unsuccessful memory is to be understood in terms of the accuracy of the relevant representations: in cases of misremembering, the satisfaction conditions of the representations are met only to some degree, and, in cases of confabulation, they are not met at all.

Representationalism thus allows us to avoid not only traditional disjunctivism but also extreme disjunctivism. Because it places memory in one category and other forms of episodic hypothetical thought in another, direct realism is bound to treat episodic memory as something exceptional. In contrast, because it says that the direct object of any instance of episodic hypothetical thought — whether episodic memory, episodic future thought, or episodic counterfactual thought, and whether successful or unsuccessful — is a representation, representationalism provides a unified treatment of episodic hypothetical thought as a whole.

3.3.2 The silence of episodic hypothetical thought

Though it therefore has considerable appeal, indirect realism faces a serious problem. It is the role that it assigns to satisfaction conditions that allows representationalism to provide a unified account of episodic hypothetical thought as a whole. But it is also this role that gets representationalism into trouble, for there is good reason to think that episodic hypothetical thoughts do not in fact establish their own satisfaction conditions.

We build here on Travis's (2004) argument for the view that perceptual states are "silent" with respect to their capacity to establish satisfaction conditions. As Travis argues, perceptual states cannot be assessed for accuracy alone because there is more than one state of affairs that will satisfy how things appear to the subject. The satisfaction con-

ditions are determined, instead, by things that are external to perceptual states, such as the subject's doxastic states.⁴ Consider a pair of scenarios. In the first, you see an apple in the kitchen and decide to eat it, only to discover that it is a wax apple. In the second scenario, you see an apple in the kitchen but decide not to eat it, because you know that it is one of the wax fruits your partner bought last week. The visual experiences at issue in these two scenarios may appear exactly alike to you, but their satisfaction conditions are arguably different: the former experience is about a real apple, while the latter is about a fake apple. Because how the experiences appear to you is the same while the satisfaction conditions are different, we may conclude that satisfaction conditions are determined by something external to perception — in this case, your beliefs.

Episodic hypothetical thought, we suggest, is similarly silent. Silence here refers, in particular, to the capacity of the mental state to determine the temporal location (past/future) and modal status (actual/possible) of the event. Consider another pair of scenarios. In the first, you have an episodic hypothetical thought that we would intuitively describe as a memory of your tenth birthday party. It seems clear to you that you are remembering an event that occurred in the actual past and not merely imagining a possible future event. This seems clear to you because you have a set of beliefs that suggest it. You believe, for instance, that you are currently an adult, that the child presented in the thought is a younger version of you, and that human development is an irreversible process that goes from childhood to adulthood. What we suggest is that it is these and other relevant beliefs that determine that your thought is about an actual past event. Compare this first scenario to a second, in which you wake up one morning with an unusual set of mistaken beliefs. You believe, for example, that humans are born adults and become children only later in their lives. Now, suppose that you have the same episodic hypothetical thought as in the first scenario. Due to your mistaken beliefs, the thought is now about a possible future event. Despite still believing that you are an adult and despite taking the person presented in the thought to be you, your belief about how human individuals develop has changed, thus changing the temporal reference of the thought. In general, the idea is that, if the same

⁴Travis himself does not explicitly endorse a doxastic account, but this has been one important development of his argument (see, e.g., Antony 2011; Genone 2014).

thought can refer either to the past or to the future depending on the beliefs that accompany it, episodic hypothetical thought must be silent with respect to temporal location and modal status.

Note that the suggestion is not that subjects must be consciously aware of the beliefs that accompany their episodic hypothetical thoughts and that determine the satisfaction conditions of the latter, for beliefs can influence thoughts both at the personal level, when the subject is or can be conscious of them, and at the subpersonal level, when the subject is not necessarily conscious of them. Thus, even if it seems to the subject, at the personal level, that there are no relevant beliefs that accompany his thought, it does not follow that these are not present at the subpersonal level. We come back to this point below.

The silence of episodic hypothetical thought challenges the representationalist approach to the co-reality problem and the problem of confabulation and misremembering. Consider the co-reality problem. The fact that the objects of episodic hypothetical thought are representations is no longer sufficient to an explanation of temporal reference and of its modal status, for these are now explained by the beliefs that accompany episodic hypothetical thoughts. Because, when considered alone, representations can be satisfied by past/future and actual/possible events equally, and because the reference to past or future and the modal status of events are established by things external to representations, such as the doxastic states, temporal reference and modal status cannot be internal features of representations. Consider the argument from confabulation and misremembering. Representationalism similarly has difficulty explaining confabulation and misremembering, for the fact that the objects of episodic hypothetical thought are representations is also no longer sufficient to an explanation of which particular event serves as the satisfaction condition for a given occurrence of episodic hypothetical thought. Because assessing the accuracy of a representation requires the identification of the particular event that is being assessed, and because such an identification requires the determination of temporal reference and modal status, a given episodic hypothetical thought is said to be an occurrence of misremembering or an occurrence of confabulating only in relation to the beliefs that accompany those thoughts.

This concludes the second part of the negative stage of our argument. To summarize:

The core claim of this section is that episodic hypothetical thoughts are silent with respect to the temporal locations and modal statuses of the relevant events. It is not episodic hypothetical thoughts themselves but rather the beliefs that accompany them that determine those features. If correct, the “silence argument” undermines the motivation for adopting representationalism. We conclude that representationalism, like direct realism, no longer appears to be a viable option.

3.4 The pragmatist alternative

This brings the first, negative stage of our argument to a close. If our reasoning in sections 3.2 and 3.3 is on the right track, neither direct realism nor representationalism can provide a viable account of the objects of episodic hypothetical thought. An alternative to both accounts is thus required. The purpose of the second, positive stage of our argument is to develop such an alternative.

The alternative that we propose draws inspiration from classical pragmatism. An important aim of pragmatism has always been to overcome dichotomies created by classical philosophical systems. Among these dichotomies are the opposition between mind and body (see Dewey 1958; McDowell 1996; Godfrey-Smith 1998, 2010) and the dispute between direct and indirect theories of perception (see Haack 1994; Bergman 2007; Legg 2014b; Wilson 2016). The latter was of particular interest to Peirce, and, while he did not himself address the topic of episodic hypothetical thought, Peirce’s attempts to bridge direct and indirect theories of perception within a pragmatist framework suggests that pragmatism may contain insights relevant to the question of the objects of episodic hypothetical thought. Our strategy in the remainder of the paper will thus be to focus on one central Peircian idea, using it as a starting point for answering that question. This idea is the *pragmatic maxim*, Peirce’s claim that the meanings of concepts should be analyzed in terms of the conceivable practical effects that their referents have or would have. For example, the concept HARD, which Peirce defined as that which would not be scratched by many other things, is understood in terms of the effects hard things would have if we were to interact

with them. We begin (section 3.4.1) by motivating pragmatism. We then develop (section 3.4.2) a pragmatist account of the objects of episodic hypothetical thought, showing how the account handles temporal reference and distinguishes among different kinds of episodic hypothetical thought in a manner consistent with the silence of episodic hypothetical thought as described in section 3.3. Finally, we show (section 3.4.3) how the pragmatist account deals with the problems raised in section 3.2 — the co-reality problem and the problem of confabulation and misremembering — thereby avoiding disjunctivism.

Before proceeding, we acknowledge, again, that readers unfamiliar with pragmatism may find our account to be outlandish in certain respects. We ask such readers to bear with us, suspending judgement until they have seen what pragmatism allows us to say about the objects of episodic hypothetical thought. Moreover, we emphasize, again, that the arguments for the negative conclusions given in sections 3.2–3.3 does not rest on pragmatist assumptions. Readers unpersuaded by our positive argument in section 3.4 are thus free to accept our negative arguments and to develop nonpragmatist alternatives to our account.

3.4.1 Motivating pragmatism

We begin with some terminology. By “pragmatic objects”, we will mean objects of mental states that are immediately available to the mind, regardless of whether they exist. Objects such as tables and unicorns are pragmatic objects, because they are immediately available to us in thought, perception, memory, and so on. But things that we do not usually think of as “objects” can also qualify as pragmatic objects: mental states, such as a pain, and physical events, such as the World Cup final, can be the objects of our mental states — e.g., one can hope that one’s pain goes away or wish to attend the World Cup final. *Pragmatic* objects thus resemble *intentional* objects, as originally introduced by Brentano (1973), in the sense that they are the things that mental states are directed at. However, while Brentano took intentionality to be what distinguishes the mental from the physical, the pragmatist account defended here will treat pragmatic objects as ontologically neutral.

Pragmatic objects are also distinguished from intentional objects by the fact that they are always seen as objects *for* action; they are, that is, conceived of in relation to how they

prompt, or would prompt, subjects to act in virtue of being related to them. By “actions”, we will mean both physical actions and mental actions resulting from the relation between a subject and a pragmatic object. For example, picking up a stone on the beach as a souvenir counts as a physical action because the pragmatic object (the stone) prompts certain bodily movements. Similarly, thinking of the stone as a good souvenir counts as a mental action because the pragmatic object (the stone) prompts the occurrence of a certain mental state. (It is, of course, also possible for pragmatic objects to cause physical actions and mental actions at the same time, such as when one forms the belief that the stone would be a good souvenir and therefore moves one’s body to pick it up.)

Because pragmatic objects are identified at the phenomenological level, they are defined in ontologically neutral terms: they are objects of the mind that prompt action, meaning that they might or might not exist and might or might not be physical or material. The strategy employed here is thus similar to the phenomenological reduction proposed by Husserl (1982; 1988). In order to identify pragmatic objects, we suspend our previous knowledge of the world and focus solely on what is immediately present to the mind (see Gallagher and Zahavi 2013, 24). For example, when one hallucinates a dog, one’s hallucinatory experience has a pragmatic object because it prompts one to act. However, when we consider this pragmatic object in relation to other things that we know about the world, we might infer that it does not exist concretely, because it is not the kind of thing that would be perceived by other subjects in the same situation.⁵

The ontological neutrality of pragmatic objects raises a question: how can we distinguish between *existent* pragmatic objects, such as chairs and stones, and *non-existent* pragmatic objects, such as possible objects and possible events? This is where the pragmatic maxim comes into play. The maxim says that the meaning of a concept is determined by the practical effects that it would have in experience. In line with the maxim, we propose to distinguish between existent and non-existent pragmatic objects in terms of the practical effects these would have in experience. The idea is that pragmatic objects prompt actions, which, when they enable subjects to deal successfully with their environments, become ha-

⁵For more detailed discussions of these issues, see Chapter 2, Sect. 2.3.2 and 2.3.3.

bitual. For example, driving a car for the first time is a complex task that requires constant attention. Inexperienced drivers find it challenging to change gears, but, as they get used to driving, this becomes an automatic task. Through repeating the same actions to deal with the same complex task, drivers develop what we will call a “habit of action”, which is the ability to reproduce, in current situations, actions that enabled them to deal successfully with earlier similar situations. We can distinguish between two kinds of habits of action. “Teleological” habits of action are developed over the course of an individual’s life to achieve goals that are relevant to the individual — for example, driving. “Teleonomic” habits of action are evolved and consequently shared among individuals of the species — for example, the tendency of the squid to release ink when approached by a predator. This distinction enables us to distinguish between existent and non-existent pragmatic objects, for these differ with respect to the kind of habits of action that they recruit.

Consider the visual experience of seeing a dog in a lab. This experience will prompt certain actions. Suppose that you are surprised to see the dog and ask your colleagues why it is standing there. They tell you that there is no dog in the lab and that you must be hallucinating. In this case, you perceive a pragmatic object, but the actions generated by the object fail to be coordinated with the actions of your colleagues in that situation. In other words, the effects generated by the pragmatic object of your experience are different from the effects generated by the pragmatic objects of your colleagues’ experiences. Thus, despite the fact that your experience has a pragmatic object, the object of your experience does not exist in the sense that it fails to generate the practical effects that would coordinate your experience with the experiences of your colleagues. On our account, the difference between existent and non-existent objects is one of degree and not one of kind. The difference of degree is explained in terms of the kinds of habits of action that are recruited by an object. A pragmatic object exists if it recruits teleonomic habits of action in multiple subjects, thus allowing for the coordination of actions in meaningful ways among those subjects. For example, to complete a puzzle, humans can coordinate their actions with each other in meaningful ways because they can identify shape and color in similar ways, which is only possible because they share certain evolved perceptual capacities. In contrast, a

pragmatic object is non-existent when it recruits only teleological habits of action, which usually do not allow for intersubjective coordination of actions. For example, a thought of a tenth birthday party will recruit different habits of action in different individuals, as those habits will depend on how each individual interacted with birthday parties in the past. We thus avoid using ontological notions, such as “physical” and “mental”, to distinguish between existing and non-existing things. This makes explicit a crucial commitment of the pragmatist account, which is that, because the objects of different mental states are all pragmatic objects, they are all fundamentally of the same kind. This commitment is crucial because, as we will see, it allows us to say that the relationship of mental states to existing and non-existing objects is of the same kind, which is an important step towards avoiding disjunctivism.

The motivation that we have offered here for our account of pragmatic objects as the objects of thought may strike some as insufficient. Supplying a fuller motivation for the account would require spelling out the particular version of pragmatism underlying it in more detail. But our intention here is not to provide a general argument for pragmatism but, more modestly, to show that a broadly pragmatist view can provide useful insights into the relationship between the mind and its objects. What the pragmatist account offers is a triadic understanding of mental relations, in which a mental relation is viewed as being irreducibly constituted by, first, a mental state, second, its pragmatic object, and third, the actions prompted by their interaction. This “triadic view” offers an alternative to two-place or dyadic view of mental relations (see Rowlands 2017), according to which they are constituted solely by how mental states stand with respect to their objects. The key novelty introduced by the triadic view, in contrast to the dyadic views which are implicitly assumed by both representationalism and relationalism, is that it gives actions a constitutive role in the determination of mental relations. In doing so, it offers a new perspective on the problems for relationalism and representationalism outlined in sections 3.2 and 3.3.

We apply the pragmatist framework to these problems in the remainder of section 3.4. Our plan of attack is the following. In section 3.4.2, we propose an account of the temporal reference of episodic hypothetical thought and of the actuality or possibility of its objects

in terms of the habits of action that are recruited by pragmatic objects. We build on this account to address the problems raised above for direct and indirect realism, considering, first, the problem of confabulation and misremembering and, second, the co-reality problem. With respect to the former, we employ the notion of habits of action to argue that confabulation and misremembering can be distinguished based on the inferences that subjects make when they are related to pragmatic objects. With respect to the latter, we rely on the characterization of pragmatic objects as ontologically neutral to argue that their existence (or non-existence) is determined only in relation to the actions that they prompt in subjects. In section 3.4.3, we conclude that the account therefore has the potential to avoid extreme disjunctivism.

3.4.2 Applying pragmatism

We begin with the temporal reference of episodic hypothetical thought. We then consider the actuality/possibility of its objects.

Temporal reference

As suggested above, the temporal reference of an episodic hypothetical thought — its reference to the past or the future — depends on the habits of judging that are recruited by its pragmatic objects. A habit of judging is constituted when subjects interact with pragmatic objects and form beliefs about those objects. For example, when eating an apple for the first time, a subject might satiate his hunger and have a pleasant experience, which will lead to the formation of judgments about apples, such as “apples are nutritious” and “apples are tasty”. These judgments will then serve as guides for action in future interactions with apples. If faced with an apple again, the subject will be inclined to perform certain actions (e.g., eat the apple) rather than others (e.g., throw it away). In short, habits of judging are beliefs that influence the subject to act in a certain way based on the outcome of previous experiences. Note that, although one is often aware of some of the beliefs that influence one’s actions, our account is compatible with the possibility that many habits of judging influence thoughts sub-personally, i.e., without one’s being aware of their presence.

On the pragmatist account, episodic hypothetical thought establishes temporal reference in virtue of habits of judging of this sort. To illustrate, suppose that you are thinking about how your life will be in ten years. The pragmatic object of your thought is an event in which there is a person who looks reasonably like you, is ten years older than you are, and lives in a beach house. Despite individuating a state of affairs, the pragmatic object does not establish a particular temporal reference. The reference to the future established by the thought is due rather to the habits of judging that accompany it. For instance, the beliefs that the person in question looks reasonably like you, that you are not currently that old, and that situations in which you are older or where you have a beach house are still to come, allow you to infer that the thought is about a possible event in the future. The same goes for thoughts about the past. Suppose that you are thinking about your tenth birthday party. The pragmatic object of your thought is an event in which there is a child who looks like you when you were ten. Again, despite individuating a state of affairs, this pragmatic object does not by itself establish a particular temporal reference. The reference to the past established by this thought is due to the habits of judging that accompany it. For instance, the beliefs that the individual in question is you, that you are no longer ten, and that situations in which you are younger have already happened or are no longer possible allow you to infer that the thought is about the past.

An obvious objection to this account of temporal reference is that it is inconsistent with the phenomenology of episodic hypothetical thought. When we think about past or future events, we do not usually engage in higher-order cognitive processes in order to form judgments about their temporal locations. In most cases, it seems to one that those thoughts simply present themselves to one as being about the past or the future. We discuss this objection in detail in section 3.5, where we argue that the fact that temporal reference is sometimes built into the phenomenology of episodic hypothetical thought is therefore not incompatible with the claim that it is underwritten by habits of judging.

Varieties of episodic hypothetical thought

With an account of temporal reference in place, we turn to modal status. On the pragmatist account, episodic hypothetical thoughts differ with respect to two kinds of habits of action recruited by their pragmatic objects. The first is responsible for presenting objects as being located in the past or future. We have just discussed habits of this kind. The second is responsible for presenting objects as being actual or merely possible. We now discuss habits of this kind, showing how the two kinds of habit in conjunction enable us to distinguish among different kinds of episodic hypothetical thought.

An episodic hypothetical thought is an *episodic memory* when its pragmatic object recruits habits of judging that place it in the past and habits of action that prompt the subject to act as if the object has been the case before. To illustrate, consider a first scenario. Some years ago, John decided to save five percent of his salary to attend a music festival that happens every year. In 2018 (the present), John remembers going to the 2017 festival. The pragmatic object of John's thought is an event involving an individual who looks reasonably like him going to the 2017 festival. John takes this thought to be about the past because he believes, for example, that the thought is about himself and that situations where he is younger are in the past. Moreover, because the pragmatic object makes John react to the thought in a certain way — e.g., he feels nostalgic about the festival and thinks it was a good idea to have saved money to attend it — the object is presented to him as being actual, i.e., as being something that happened before.

Similarly, an episodic hypothetical thought is a (*past-oriented*) *episodic counterfactual thought* when its pragmatic object recruits habits of judging that place it in the past and habits of action that prompt the subject to act as if the object was once, but no longer is, possible. To illustrate, consider a second scenario. In 2018, John thinks about how the 2017 festival would have been if it had not been definitively cancelled in 2016. The pragmatic object of his thought is an event involving an individual who looks reasonably like him in a festival occurring in 2017. John takes this thought to be about the past because he believes, for example, that the thought is about himself and that situations where he is younger are in the past. Moreover, because the pragmatic object makes John react to the thought in a

certain way, e.g., he does not feel nostalgic and he is not disposed to save part of his salary, the object is presented to him as having been once possible, but not as being possible at the moment of his thought.

An episodic hypothetical thought is an *episodic future thought* when its pragmatic object recruits habits of judging that place it in the future and habits of action that prompt the subject to act as if the object is possible. This is illustrated in a third scenario. In 2018, John thinks about how the 2019 edition of the festival will be. In this scenario, the festival was not cancelled in 2016. The pragmatic object of John's thought is the event describing an individual who looks reasonably like him going to the 2019 festival. John takes this thought to be about the future because he believes, for example, that 2018 is the present year and that 2019 will follow it. Moreover, because the pragmatic object makes John react to the thought in a certain way, e.g., he will continue to save his salary every month to attend the 2018 edition and look forward to the festival, the object is presented to him as being possible.

Finally, an episodic hypothetical thought is a *future-oriented episodic counterfactual thought* when its pragmatic object recruits habits of judging that place it in the future and habits of action that prompt the subject to act as if the object was once, but no longer is, possible.⁶ Note that future-oriented episodic counterfactual thought differs from past-oriented episodic counterfactual thought only in its temporal reference. Both are thoughts about things that were possible at a given moment but are no longer possible at the moment of thinking. This is illustrated by a fourth scenario. In 2018, John thinks about how the 2019 edition of the festival would have been if it had not been cancelled in 2016. The pragmatic object of John's thought is the event describing an individual who looks reasonably like him going to the 2019 festival. John takes this thought to be about the future because he believes, for example, that 2018 is the present year and that 2019 will follow it. Moreover, because the pragmatic object makes John react to the thought in a certain way, e.g., he will not be disposed to save his salary to go to the 2019 edition, the object is presented to him

⁶Episodic future thought and future-oriented episodic counterfactual thought are not normally distinguished from each other; one virtue of the pragmatist account is that it makes clear the importance of drawing this distinction.

as having been once possible, but not as being possible at the moment of his thought.

In short, while the particular habits of action recruited will vary from situation to situation, it is the presence of habits of action establishing temporal reference and habits of action marking objects as actual, possible, or not possible, that determine the nature of an episodic hypothetical thought.

One may worry that this account rules out the possibility of entertaining multiple thoughts about the same event.⁷ Suppose that, on several different occasions, you entertain thoughts about your tenth birthday party. Suppose further that the pragmatic object of one thought describes your friends and family as being there and you as having chocolate cake and that the pragmatic object of another thought describes your friend and family as being there but does not describe you as having chocolate cake. If the pragmatist account is right, one may worry, these two thoughts cannot in fact both be about your tenth birthday party, since their pragmatic objects have different features. The worry, in general, is that, given that there is significant variation in how we think about events over time, the pragmatist account implies that we can never think about the same event twice.

In response to this worry, we invoke the characterization, offered in section 3.4.1, of pragmatic objects as being identified at the phenomenological level. Because they are identified at this level, they are silent in the sense that they are not themselves about events in the world; reference to events, instead, obtains in virtue of the actions prompted by pragmatic objects. This means that, even if two thoughts have different pragmatic objects, as in the example above, it does not follow that they are about different events. This would follow given a dyadic view on which the features of the objects themselves would be responsible for establishing reference, but it does not follow given the triadic view. On the pragmatist account, while some similarity in terms of the phenomenological features possessed by their pragmatic objects is required in order for two thoughts to refer to the same event, that similarity alone does not determine coreference. In view of the scope of our discussion here, we have focused on how habits of action — including habits of judging — establish temporal reference and actuality/possibility. But reference is also to be explained in terms

⁷We are grateful to an anonymous referee for raising this issue.

of such habits. Consider again the example above. Despite the fact that the relevant pragmatic objects have different phenomenological features — the first describes you as having chocolate cake but the second does not — both thoughts may refer to the same event if they involve not only habits that place the event in the past and lead you to act as if it were actual but also, in consequence, habits that lead you to form the belief that the relevant event is the particular event of your tenth birthday party. Two thoughts will thus count as being about the same event as long as their pragmatic objects share some — but not necessarily all — phenomenological features, there are habits establishing the same temporal reference and modal status, and there are habits leading the subject to believe that the pragmatic objects refer to the same particular event. The issue of how episodic hypothetical thoughts refer to particular events is a complex one, and addressing it in greater detail is beyond the scope of this paper. We hope, however, that these considerations suffice to show that the pragmatist account is capable in principle of addressing it.

3.4.3 Avoiding disjunctivism

We have argued that different forms of episodic hypothetical thought can be distinguished in terms of temporal location and modal status, where these are determined by the habits of action that are recruited by their pragmatic objects. The problem of confabulation and misremembering, the co-reality problem, and the issue of disjunctivism remain to be addressed. Addressing these questions will occupy us for the remainder of section 3.4.3.

The problem of confabulation and misremembering

Before we discuss the pragmatist approach to confabulation and misremembering, let us be more precise about what we mean by “remembering”, “misremembering”, and “confabulation”. We will say that a mental state is a case of *remembering* if it allows for true inferences (i.e., inferences to true conclusions) and no false inferences (inferences to false conclusions) about the past; it is a case of *misremembering* if it allows for some true and some false inferences about the past; and it is a case of *confabulation* if it allows only for false inferences about the past. However, since genuine occurrences of remembering can

involve inaccurate elements (see, e.g., De Brigard 2014a), this distinction applies only in relation to particular contexts. To illustrate, suppose that a subject is trying to remember how the weather was during his tenth birthday party. If he remembers that the party was on a sunny day, and if that happens to be true, he is remembering the event accurately even if he remembers playing with his friend John, when it is false that John attended the party. Depending on whether we evaluate the mental state in relation to the weather or to John's attendance to the party, the assessment of it as true or false will vary. For this reason, we will say that a putative memory is evaluated in the "context of thought" in which it occurs, where a context of thought refers simply to the context specifying the elements relevant to assessing the accuracy of the inferences made by subjects.

With these clarifications in place, the pragmatist account says that a given mental state is an occurrence of remembering, misremembering, or confabulation depending on the nature of the inferences that are prompted by its pragmatic object in contexts of thought. The pragmatic object of a state of remembering prompts only true inferences about the past; the pragmatic object of a state of misremembering prompts some true and some false inferences about the past; and finally, the pragmatic object of a state of confabulating prompts only false inferences about the past. These definitions are compatible with the silence argument presented in section 3.3 and with the triadic view, since it is the actions — in this case, the inferences — generated by the interaction between a mental state and a pragmatic object that explain the differences between kinds of episodic hypothetical thought.

Consider, first, remembering. When one remembers something, one thinks about an event to which one was once related in perception. The objects of perception, in turn, are themselves pragmatic objects and recruit habits of action. For example, when you see a yellow chair, you form the perceptual judgment "the chair is yellow" because the habits of judging that are recruited are the habits to judge objects with such-and-such properties as chairs and objects of that color as yellow. If these habits allow you to make true inferences about your environment, then your experience is said to be veridical. If they do not, then it is said to be non-veridical. In parallel to this account of the veridicality of perception, the pragmatist account says that a mental state is an occurrence of remembering when, in

a given context of thought, the pragmatic object recruits habits of judging that allow the subject to make true inferences about the past. A “true inference about the past” is, in turn, one that would be true of the perception of the event. For example, in the birthday case above, the subject remembers correctly that it was a sunny day because the inferences he makes on the basis of this thought would also be true of the perception of the event. Note that we are not saying that the subject needs to be able to perceive the event again but only that, if this were possible, the inferences generated by remembering would be true of his perception of the event. Similarly, we are not claiming that successful remembering must allow the subject to make all the true inferences that perception of the event would enable him to make but only that it must allow him to make those that are relevant in the context of thought.

Consider, second, misremembering. In contrast to remembering, a mental state is an occurrence of misremembering when, in a given context of thought, the pragmatic object recruits habits of judging that allow for some true and some false inferences about the past, i.e., some inferences that would be true and some inferences that would be false of the perception of the event. In the birthday case, the subject misremembers because the pragmatic object of his thought allows for both true inferences (that it was a sunny day) and false inferences (that he played with John). Consider, finally, confabulation. A mental state is an occurrence of confabulation when, in a given context of thought, the pragmatic object recruits habits of judging that allow only for false inferences about the past, or inferences that would be false of the perception of the event. In the birthday case, the subject confabulates when, in the relevant context of thought, he judges mistakenly that he played with John because this inference would be false of the perception of the event.

Two points need to be emphasized here. The first is that the account does not require that the inferences made on the basis of remembering should have been made at the time of perceiving. For example, the subject above need not have formed the perceptual judgment “it is a sunny day” during his tenth birthday party for the thought to count as remembering. The relationship here is counterfactual, that is, the inference made on the basis of remembering is true if perception of the event would have yielded a similar perceptual judgment

that would be true. So, this account does not require an actual but only a counterfactual correspondence between inferences made on the basis of remembering and perceptual judgments.

The second point is that the pragmatist account does not require that the habits of action recruited in remembering match the habits of action recruited in perception. Because perception and memory ultimately relate to pragmatic objects, if such a match were required, there would be no distinction between remembering something and perceiving something. The account only requires that the truth-value of the inferences made on the basis of remembering in a context of thought and the inferences that would have been made on the basis of a perception of the event be the same. This highlights an important difference between the pragmatist account and causal theories of remembering (e.g., Martin and Deutscher 1966; see Debus 2017 for discussion). On the pragmatist account, genuinely remembering an event does not require the preservation of a trace connecting the actual memory to a previous perceptual state (see De Brigard 2014b; Robins 2017b for discussion). The subject must have perceived the event at some point in his life, but the connection between remembering and perceiving is established in counterfactual terms, i.e., in terms of whether the inferences generated by remembering *would* be true of the past perception of the event.

One might worry that the fact that it does not require the presence of a trace connecting remembering to a past perceptual experience renders the account unable to distinguish cases of remembering from cases of veridical confabulation (see Michaelian 2016b; Bernecker 2017). A veridical confabulation is a mental state such that its pragmatic object recruits habits of action that allows for true inferences of the past but such that the subject has never been perceptually related to the relevant event. In other words, the subject makes the right kind of inferences only accidentally. For example, a subject can confabulate being at the 2002 World Cup Final and only accidentally make true inferences about the match, such as that Ronaldo scored two goals for Brazil, when in fact the subject did not go to the stadium. The pragmatic object here prompts true inferences, but these might be due to sheer guessing. So, this is a confabulatory state that only accidentally happens to be veridical.

Veridical confabulation would pose a problem for the pragmatist account only if it held

that a subject can remember an event even if he did not perceive it in the past. But this is not what the account holds. Instead, it holds that we are to look at the inferences generated by remembering and consider whether those inferences would be true of the subject's perception of the event. In the case above, the mental state would count as a case of veridical confabulation rather than a case of genuine remembering because the inferences generated by the confabulation would not be true of any of the subject's past perceptual states. One might reply, however, that there could be cases of veridical confabulation in which the subject makes true inferences based on a past perceptual state. For example, by seeing Ronaldo score two goals on television, and not at the stadium, when the subject says that he saw Ronaldo score two goals at the final, he is making a true inference about the past based on a previous perceptual state, namely, watching the 2002 Final on television. This would mean, the reply continues, that the subject is remembering and not confabulating. In response, we point out that the reply rests on an ambiguity relating to the kind of inference the subject is making when he says that he saw Ronaldo scoring two goals. In the initial case described, the subject infers that he saw Ronaldo scoring at the stadium, but this is not the same thing as seeing Ronaldo scoring on the television. So, the inference he is making on the basis of his confabulation would not be true of any past perceptual state because it is not the case that he saw Ronaldo scoring from the stadium.

The co-reality problem

With the problem of confabulation and misremembering out of the way, we turn to the co-reality problem. The challenge is to explain how events that no longer exist might be constituents of mental states. Our account appeals here to the neutrality of pragmatic objects and the triadic view discussed above. The neutrality of pragmatic objects, again, refers to the fact that they are only said to be existent or non-existent, to be in the past or in the future, in relation to the habits of action that they recruit. Building on this, the triadic view says that the nature of a mental relation is irreducibly triadic, in that it can only be determined by looking (first) at the actions generated by the interaction between (second) a mental state and (third) a pragmatic object.

By treating the objects of episodic hypothetical thought as pragmatic objects, the pragmatist account avoids the co-reality problem. Because existence is understood in terms of the actions generated by the interaction between a mental state and a pragmatic object, instead of as a property ascribed to the objects of episodic hypothetical thought, the pragmatist account is not committed to the idea that non-existent things are constitutive parts of episodic hypothetical thought. So, for example, when a subject thinks about how his graduation would have been if he had attended another university, the pragmatic object describing this event is a constitutive part of his thought, but what determines whether his thought is about an existing or non-existing event is not an intrinsic feature of the pragmatic object but rather the actions generated by it. This account avoids problems pertaining to the ontological status of those objects while simultaneously preserving the distinction between existence and non-existence, which are now understood in relation to habits of action. By the same token, the co-temporality problem, too, is avoided, for pragmatic objects need not be located in the past/future in order for them to be about the past/future. Temporal reference, as we argued above, is established by habits of action. In this respect, the pragmatist account offers an alternative to the eternalism (Bernecker 2008) discussed in section 3.2, for it does not require us to say that past events continue to exist in the present.

A worry that arises at this point is that pragmatism is just a form of representationalism in disguise because, one might say, “pragmatic object” is just a new term for the old idea of a representation. For two reasons, this worry is unfounded. First, because pragmatic objects are silent in Travis’s (2004) sense, they do not establish their own satisfaction conditions. As discussed in section 3.4.1, pragmatic objects establish reference only in the context of the actions they prompt in subjects. So, unlike representations, which do establish their own satisfaction conditions, pragmatic objects are silent. Second, traditional representationalist views are based on discontinuist views of mind and world (see Menary 2009). In contrast, the pragmatist account incorporates a continuist view, on which the objects of the mind (so-called “representations”) and the objects of the world (so-called “physical things/events”) are viewed as being fundamentally the same. In line with the pragmatic maxim introduced earlier, the difference between them is accounted for in terms of the effects they generate in

experience. Thus, “pragmatic object” is not just a new term for the old idea of a representation, as pragmatic objects are incompatible with at least two important ideas associated with representations.

Disjunctivism avoided

We are aware that the account developed in the preceding sections presupposes claims that many readers will be inclined to reject. Our primary aim, however, is not so much to develop a definitive answer to the question of the objects of episodic hypothetical thought as it is to provide an initial discussion that might serve to set the stage for future work on the question. Thus, we invite readers unprepared to sign on to the full-blown pragmatist account to take a more modest point away from our discussion: even if the details of the pragmatist account turn out to be wrong, the suggestion that a proper understanding of the nature of the objects of episodic hypothetical thought requires moving away from a narrow focus on the nature of mental states and their objects to a broader focus including elements external to this dyad, such as the actions or beliefs generated by their interactions, might be right. Nevertheless, the merits of the pragmatist approach should not be underestimated, and, in the remainder of this section, we show how the approach enables us to avoid disjunctivism.

Disjunctivism poses a problem because it postulates a fundamental discontinuity between remembering and other forms of episodic hypothetical thought (Debus 2014; Perin 2016). This discontinuity, as we saw above, forces the relationalist to adopt problematic commitments in order to save his view. The pragmatist account, in contrast, offers a framework in which the continuity between remembering and other forms of episodic hypothetical thought is preserved, thus offering relationalists an option that does not require disjunctivism.⁸

The discontinuity between remembering and other forms of episodic hypothetical thought postulated by disjunctivists is grounded in the view that the objects of remembering exist but the objects of other forms of episodic hypothetical thought do not. As Debus (2008; 2014) argues, remembering involves being in an experiential relation to events which can-

⁸For a similar proposal focusing on disjunctivism about perception, see Sant’Anna (2018b) or Chapter 2.

not possibly obtain in other forms of episodic hypothetical thought. In opposition to this view, the pragmatist account denies that remembering is distinctive because of the existence of such an experiential relation. Instead, the difference between remembering and other forms of episodic hypothetical thought is explained in terms of the habits of action recruited by their pragmatic objects. In contrast to disjunctivism, this account relies not on intrinsic features of the objects, such as whether they exist, but rather on the actions generated by them. Thus, because the pragmatist account focuses on action, even if we grant Debus's claim that remembering involves a kind of experiential relation that other forms of episodic hypothetical thought lack, it does not follow that we are dealing here with two fundamentally distinct kinds of mental states. This makes it possible to say, in turn, that all occurrences of episodic hypothetical thought relate to the same kind of thing — namely, pragmatic objects — which allows for an even stronger form of relationalism, one on which *all* occurrences of episodic hypothetical thought are relational.

A basic problem for the claim that all occurrences of episodic hypothetical thought are relational is that it is simply incoherent to say that we can be directly related to things that have ceased to exist or that never existed in the first place. This becomes a problem for the pragmatist account only if it is taken to incorporate a dyadic view of mental relations, which it does not. Dyadic views take mental relations to be two-place relations composed by a mental state and an object, such that, if one of the constituents is missing, the relation fails to obtain. So, because the objects of forms of episodic hypothetical thought other than episodic memory are viewed as non-existent objects, it is not possible for episodic hypothetical thought to relate to those objects. This view assumes, however, that the predicates “existent” and “non-existent” apply to objects *prior* to their being parts of mental relations. The pragmatist account denies this explicitly when it commits to a neutral characterization of pragmatic objects. By refraining from drawing the distinction by relying on ontological notions, “existence” and “non-existence” are applied to objects when they are considered as parts of irreducibly triadic relations; more specifically, those predicates are applied to objects in relation to the actions generated by the interaction between a mental state and a pragmatic object. Thus, because the pragmatist account commits to a triadic view of mental

relations, it does not require that episodic hypothetical thought relate to non-existent things in the sense implied by dyadic views.

It is important to note that pragmatic objects do not occupy a third and exclusive ontological category. The pragmatist account accepts the idea that things either exist or not, but it contests the possibility of drawing this distinction without taking into account our interactions with them. In a pragmatist spirit, we do not deny the importance of the distinction, but we do question the traditional use of ontological notions to frame it. The general motivation here lies in the conviction, shared by different pragmatist views, that reality is continuous and that, due to its continuity, the differences between things must be accounted for in a similarly continuous way.

This concludes the positive stage of our argument. To summarize: We have shown, in response to the problem of confabulation and misremembering, how a given occurrence of episodic hypothetical thought can be characterized as an instance of remembering, misremembering, or confabulation by considering whether its pragmatic object prompts wholly true, partly true/partly false, or wholly false inferences. We have shown, in response to the co-reality problem, how the triadic view and the pragmatist understanding of the existence and nonexistence of pragmatic objects allows us to avoid the idea that nonexistent things are constituents of episodic hypothetical thoughts. And we have shown, finally, how these features of the pragmatic approach enable it to avoid extreme disjunctivism. In doing so, the pragmatic approach offers an unconventional, but stronger version of relationalism, where *all* forms of episodic hypothetical thought are relational. This new relationalist account has the advantage that it avoids the problems raised to the more traditional versions of the relationalism — more specifically, direct realism — while preserving desirable features of representationalism, namely, the idea that all forms of episodic hypothetical thought are occurrences of the same kind.

3.5 Pragmatism and auto-noesis

With this, our main argument is at a close. The negative stage of the argument (sections 3.2–3.3) is meant to demonstrate that neither direct realism nor representationalism offers an adequate account of the objects of episodic hypothetical thought. The positive stage of the argument (section 3.4) is meant to demonstrate that the pragmatist approach represents a promising alternative. While we acknowledge that readers antecedently sceptical of pragmatism may not be persuaded by the positive stage of our argument, we reiterate that such readers may nevertheless be persuaded by the negative stage, which does not itself rest on pragmatist assumptions. We thus look forward to the development of nonpragmatist accounts of the objects of episodic hypothetical thought.

In the meantime, we offer, by way of conclusion, a response to (what we take to be) the most pressing objection to the pragmatist account. One might, as we have acknowledged, object to the account itself on the ground that the metaphysical status of the pragmatic objects that it posits is unclear. But one might also object to the argument by means of which we have motivated the account, and it is to an objection of this sort that we want to respond in this final section of the paper.

Our overall argument for the pragmatist account depends on the argument for the silence of episodic hypothetical thought developed in section 3.3. The key element of that argument is the claim that there is nothing internal to episodic hypothetical thoughts that suffices to establish satisfaction conditions for them. The objection to this claim is simply that, while perception may indeed be silent in Travis’s sense, episodic hypothetical thought is not, for it has an *auto-noetic* character. The view that episodic hypothetical thought — or at least episodic memory — has such a character has been particularly prominent in recent years (e.g., Mahr and Csibra 2018), but it has long been standard in both philosophy and psychology. James, for example, argued that memory involves a “feeling of the past direction of time” (1890, 650), while Russell saw memory as involving a “feeling of pastness” (1921, 161–162). Indeed, Klein (2015) has argued that this feeling — which has come, following Tulving’s (1985b) introduction of the term, to be known as auto-noetic consciousness — has

traditionally been taken to be definitive of episodic — as opposed to semantic — memory. In line with this tradition, Dokic, for example, has recently proposed an account on which episodic memory is characterized by an episodic feeling of knowing (Dokic 2014). Even those who are uneasy with the notion of a feeling of pastness typically simply relocate pastness from the phenomenology of episodic memory to its content. Fernández (2016), for example, characterizes the content of episodic memories as informing the subject that they originate in his past experience, and Martin (2001) and McCormack and Hoerl (1999) have defended similar views. Regardless of whether auto-noesis is characterized in terms of phenomenology or in terms of content, its involvement in episodic memory suggests that episodic memories are not in fact silent. If this is right, then episodic memories, and perhaps episodic hypothetical thoughts of other kinds — Michaelian (2016c) has argued that episodic future thought involves a feeling of futurity analogous to the feeling of pastness — would, contra our argument, seem to be able to establish their own satisfaction conditions.

It is not clear, however, whether this objection poses a problem for our account, as there are reasons to think that subjects might have episodic memories that do not involve auto-noesis. Klein and Nichols (2012) report the case of patient R.B., who, after an accident resulting in head trauma, could arguably remember events from his personal past without the sense of ownership that is usually ascribed to episodic memories. For example, in describing his recollection of studying with friends at MIT earlier in his life, R.B. claims that, besides seeing “the scene in [his] head” and being “able to re-live it”, he has a “sense of being at there, at MIT, in the lounge” (Klein and Nichols 2012, 687). However, he adds, it does not feel like he owns the memory but rather like he is “imagining [or] re-living the experience [...] described by someone else” (687). As Klein and Nichols describe R.B.’s case, “he knows [the memories] are his in some sense, but he feels as though they are not ‘owned’” (2012, 688). On Klein’s and Nichols’ (2012) description, R.B. possesses information about the what, when, and where of the events in question and is also capable of “re-living” them. Moreover, he is able to recall them at will. The only thing missing is the sense of ownership of those memories. As he puts it, “I could answer any question about where I lived at different times in my life, who my friends were, where I went to school,

activities I enjoyed, etc. But none of it was ‘me’.” (Klein and Nichols 2012, 686). Thus, it seems that R.B.’s mental states resemble occurrences of episodic memory in normal subjects, except for the fact that they lack auto-noesis. This suggests that, although auto-noesis is pervasive in episodic memories in normal subjects, it is not a requirement for their occurrence. By mentioning R.B.’s case, we do not expect to adjudicate the debate, but only to show that the claim that episodic memory necessarily involves auto-noesis is not uncontroversial and, therefore, that auto-noesis does not straightforwardly imply that the silence argument is wrong.

However, even if we set this issue aside and grant that episodic memory necessarily involves auto-noesis, auto-noesis can be incorporated into our view. Consider, first, phenomenological accounts. The fact that a memory seems to the subject to be about the past does not threaten the pragmatist account. We have argued that beliefs can accompany episodic hypothetical thoughts in two ways, personally and sub-personally. While temporal reference can feature in the phenomenological character of a thought, which is defined in personal terms, it does not follow that what explains the occurrence of temporal reference must itself be a personal-level process. Thus, the claim that it seems to subjects that their thoughts are by themselves about the past is consistent with the claim that temporal reference is due to beliefs operating at the sub-personal level. This account parallels doxastic accounts of perceptual error, in which perceptual errors are viewed as a result not of perception but of the doxastic states accompanying it. On such views, subjects need not be conscious of the doxastic states that cause perceptual errors in the same way that, on our view, subjects need not be conscious of the doxastic states that determine temporal reference. Again, this is because doxastic states can operate at the sub-personal level. The upshot is that phenomenological features of our mental states including temporal reference or perceptual error can be consistently explained by sub-personal processes, which is compatible with the pragmatist account.

Consider, second, content-based accounts. On these accounts, it is possible for auto-noesis to be built into the content of episodic hypothetical thought but not necessarily reflected in its phenomenology. Content-based accounts, we claim, are not incompatible with the

pragmatist account. The pragmatist account denies that episodic hypothetical thought possesses content in isolation from other mental states, but it does not deny that it can have content when considered in relation to those mental states. Once we consider episodic hypothetical thoughts in relation to the beliefs that accompany them, they are no longer silent — i.e., they acquire content or satisfaction conditions. So, as long as the possession of satisfaction conditions, and therefore of temporal reference, is explained by this relation, the pragmatist account can accommodate the idea that episodic hypothetical thought can have content.

One final objection is that this response is inconsistent with the silence argument, the central claim of which is that episodic hypothetical thought is not representational and therefore not contentful. This apparent inconsistency can be resolved by noting that Travis's original silence argument does not rule out the possibility that perception may have content, in the sense of possessing satisfaction conditions, but only that perception itself can be assessed for accuracy. As long as we consider perception in relation to the things that ascribe satisfaction conditions to it, we can talk about it as being contentful. Analogously, the silence argument provided here does not rule out the possibility that episodic hypothetical thought may have content but only the possibility that episodic hypothetical thought itself can be assessed for accuracy. However, if we consider episodic hypothetical thought in relation to beliefs, nothing prevents it from acquiring content from those beliefs. Thus, the fact that the pragmatist account is compatible with content-based accounts of auto-noesis is not inconsistent with the silence argument.

Chapter 4

Failing to remember

Introduction

This chapter discusses the relational view of memory in relation to cases of unsuccessful remembering. I start by, first, providing a relational analysis of remembering that is compatible with the most prominent relational view in the literature, which is proposed by Debus (2008), and second, by considering how such analysis of remembering deals with unsuccessful remembering. I argue that the relational view, in its current form, cannot properly distinguish between successful remembering and different forms of unsuccessful remembering. I proceed by distinguishing between two ways in which unsuccessful remembering can happen: misremembering and confabulation. I argue that the requirement for the presence of an experiential relation, as defended by Debus (2008), is neither necessary nor sufficient for remembering. This is because there are, on the one hand, cases of remembering that do not instantiate the relevant experiential relations, and, on the other hand, cases of confabulation and misremembering that do instantiate the relevant experiential relations. For this reason, if the relational view is to be successful in accounting for unsuccessful remembering, it will need to provide extra considerations to show how it can deal with such occurrences.

Chapter 4 builds on an important issue that appeared in Chapter 3. This is the problem of how relational views of memory deal with unsuccessful or non-veridical memory. Just

as with relational views of perception, it is not entirely clear how memory errors can be accounted for in a relational framework. The decision to have a chapter focusing on this particular issue was motivated by an important gap in the memory literature. While relational accounts of perception are prominent today, relational accounts of memory have been largely ignored, and hence there have been very few discussions of them recently. The only philosopher to clearly propose a relational view in recent years is Dorothea Debus (2008). Despite her efforts, other philosophers of memory have not felt tempted by relational views. There are, as I discuss in the general introduction and also in this paper, several reasons for this. However, the problem of how relational accounts of memory deal with error is perhaps the most pressing one. Thus, while my project in this chapter is mainly critical of Debus's particular version of the relational view, the ultimate goal is not to discredit relational views, but rather to call the attention of those who might be sympathetic to relationalism to the fact that they need to engage more effectively with issues pertaining to memory errors.

A second and more instrumental reason for focusing on the problem of relational views and memory errors is because this will play a crucial role in the argument of Chapter 5. In my discussion of a hybrid view of memory, I argue that, despite the problems facing relational accounts of memory, there are good reasons to consider them seriously. Thus, in Chapter 5, I offer an account of memory that incorporates the attractive aspects of relational views while avoiding its most pressing problem, namely, the problem of error.

In terms of the overall structure of the thesis, one important thing to note in relation to Chapter 4 is that pragmatism becomes less present in my discussion of memory. This is explained by the fact that my approach to the philosophy of memory is largely exploratory here. That is, given the similarity of the issues in the perception literature and the memory literature, and given the fact that not many explicit discussions of those issues have appeared in the memory literature lately, my efforts were dedicated to using my knowledge of the perception literature to advance some debates in the memory literature. The originality of the discussions in Chapter 4 (and Chapter 5 too) is not, therefore, so much due to the proposal of sharp positive views, but rather to the establishment of meaningful theoretical connections between two different areas and the proposal of new topics for discussion in

future works.

Outline

This chapter discusses the challenges posed by unsuccessful remembering to relational accounts of memory. Section 4.2 attempts to provide a relational analysis of remembering that is compatible with one of the most prominent relational accounts in the literature. Section 4.3 discusses the problems that unsuccessful remembering raises to relational views: Section 4.3.1 focuses on cases of confabulation and Section 4.3.2 focuses on cases of misremembering. I argue that the relational analysis extracted from relational accounts cannot properly distinguish between successful remembering and these two different forms of unsuccessful remembering.

4.1 Introduction

The dispute between indirect and direct realist views of memory has occupied a central place in historical debates. Philosophers such as Locke (1975), Hume (2011), and more recently, Russell (1912), proposed indirect realist views, where in episodic or recollective memory — i.e., memories of events that we experienced previously — one is said to be only indirectly aware of the events that one remembers by being directly aware of ideas (or representations) that are about those events. In contrast, direct realist accounts, also offered by different authors, such as Reid (2000), Laird (1920), and Russell (1921), claim that recollectively or episodically remembering something is a matter of being directly aware of the past events themselves. For direct realists, there is no intermediary entity connecting one's memory to the past events themselves.

The quarrel between these two views, however, is not only about whether memory is direct or indirect. Another way to see the disagreement is suggested by some recent works related to a similar dispute in the perception literature (see, e.g., Crane 2006; Brogaard 2014; Locatelli and Wilson 2017). This literature suggests that, in the case of perception, direct and indirect realists disagree about what perception is at the most fundamental level.

For indirect realists, perception is fundamentally a matter of representing the world.¹ For direct realists, in contrast, perception is fundamentally a matter of being in certain relations to the world. For this reason, these views have been called *representationalism* and *relationalism* respectively. In a similar fashion, one can see the dispute in memory as being about whether memory is, at the most fundamental level, a matter of *representing* the past, or whether it is a matter of standing in an appropriate *relation* to it. As an attempt to provide an updated characterization that is more in line with the contemporary terminology, I shall adopt the terms used by philosophers of perception and refer to indirect realist views as *representational views of memory* and to direct realist views as *relational views of memory*.²

Despite its importance in historical debates, and despite the prominence of its counterpart in the perception literature, the dispute between representationalism and relationalism has not received a lot of attention in contemporary philosophy of memory. There are, I think, two important reasons for this. The first reason is that the causal theory of memory, which was initially proposed by Martin and Deutscher (1966), has been predominant in the last four decades. The causal theory says, among other things, that a subject counts as remembering if and only if the subject's current mental state is causally related, in an appropriate manner, to the past events themselves. The view assumes, as a result, a form of indirect realism or representationalism about memory. That is, it requires that subjects

¹It is important to note that, despite believing that perception is, indeed, a form of representing the world, most representationalists do not see themselves as indirect realists. Searle (1983), for example, offers a representational theory of perception that, according to him, is not incompatible with more traditional forms of direct realism. For him, despite representing the world when we perceive it, we are not aware of our mental states, but of the objects out there in the world (see also Harman 1990). Representations, on this view, are not the objects of perceptual awareness, but they are the means by which we become aware of external objects. Bernecker (2008) offers a similar argument in the case of memory. Since the claim that representationalism is a form of direct realism has been the topic of controversy (see Locatelli and Wilson 2017; Travis 2017; Martin 2017), I shall not go into any more details here as it will not influence my argument.

²The option for the contemporary terminology is due to three different reasons here. The first is that it helps avoiding confusion about what representational views are in the case of memory. Indirect realism, at least in the twentieth century, has often been associated with sense-data theories (see, e.g., Russell 1912; Broad 1925; Ayer 1956), which are reminiscent of early modern empiricism. In the contemporary debate on memory, however, those who take themselves to be representationalists are not necessarily committed to such views, although they will sometimes accept the label of indirect realists. The second reason is that Dorothea Debus (2008), who is the most prominent contemporary defender of direct realism and whose work will be discussed later, calls her own view "relational" in allusion to Martin's (2004) relational view of perception. And the third reason, which is more general in nature, is that I do not want to suggest that there are any deeper philosophical connections between representational and relational views and indirect and direct realist views other than one of historical kinship.

be capable of having mental representations that are sufficiently similar³ and are causally related to the past events represented in memory. Consequently, these representations serve as the intermediary entities connecting current memories to the past events themselves.

The second reason is that a form of representationalism seems to be better supported by empirical research on memory. For example, it has been suggested by a number of different researchers that memory is a result of a more general cognitive capacity to think about events, whether or not those events occurred (see, e.g., Suddendorf and Corballis 2007; Schacter et al. 2007, 2012; De Brigard 2014a; Michaelian 2016c). Thus, because relationalism seems to require the existence of the objects of memory — and hence of other forms of episodic thinking — the view simply appears unappealing to most people when considered in relation to these results. Another issue has to do with the fact that it is now widely known that, rather than simply retrieving information from the past, memory is constructive, in the sense that not all the information that one remembers needs to be originated in one's past experience of the relevant event (see, e.g., Schacter et al. 2007; Michaelian 2011; Robins 2016b; Cheng and Werning 2016). This suggests that memory need not always be able to reach back to past experiences or events, but rather construct actual representations of what happened. For this reason, a form of representationalism seems to be the most simple and intuitive option, as representations can be constructed by the brain in the absence of the represented objects. A third issue is that representationalism seems to accommodate memory errors better. By memory errors, I mean occurrences of *misremembering* (see, e.g., Robins 2016a; Michaelian 2016b), such as when one remembers some details of an event incorrectly — e.g., remembering having strawberry cake at your birthday party, when you had chocolate cake — and occurrences of *confabulation* (see, e.g., Hirstein 2005; Michaelian 2016b; Bernecker 2017; Robins 2017a), such as when one remembers events that did not happen — e.g., remembering having a birthday party when

³Martin and Deutscher (1966) conceive of this similarity in terms of a structural analogy holding between the past representation and the current representation. They say that “the past experience must constitute a *structural analogue* of the thing remembered, to the extent to which he can accurately represent the thing” (191, my emphasis). It is not entirely clear, however, where the structural analogy is to be found. The most natural interpretation seems to be that the content of the past representation must have the same kind of structure as the content of the current representation, but they do not say anything as to what the structure of those contents are supposed to be. Since this point will not play a major role in my discussion, I will simply take this interpretation for granted here.

there was no such party. Since representations can *mis*-represent existing things, or even represent things that do not exist at all, they provide a good theoretical resource to explain memory errors.

Despite the apparent advantage of representationalism, the prospects of relational views in relation to empirical research on memory have not been addressed in enough detail. This is, in part, due to the fact that there are not many relational views out there, and also because those who have proposed relational views have not addressed this topic in detail. Consider, for example, the problem of memory errors and the relational view offered by Debus (2008), who is arguably the most prominent contemporary defender of relationalism. According to Debus, one episodically or recollectively remembers something only if one stands in an *experiential relation* to the relevant past events. This experiential relation, she says, is to be understood in terms of the causal, temporal, and spatial relationships holding between subjects and the past events. Although she offers one important necessary criterion to say that a subject is remembering, Debus is not very clear on what a complete relationalist analysis of remembering looks like. In other words, unlike the classical analysis provided by Martin and Deutscher (1966), she does not give us a clear set of necessary and sufficient criteria for remembering. As a result, it is not clear when, for the relationalist, a subject fails to remember, and, therefore, when he misremembers or confabulates.

In light of the recent re-emergence of relationalism in the contemporary philosophical debate and of the growing empirical literature on memory errors, this paper has two goals: the first is to discuss what a relational analysis of remembering that takes Debus's (2008) relational view as a starting point would look like (Section 4.2); and the second is to see whether such relational analysis of remembering can accommodate the occurrence of unsuccessful remembering or memory errors (Section 4.3). As my positive contribution to this topic, I shall argue that the relational view of memory faces serious problems to accommodate unsuccessful remembering. In particular, the argument is that the presence of an experiential relationship alone cannot distinguish properly between successful and unsuccessful occurrences of remembering. I should note that, while some might view my proposal as being an implicit endorsement of traditional representationalism, I shall not take

a firm stand on the subject here. Recent work on the dispute between representationalism and relationalism in philosophy of perception suggests that they are not mutually exclusive views and hence that “hybrid” views are possible (see Schellenberg 2010, 2011; Siegel 2010; Logue 2014; Hanna 2015; Sant’Anna 2018b). Despite there being no such attempts of developing hybrid views in the case of memory, it has been already pointed out by some that this is a potential topic of interest (see Michaelian and Sutton 2017). For this reason, even if one finds relationalism to be problematic, it does not mean that one is necessarily committed to traditional versions of representationalism. As it stands, the argument that I will develop here leaves open the question of whether some form of relationalism can be preserved by reconciling it with some form of representationalism.

4.2 The relational analysis of remembering

The relational view developed by Debus (2008) says that what is characteristic of remembering is that subjects stand in an *experiential relation* to the past events. The term “experiential relation” is a technical one, and since it is central to her account, it will be helpful to examine it more carefully. As Debus points out at the beginning of her paper, she is concerned with *recollective memories*, which, according to her, are memories that have experiential characteristics. Recollective memories are, roughly speaking, equivalent to what psychologists call today *episodic memories* (see Tulving 1972, 1985a). In other words, they are memories of particular events that one experienced in the past, such that, when one remembers, it seems to one that one is “re-experiencing” or “re-living” the relevant event (Tulving 2002). Thus, those memories have similar characteristics to sensory perception, in the sense that we seem to “see”, “hear”, “smell”, and so on, the relevant features associated with the original event, but now “inside” our minds.

Being able to “re-experience” or “re-live” an event requires, in turn, that one experienced the relevant event in the past. It is important to note that, when one remembers, one is not literally having another perceptual experience of the event, but is rather undergoing a different kind of mental state that has a distinctive relationship to the past event, namely,

a specific kind of experiential relationship. So, while it is true that both perception and memory have experiential characteristics, and while it is arguably the case that both put us in an experiential relation to their objects, the relevant experiential relation instantiated in each case is different. For Debus, in the case of memory, the experiential relation is characterized as one that supervenes on specific causal, temporal, and spatial relations that hold between subjects and the past events. The relation is *causal* because, as she points out, when one remembers, certain neurophysiological events occur in one's brain. And these events are, accordingly, caused by earlier neurophysiological process which were, in their turn, ultimately caused by the past events at the time of the original experience (Debus 2008, 411). Similarly, the relation supervenes on *spatial* and *temporal* relations because, as I mentioned initially, remembering requires that the subject experienced the relevant event before. That is, the original experience must happen earlier than the memory and there must be a causal connection with the original experience. Debus frames this in terms of the original experience and the memory being in the same "spatiotemporal path" traced by the subject in the world. That is, according to her, by inhabiting the world, subjects trace a certain "spatiotemporal path", such that, to count as remembering something, the spatiotemporal point in which the memory occurs must precede the spatiotemporal point in which the experience occurred, and both must belong to the same "spatiotemporal path".

For now, I shall not dispute whether there is such experiential relation or whether it is necessary for recollective memory. Instead, I will take both these things for granted. My aim in this section is, rather, to provide a relational analysis of remembering that is compatible with Debus's relational view. While there are some indications of what such an analysis might look like in her discussion, she does not say this quite explicitly. Thus, as a first step towards providing a relational analysis, I will take the traditional analysis of remembering provided by Martin and Deutscher (1966) as a starting point and then modify it to fit with the relational view. For Martin and Deutscher, a subject *S* will count as remembering an event *e* iff:

- (1) *S* has previously experienced *e*;
- (2) *S* has a present mental representation of *e*;

- (3) *S* stands in an appropriate causal relationship to *e*;
- (4) The content of *S*'s present representation is sufficiently similar to the content of *S*'s previous experience.⁴

To clarify, consider my putative memory *M* of my tenth birthday party. According to Martin and Deutscher, *M* is an instance of remembering because: (1) I had a *previous* experience of my tenth birthday party, which led me to form representations of this event; (2) I have a *current* representation *M* of my tenth birthday party; (3) my actual representation *M* is caused, perhaps by means of a memory trace, by my past experience of the party; and (4) the content of *M* is sufficiently similar to the content of the past experience — e.g., I remember my parents being there, I remember having chocolate cake, and so on, which were all represented in my previous perceptual experience of the party. (1)–(4) are, therefore, jointly necessary and sufficient conditions to say that a subject count as remembering.

Now, can (1)–(4) be incorporated into a relational analysis of remembering? Let's consider (1) first. Since Debus says quite explicitly that successful remembering requires that subjects have experienced the relevant events in the past, and this is central to make sense of the notion of an experiential relation, it seems that (1) can be straightforwardly incorporated into the relational analysis. The same is true of (3). That is, the experiential relation is said to supervene on causal and spatiotemporal relations holding between subjects and events. However, (2) and (4) do not seem compatible with the relational view. Debus is not very explicit about whether or not she thinks that remembering can involve some sort of representation. However, there are good (implicit) reasons to think that it does not. The first reason is that, later on in the paper, Debus says that the experiential relation is supposed to make subjects *directly* aware of the past events. So, when one remembers, it is not the case that one's awareness of the event is mediated, which eliminates the need for representations.⁵ The second reason is that Debus frames her view in relation to Martin's (2004) relational view of perception, which is a form of direct realism that avoids repre-

⁴This formulation is adapted from Bernecker (2010, ch. 1). See also Bernecker (2015, 302).

⁵Although see Bernecker (2008, ch. 5), who claims that the presence of representations in memory is not incompatible with a form of direct realism. I will leave this issue aside as it is not clear whether Bernecker's view is a genuine form of direct realism.

sentations altogether. While Debus is careful to make explicit that her understanding of mnemonic awareness differs in important respects from what is usually understood by perceptual awareness, it looks like that the same, or at least very similar, motivations drive both the relational view of memory and relational views of perception. While this does not make a conclusive case against the presence of representations, it gives us good reasons to cast doubt on the inclusion of (2) and (4) into the relational analysis.

Thus, by looking at the analysis proposed by Martin and Deutscher (1966), we can use (1) and (3) as starting points to conceive of a relational analysis of remembering. Considered alone, however, (1) and (3) provide only a partial picture of a relational analysis. The reason is that, for the relational view, the relationship that subjects have to past events is not merely causal. In order to provide a more complete picture, the relational analysis needs to require that subjects stand in an *experiential* relation to the relevant events, which takes into account not only a causal relation, but spatial and temporal relations as well. On this view, then, the relational view of memory offered by Debus seems to suggest that *S* remembers *e* iff:

(R1) *S* has previously experienced *e*;

(R2) *S* is presently experientially related to *e* by undergoing the mental state *M*; where

(R2.1) The experiential relation supervenes on specific causal, spatial, and temporal relations obtaining between *S* and *e*.

Thus, to use the same example as above, the relational view says that I remember my tenth birthday party because (R1) I had a previous perceptual experience of the party, and (R2) I am experientially related to the party in the sense that my memory is caused by it and the party is located at a spatiotemporal point that precedes the actual spatiotemporal point in which I have the memory. On this view, then, (R1) and (R2) are viewed as jointly necessary and sufficient conditions to say that a subject count as remembering.

Before moving forward, one important clarification refers to the nature of the mental state *M* in (R2). Since, on the relational view, *M* cannot be understood as a form of representation, one natural question would be what *M* is according to the relational view. The

traditional move by relationalists in perception is to understand M (in the case of perception, its equivalent P) as being a non-representational form of awareness of particular objects in the external world, in which those objects are constitutive parts of the perceptual states (see Martin 2004; Fish 2009). Similarly, in the case of memory, one might understand M as a non-representational form of awareness of the past, in which the past events are constitutive parts of the mnemonic states. This seems to be what Debus has in mind when she talks about memory making us directly aware of the past, although, I should note again, she is careful to make explicit that the kind of awareness involved in memory differs in important senses from the kind of awareness present in perception (see Debus 2008, 408).

The relational analysis stated above is intended as a response to the general question of what is it to remember, as opposed to, e.g., the questions of what is it to perceive or what is it to think. However, in providing an analysis of a given phenomenon, one might ask not only whether the analysis provides a good way of distinguishing the target phenomenon from other phenomena, but also whether it distinguishes between different ways in which the same phenomenon can happen. In the case of memory, in particular, it is possible to distinguish between *successful* and *unsuccessful* occurrences of remembering. Successful remembering refers to, roughly speaking, cases where we remember things as they happened. Unsuccessful remembering, in contrast, can take two different forms. The first is *misremembering* (Robins 2016a), which corresponds to cases where we get some details of an event right, but get some other details of an event wrong.⁶ The second is *confabulation* (Hirstein 2005), which corresponds to cases where we get all the details wrong — i.e., the event remembered did not happen. Since the presence of an experiential relation is central to the relational analysis, one would expect that the presence or the absence of such relation would provide a meaningful way to distinguish successful and unsuccessful remembering. In fact, this seems to be the suggestion behind Debus's (2008) overall proposal. However,

⁶I should say that this is not necessarily what Robins understands by misremembering. For Robins, “[m]isremembering is a memory error that relies on successful retention of the targeted event” (2016a, 433), which is paradigmatically illustrated by the DRM effect, where the content of what is misremembered is non-trivially related to the content of what is remembered or retrieved — e.g., subjects are more likely to misremember seeing particular fruits, such as bananas, if the content of what is remembered or retrieved also involves fruits, such as apples and oranges. The definition of misremembering that I am using here does not require such a non-trivial relation. According to this definition, one misremembers if the content of what is misremembered was not originally experienced.

it is not entirely clear whether this will work. In the next section, I will argue that (R2), or the requirement for the presence of an experiential relation, is neither necessary nor sufficient to distinguish successful and unsuccessful remembering, and thus that the relational analysis, as formulated in here, is at best incomplete.

4.3 Unsuccessful remembering

The problem of unsuccessful remembering and its relationship to the relational view of memory view has not been explored in much detail. In the perception literature, however, forms of what we might call unsuccessful perception — i.e., illusions and hallucinations — have proven to pose real challenges to relational or direct realist views of perception (see Smith 2002; Crane and French 2017). The main reason for this is that successful or veridical occurrences of perception can be indistinguishable, from the point of view of the subject, from unsuccessful or non-veridical ones. Thus, when relationalists or direct realists say that perception requires the presence of a relation to the objects that are perceived, they need to explain how and why there can be phenomenologically indistinguishable perceptual experiences either in the presence or in the absence of the objects perceived.

The same problem, I think, arises for relational or direct realist views of memory. Successful and unsuccessful remembering can often be phenomenologically indistinguishable, which questions the necessity of there being an experiential relation to the past events when one remembers. In the same way that I can see a spider on the wall and fail to realize that there is not a spider on the wall, I can remember an event, e.g., my tenth birthday party, and fail to realize that this event did not happen. However, if I can seemingly remember an event when it did not occur, and hence undergo the same mental state as the one I would undergo had the event been the case, then it does not seem to make sense to postulate a relationship to such event as necessary for remembering. That is, the fact that one can seemingly remember an event when the event did not happen “screens off” (Martin 2004) any explanatory role played by the past events in one’s theory of memory.

One strategy popular among relationalists about perception is to resort to a disjunctivist

view of perception. There are different varieties of disjunctivism (see Byrne and Logue 2008), but one popular characterization is that, at the most fundamental level, successful or veridical and unsuccessful or non-veridical perception do not share any essential features. In other words, they are viewed as occurrences of two different kinds (see Snowdon 1980; Martin 2004; Fish 2009). Thus, if disjunctivism is true, it is no longer a problem to say that veridical perception is relational, while non-veridical perception is not. In her discussion of “memory hallucinations”, which I have called confabulations here, Debus (2008) employs the same strategy. Inspired by Martin’s (2004) disjunctivist account of perception, she says that the fact that two mental occurrences are indiscriminable from a subjective point of view does not imply that they are occurrences of the same kind. For Debus, the absence of the experiential relation in the case of the so-called memory hallucinations postulate a fundamental separation between them and successful remembering (2008, 414). This explains why, despite being phenomenologically indistinguishable, successful memories are relational and memory hallucinations, or confabulations, are not.

Thus, if combined with disjunctivism, the relational analysis can provide a principled way to distinguish successful remembering from at least one kind of unsuccessful remembering, namely, what Debus calls “memory hallucinations”, or what I have called “confabulations”. However, I think that, even with disjunctivism, the relational analysis fails to provide a satisfactory way to distinguish between successful and unsuccessful remembering. The main reason, as I will argue in section 4.3.1, is that memory hallucinations or confabulations do satisfy (R2) above, so there is no principled way to say that they are fundamentally different from successful remembering. But before I turn to this question, let me define more precisely what I will understand by unsuccessful remembering. As I said before, unsuccessful remembering can take two forms. The first form, *misremembering*, is when we remember incorrectly one or more details of an event that occurred.⁷

The second form, *confabulation*, is when we remember an event that did not occur at all. To illustrate, consider misremembering first. Consider, for instance, my memory of my

⁷For more discussions about misremembering and its relationship to current theories of remembering, see Robins (2016a) and Michaelian (2016b). For discussions about confabulation, see Hirstein (2005), and for discussions about the relationship of confabulation to current theories of remembering, see Michaelian (2016b), Bernecker (2017), and Robins (2017a).

tenth birthday party. I remember my parents being there, I remember it being a sunny day, and I remember having chocolate cake. While this memory is about an event that happened, at least one of its elements incorrectly describe the event; that is, it was a rainy rather than a sunny day. So, despite being about an event that happened, and despite getting some details right — i.e., my parents were there and I had chocolate cake — I *mis*-remember the event in question. Consider, now, confabulations. Suppose that I remember going to Paris when I was a teenager. Since I have never been to France, this memory will count as an occurrence of confabulation because the event in question did not happen.

Now that we have a clearer understanding of what unsuccessful remembering is, I will move on to show why it poses a problem to the relational view of memory. Section 4.3.1 will discuss confabulations and Section 4.3.2 will discuss misremembering. My main argument will be that both forms of unsuccessful remembering instantiate the relevant experiential relation to past events and thus that the relational analysis of remembering fails to properly distinguish them from successful remembering.

4.3.1 The problem of confabulation

Can the relational analysis of remembering rule out cases of confabulations? To answer this question, we should go back to (R1) and (R2) above. For the relational view, *S* remembers *e* iff:

(R1) *S* has previously experienced *e*;

(R2) *S* is presently experientially related to *e* by undergoing the mental state *M*; where

(R2.1) The experiential relation supervenes on specific causal, spatial, and temporal relations obtaining between *S* and *e*.

The question is, then, whether there are cases of confabulation that satisfy (R1) and (R2). The most obvious way to rule out cases of confabulation would be to appeal to (R1). Since the events that one confabulates have not happened, it is not possible that one experienced those events in the past. Therefore, confabulations do not satisfy (R1) and are thus distinct

from occurrences of successful remembering. While I think this provides a satisfactory way to deal with cases of confabulations, it does not seem to reflect the motivation behind the relational view. When she introduces disjunctivism to deal with confabulations, Debus argues that confabulations are distinct in kind from successful remembering because, in the former case, the relevant experiential relationship fails to obtain. So, at least on Debus's version of the relational view, it is the presence of (R2), and not of (R1), that provides the means to distinguish confabulations from successful remembering.

I will come back to (R1) later on. For now, let us focus on (R2), as it is the central element of Debus's view. The question now is, therefore, whether (R2) can provide a meaningful way to distinguish successful occurrences of remembering from confabulations. To start with, it is important to note that (R2) will be successful in providing such distinction only if we make the further assumption that, to successfully remember, subjects must be in a *unique* experiential relation to a *unique* past event. That being the case, the claim that confabulations do not count as successful remembering becomes straightforwardly true, for there is no unique experiential relation holding between the subject and the confabulated event, as the event in question did not happen. However, the relational view does not require us to make that assumption. The claim is simply that successful remembering instantiates the relevant experiential relation, while confabulations do not. Thus, if one does not make this assumption, it is not hard to envisage cases where confabulations instantiate experiential relations to past events in the way described by (R2).

To make this clear, let me introduce the notion of a *partial experiential relation*.⁸ A partial experiential relation is similar to the one describe in (R2), that is, it supervenes on the causal, spatial, and temporal relations between a subject *S* and a past event *e*. However, because of its partiality, the partial experiential relation only makes the subject aware of *some* but not all the constituents, or properties, of *e*. For example, I can be in a partial experiential relation to my tenth birthday party, such that I can remember having chocolate

⁸It is not completely clear whether Debus (2008) is committed to the claim that remembering can only involve full experiential relations, as opposed to partial experiential relations, as I will discuss here. However, as it will become clear later on, accepting that remembering can involve partial experiential relations is the main source of problems for the relational view. For this reason, I will assume that Debus is at least implicitly committed to the idea that remembering requires the presence of a full experiential relation.

cake next to my parents, but fail to remember other elements, including that this memory is of my tenth birthday party. Suppose, now, that I remember an event that did not happen. Because the event did not happen, the relational view says that there cannot possibly be an experiential relation between the subject and the event, thus making this an occurrence of confabulation. But, it does not follow from this that the confabulation fails to put the subject in partial experiential relations to past events. It could be the case — and I will argue later that there are good reasons to believe that this is actually the case — that the confabulation puts the subject in two or more partial experiential relations to two or more past events, such that he becomes aware of different constituents or features of multiple events as belonging to a unique event. Consider, for instance, my memory of having pizza at the beach. While I have never experienced such an event, it is true that I have previously experienced the following two different events: one in which I had pizza and another one in which I was at the beach. So, if the notion of a partial experiential relation is sound, there is no reason why I could not be aware of the confabulated event by being partially experientially related to two or more events that I experienced previously. However, if that is the case, then nothing prevents confabulations from putting us in experiential relations to past events.

One way for relationalists to avoid this problem would be to simply bite the bullet and add the additional clause that a mental state counts as successful remembering only if it makes the subject experientially related to a *unique* event. However, I think this move should be resisted. If correct, it would lead to a restrictive, and hence counterintuitive, account of remembering. To see this, suppose that I have a putative memory *M* of my tenth birthday party, where I remember having chocolate cake with my parents and my friends. Since this event happened, and since I experienced it, (R1) is satisfied. However, *M* is not brought about by my being experientially related to the event of my tenth birthday party, but instead, by my being partially experientially related to different events (see Michaelian 2011; Robins 2016b). In this particular case, *M* puts me in a partial experiential relation to an event e_1 , where I had chocolate cake, to an event e_2 , where I had a meal with my parents, and to an event e_3 , where I was hanging out with my friends. In other words, my awareness

of the three distinctive constituents of M, namely, <having chocolate cake>, <having a meal with my parents>, <being with my friends>, are awareness of constituents of *different* past events that are brought together by M.

If this is right, the claim that remembering requires a unique experiential relation will require us to say that I am *not* remembering my tenth birthday party. But this is problematic for at least two reasons. The first is that M above satisfies (R1) and, moreover, it is experientially related, in a relevant sense, to past experiences. From a common sensical point of view, this seems to be enough to say that one is successfully remembering, as one gets the relevant details right and one does so in virtue of being (partially) experientially related to the past. The only thing that is missing, arguably, is that the event that M purportedly makes me aware of, namely, my tenth birthday party, is not operative — i.e., it is not the ultimate cause — in producing M (see Martin and Deutscher 1966, p. 166). Thus, it might be argued that, despite seeming so, M does not count as a successful occurrence of remembering because none of the experiential relations trace back to the original event, which is precisely what is required to say that M allows me to be aware of my tenth birthday party.

This answer will, however, provide only a temporary solution. The reason is that, while the relational view might be able to rule out cases where none of the partial experiential relations trace back to the past event itself, it will not be able to rule out cases where at least one of the partial experiential relations establish such relation. Consider the case above again. Suppose that my partial awareness of <having chocolate cake> is, in fact, awareness of myself having chocolate cake in my tenth birthday party. However, my awareness of <having a meal with my parents> and of <being with my friends> are partial awareness of constituents of different events, i.e., e_2 and e_3 . In this case, then, there is at least one experiential relation that traces back to the relevant past event, and thus M should count as a successful occurrence of remembering even if there is more than one experiential relation connecting M to multiple past events. Requiring the presence of a unique experiential relation is, therefore, too restrictive, which provides a counterintuitive account of memory.

So far, I have attempted to show that, from an intuitive point of view, requiring a unique experiential relation is implausible. However, relationalists might resist this by saying that

they do not share our intuitions. To put it sharply, the claim here would be that, for relationalism, if there is not a unique experiential relation holding between the subject and the event remembered, then, intuitively, the subject's putative memory will not be an occurrence of successful remembering. This leads us to the second problem with the idea that remembering requires a unique experiential relation. While I think that this response is, in principle, open to relationalists, I do not think it is a promising one. The reason is that it provides an analysis of remembering that starts by stipulating what the phenomenon must be, instead of describing or explaining what the phenomenon actually is. The problem with this approach is that our theoretical intuitions do not always correspond to how the phenomenon actually is, which runs the risk of making our initial stipulations empty or explanatorily uninteresting.

And this, I think, is the problem with this strategy. In other words, it has been shown by empirical research that memory is *constructive*, in the sense that our memories are constructed from different sources other than the original experience. Some people, such as Michaelian (2011), have suggested that while some causal link must be preserved to the past experience that the memory is about, not all the elements or constituents that figure into that memory need to be derived from the original experience. Such elements can come from different sources, such as semantic information (Cheng and Werning 2016; Cheng et al. 2016) or background knowledge (Suddendorf and Corballis 1997) present in retrieval (see also De Brigard 2014a and Michaelian 2016c). Based on this and other empirical findings, De Brigard (2014a) has suggested, for example, that it is not unsurprising that most of the occurrences of remembering that we call "successful" will have elements that are not drawn from the original experience. Setting the details of this debate aside, the point is that how memory works, or how remembering happens in the world, does not seem to support the prior relationalist theoretical intuition that a unique experiential relation to the past event must be characteristic of successful remembering. If that was the case, most of the occurrences of remembering that we call "successful" would be actually unsuccessful. So, again, appealing to the presence of a unique experiential relation provides a restrictive, and hence inadequate, account of how memory works.

I have argued that (R2) does not provide a sufficient criterion to distinguish successful remembering from one form of unsuccessful remembering, i.e., confabulations, as is suggested by Debus's relational view. Before I move on to consider misremembering, let me come back to the relationship between (R1) and confabulations. As I said initially, I think that relationalists can satisfactorily rule out confabulations by appealing to (R1), as subjects have not experienced the relevant events. The argument I gave so far only provides reasons to deny that (R2) is neither necessary nor sufficient to draw the relevant distinction between successful remembering and confabulation. The problem with placing the burden of the distinction on (R1) is that it undermines the whole motivation for providing a relational account of memory. In Debus's (2008) version of the view, (R2) plays the explanatory role of distinguishing successful remembering from unsuccessful remembering. This is suggested, among other things, by the radical move that Debus makes to disjunctivism. But if, as I argued, (R2) is not required to draw the distinction, the question that poses itself is why give it a central place in one's account of memory, so as to call it a *relational* view of memory, when, in reality, it is only playing a marginal role in one's analysis of remembering.

But even if we set this worry aside, relationalism would still provide an incomplete account of confabulations. I have assumed throughout my discussion that confabulations are falsidical occurrences of remembering, that is, that they are cases where subjects get things wrong. However, Michaelian (2016b) has recently pointed out to the possibility of *veridical confabulations*, that is, cases where subjects get things right by mere accident. The most intuitive analogy here, as Michaelian himself makes explicit, is to veridical perceptual hallucinations (see Lewis 1980). Veridical hallucinations refer to cases where subjects perceive the environment accurately, but the cause of their perceptual experiences is not the environment, but something else, such as evil scientists or some misfiring happening at the neuronal level. Similarly, veridical confabulations are cases where subjects remember the past correctly, but this is not due to their memories being adequately related to the past, but due to something else, such as the malfunctioning of the underlying mechanisms responsible for memory, or to the occurrence of processes not ordinarily associated with remembering, such as guessing. Assuming that veridical confabulations are possible, and

hence that a complete analysis of remembering must be able to distinguish them from falsidical confabulations, simply appealing to (R1) would not provide such a distinction. Both veridical and falsidical confabulations fail to satisfy (R1); and, for this reason, the relational view would mistakenly place these two different occurrences of confabulations under the same category.

In summary, I do not think that the above provides a definitive case against pursuing relational views. However, I think that it gives us good reasons to think that appealing solely to the notion of an experiential relation might not be the best way to go.

Before I move forward, I should address one objection that is likely to arise in the context of the introduction of partial experiential relations. Some might argue that, even if there is such a thing as a partial experiential relation, it cannot be used to characterize confabulations appropriately. In particular, the claim is that partial experiential relations cannot explain the phenomenology of confabulation. Like successful remembering, confabulations seemingly make us aware of unique events; however, if they can be partially related to multiple events, as I have suggested, it would follow that it is impossible for us to experience confabulations as involving awareness of unique events. Instead, it would be more plausible to say that they make us aware of a conjunction of two or more events. This would conflict, however, with how we experience confabulations.

This worry can be resolved by noticing that, in partial experiential relations, we become aware of parts or properties of events (Sant'Anna 2018a; Chapter 5, 171). For example, the event of my tenth birthday party has different properties, such as “having chocolate cake”, “having a meal with my family”, and so on. Being fully or wholly experientially related to this event would entail that I would be aware of *all* the properties instantiated by this event. In contrast, in cases where I am only partially experientially related to this event, I am aware of one or more (but not all) of its properties, e.g., “I remember having chocolate cake”. Being aware of this particular property does not, however, make me automatically aware of the event as a whole. So, when confabulations put a subject in multiple partial experiential relations to different events, it is not required that the subjects experience the confabulations as being composed by multiple events. Instead, they are said to be aware

of properties of multiple events which are somehow brought together to consciousness. In a representationalist framework, the equivalent claim here would be that different bits of representational content are brought together to form a new representation of a unique event.

In addition, rather than being incapable of explaining the phenomenology of confabulation, one might argue that, if there is no such thing as a partial experiential relation, relationalists cannot offer a positive account of the phenomenology of confabulation. Because, in confabulations, we are not fully or wholly aware of any events, it would not be possible for us to experience those mental states as making us aware of events at all. One might, of course, retreat to a negative form of disjunctivism here and simply claim that it is not the relationalist's business to explain confabulations, as those are different in kind from successful remembering. However, if discussions about perceptual error serve as any guide, this strategy is unlikely to be satisfactory. Negative disjunctivist accounts have been met with skepticism and some disjunctivists have recently made efforts to provide positive accounts of error (Fish 2009; Brewer 2011). Thus, unless relationalists are prepared to offer a positive account of confabulation that does not involve partial experiential relations, the prospects of relationalism will be even dimmer if it does not adopt the notion of partial experiential relations.

4.3.2 The problem of misremembering

I have argued that the relational analysis of remembering faces trouble to distinguish successful remembering from the first form of unsuccessful remembering, which I have called confabulations. I will now consider the second form of unsuccessful remembering, misremembering, which poses an even more pressing problem to the relational view. The argument proposed in this section is somewhat independent from the argument proposed in 4.3.1. I will argue that, even if we put the problems pertaining to confabulations aside, the relational view still fails to distinguish between successful remembering and misremembering.

As I discussed before, in cases of misremembering, one is aware of an event that one

experienced, but one gets some details of the event wrong. Consider the birthday party example above. Suppose that, instead of remembering having chocolate cake, I remember having strawberry cake. I get all the other details about the event right: I remember my parents and friends being there, I remember it being a rainy day, and etc., but I remember the flavor of my birthday cake wrong. In such case, arguably, I am aware of an event that happened, but remember one detail wrong — i.e., the flavor of my birthday cake.

How does the relational view distinguish successful remembering from misremembering? To approach this question, we can employ the same strategy used in the case of confabulations, that is, we can ask whether there are cases of misremembering that satisfy (R1) and (R2). Since the events that are objects of misremembering have been experienced in the past, it seems uncontroversial that (R1) is satisfied. However, and this is where the problem that misremembering poses to the relational view comes to the surface, because those events exist and were experienced, it looks like that (R2) is also satisfied, for subjects stand in the right kind of experiential relation to the events. But, if that is the case, then occurrences of misremembering satisfy the criteria to be classified as successful remembering, which is implausible from the perspective of the relational view.

There is one obvious way for relationalists to avoid this problem. This consists in denying that, when I remember having strawberry cake at my tenth birthday party, I am misremembering a past event that actually happened. Instead, the suggestion is that, rightly because I remember myself having strawberry cake, and not chocolate cake, I am *not* remembering my tenth birthday party, but some *other* event that did not happen. Thus, because the event in question did not happen, and hence was not experienced by me, the putative memory does not qualify as an occurrence of successful remembering because the relevant experiential relation fails to obtain. In other words, my putative memory fails to satisfy both (R1) and (R2).

This solution faces three important problems. The first problem is that it will sound counterintuitive to some. As I discussed before, it has been suggested by different researchers that it is not uncommon for our memories to have elements that are not derived from a unique past experience (see, e.g., Michaelian 2011; De Brigard 2014a; Cheng and

Werning 2016). Thus, it is not unsurprising that those memories will often be composed by elements that were not present in the original experience. However, it seems too strong to conclude from this that we are systematically failing to remember the events in question. This would mean that, whenever our memories get some detail wrong, which, as De Brigard (2014a) points out, happens quite often, the brain mechanisms responsible for memory would be malfunctioning. However, a system that is systematically malfunctioning is problematic from an evolutionary point of view, as one would need to explain why it has been preserved despite failing to realize its function and why it appears to be useful, as seems to be the case with memory. So, to avoid these worries, it seems more reasonable to say that, in misremembering, we are related to the events that happened, but for some reason or another, we get one or more of the details wrong. This would allow us to say that, overall, the system is functioning as it should, while still allowing for error to be present in its outputs, i.e., in the memories produced.

The second problem this solution faces relates to the issues raised in 4.3.1. We could, in response, point out that, while it might be true that occurrences of misremembering fail to instantiate unique experiential relations to events, nothing would prevent them from instantiating multiple partial experiential relations to different events. Thus, relationalists would need to deal with the same problems raised in the context of confabulations. In other words, they would need to explain how it is possible for unsuccessful remembering to instantiate experiential relations without making the relational view trivial. And finally, the third problem is that, even if we set these worries aside, and grant relationalists that, when I misremember, my putative memory relates to an event that is different from my tenth birthday party, the relational view would still provide an incomplete analysis of remembering. The reason for this is that it would conflate confabulations and misremembering, thus failing to make sense of the differences between these two forms of unsuccessful remembering. In other words, while the relational view might be able to distinguish between successful from unsuccessful remembering, it will not be able to distinguish between forms of unsuccessful remembering because both misremembering and confabulations are understood in terms of the absence of the relevant experiential relation. Therefore, the relational view provides

an inappropriate taxonomy of memory errors, and, consequently, an incomplete analysis of remembering.

In response to the third problem, relationalists might want to deny that there is a real distinction between forms of unsuccessful remembering. The claim here would be that there is no real difference between getting some and all of the details of a given event wrong; in other words, the same process is taking place in both cases, i.e., the subject is ultimately *unsuccessfully* remembering events. Alternatively put, what I have been called misremembering so far is nothing but a less extreme case of confabulation. Thus, if this is right, relationalism is not obliged to provide a taxonomy of memory errors along the lines proposed here.

This move, I think, faces similar problems to the move discussed in 4.3.1 that attempts to deny that, by definition, there cannot be occurrences of successful remembering that instantiate partial experiential relations. In other words, it runs the risk of forcing our theoretical intuitions into our analysis of the phenomenon in question. While it is not logically incoherent to say that misremembering is nothing but a form of confabulation, this does not seem to reflect how misremembering and confabulations actually happen in the world. As De Brigard (2014a) points out, occurrences of what I have called misremembering should not be seen as “bad” products of our memory systems; instead, they are the expected outcomes given how those systems work. Confabulations, in contrast, occur in situations where something has clearly gone wrong, such as when subjects suffer head injuries that lead to some sort of brain damage (see Hirstein 2005 and Bernecker 2017 for more discussion). As Michaelian (2016b) points out, while both misremembering and confabulations share the feature of being both erroneous — in his terminology, they are both “inaccurate” — unlike misremembering, which is produced by mechanisms that are working in a reliable manner, confabulations are produced by mechanisms that are working unreliably (see also Hirstein 2005 for a similar account of confabulations). So, it looks like that the workings of the systems responsible for misremembering and confabulation differ in important senses.

Another problem, which is epistemological in nature, is that misremembering, but not confabulation, seems to put us in a position to know at least some things about the past in

a reliable manner. In her recent account of misremembering, Robins (2016a) notes that, “[w]hereas misrememberings result from the distortion of retained information, confabulations are wholly inaccurate, reflecting no influence of retained information from a particular past event” (434). In other words, despite involving distortion, misremembering can still provide us with some useful information about the relevant events because there is some retention of information from the past, which is not the case with confabulations. Thus, as Robins puts it, “[t]he distinction between misremembering and confabulation is particularly important” (434).

In summary, the main point is that, if relationalists want to pursue this path, they will need to provide independent reasons, which are not derived from their initial theoretical motivations, to consider misremembering as being just another form of confabulation. It is not clear, however, whether there are such reasons. In fact, recent work on the subject seems to suggest that it is the other way around. In the absence of such reasons, then, relationalism becomes vulnerable to criticism relating to its failure to provide a picture of remembering that is sensitive to research being done in the empirical sciences (see Section 4.1). And this is, as I pointed out at the beginning, one important reason why relational views have not enjoyed much popularity recently. So, denying that there is a real distinction between misremembering and confabulation is likely to bring more problems than solutions to relationalists.

To conclude, let me discuss a worry that might arise in relation to my discussion of unsuccessful remembering. One might (rightly) point out that my characterization of unsuccessful remembering implicitly assumes that the erroneous or distorted elements of memory have to be derived from different experiential sources. However, that is not necessarily the case, as the erroneous elements in some forms of misremembering, such as DRM cases, need not derive from any experiential source.⁹ In response, I would like to clarify that, despite focusing on those cases here, I am not committed to the idea that memory errors are necessarily cases where the erroneous elements are derived from experiential sources. Indeed, in DRM cases, it looks like that appealing to non-experiential elements, such as

⁹I'm grateful to Sarah Robins for calling my attention to this issue.

semantic information influencing retrieval, is more adequate to explain the occurrence of those errors. So, my argument here is not meant to apply to those cases. However, when we look at them more closely, it looks like that the relational view will also have trouble to explain the occurrence of those forms of misremembering. The reason why accounting for error in terms of the influence played by semantic information present at the time of retrieval is appealing is that a form of representationalism is assumed beforehand. In other words, because memory is taken to be a form of representation, there is no mystery in how the content of semantic memories, or semantic information more generally, can distort the content of episodic memories. This alternative is not, however, available to relationalists, for it is not clear whether talk of representation can be coherently integrated into their accounts. So, relationalism will also be required to provide an account of those cases in addition to the ones that were discussed here.

4.4 Conclusion

This paper had one major goal, which was to develop our understanding of the issues surrounding the relationship between relational views and memory errors. This has been accomplished by, first, producing a relational analysis of remembering compatible with the most prominent relational view in the literature, and second, by considering how such analysis of remembering deals with unsuccessful remembering or memory errors. I argued that the relational view, in its current form, cannot properly distinguish between successful remembering and different forms of unsuccessful remembering. I did that by distinguishing between two ways in which unsuccessful remembering can happen: misremembering and confabulation. I argued that the requirement for the presence of an experiential relation, as defended by Debus (2008), is neither necessary nor sufficient for remembering. This is because there are, on the one hand, cases of remembering that do not instantiate the relevant experiential relations, and, on the other hand, cases of confabulation and misremembering that do instantiate the relevant experiential relations. For this reason, if the relational view is to be successful in the context of empirical research on memory; more specifically, in the

context of empirical research on memory errors, it will need to provide extra considerations to show how it can deal with such occurrences.

Chapter 5

The hybrid contents of memory

Introduction

This chapter proposes a novel account of the contents of memory. By drawing on insights from the philosophy of perception, I propose a hybrid account of the contents of memory designed to preserve important aspects of representationalist and relationalist views. The hybrid view I propose also contributes to two ongoing debates in philosophy of memory. First, I argue that, in opposition to eternalist views, the hybrid view offers a less metaphysically-charged solution to the co-temporality problem. Second, I show how the hybrid view conceives of the relationship between episodic memory and other forms of episodic thinking. I conclude by considering some disanalogies between perception and memory and by replying to objections. I argue that, despite there being important differences between memory and perception, those differences do not harm my project of tracing an analogy between them.

Chapter 5 brings my discussion of memory to an end, and with it, the thesis is concluded. I said at the beginning that the overall argument of the thesis is that hybrid accounts of perception and memory should be preferred. Chapter 5 thus relies on a prominent hybrid account of perception to develop a hybrid account of memory. As such, it provides the final, and perhaps the most fundamental, element to connect the work done in Part I and Part II of the thesis. Moreover, Chapter 5 provides an important contribution to the memory liter-

ature. Besides developing an original hybrid account of memory, which, to my knowledge, no one has done so far, the chapter brings relationalism back to the main discussions about memory, which, as I discussed in the general introduction, is dominated by representational views. However, the re-introduction of relationalism, so to speak, is made in a different setting. Instead of framing relationalism as being opposite to mainstream representationalism, I attempt to show how representationalists can consistently and coherently accept relationalist insights. So, while Chapter 4 was mainly critical of relationalism, Chapter 5 can be seen as a first step towards a positive engagement with it.

Despite my final discussions of perception and memory having resulted in hybrid views, the reader will notice that there are important theoretical and practical differences between the discussions offered in each domain. I would like to discuss three of those differences here. The first of these issues relates to my reliance on Schellenberg's hybrid account of perception to provide a hybrid account of memory in this chapter. In Chapter 2, I formulated my own pragmatist-inspired hybrid account of perception in opposition to Schellenberg's account, which makes the proposal of Chapter 5 seem unmotivated or odd. The decision to rely on Schellenberg's account was based on a set of different considerations. The first is that, as I discussed in the introduction to Chapter 2, my own pragmatist-inspired hybrid account requires taking on unfamiliar commitments to many readers, which might be reason enough to put some readers off. And, since the main goal of this chapter was to introduce hybrid accounts to the memory literature, it seemed to make more sense, from a practical point of view, to use a more familiar and established view in the perception literature as a starting point. The second important consideration is that this chapter started off as an exploratory work. I was not entirely sure whether there could be hybrid views of memory, and since Schellenberg has provided fairly detailed discussions of her hybrid view of perception, it seemed to make more sense to use it as a starting point. I think, however, that my hybrid account of memory is also biased towards representationalism, which will likely raise similar issues to the ones I discuss in Chapter 2. In particular, one topic that I have not explored here, but that is relevant in this context, is whether there can be a "Hybrid Triad" of memory. A Hybrid Triad of memory would require an account that, following the

suggestion made in Chapter 2, firstly recognizes the representationalist insight that memory depends on how subjects represent the world to be; secondly, preserves the relationalist insight that events shape our memories in a constitutive sense; and thirdly, offers a direct realist view of memory. Again, while I think one could make a good case for incorporating these claims into one's theory of memory, it is not clear to me how a Hybrid Triad of memory can be addressed properly.

One natural suggestion here would be to develop a pragmatist-inspired hybrid account of memory. In the same way in which Schellenberg's framework can be adapted to memory, one could argue that my own framework of Chapter 2 could be adapted to memory. This brings us to the second difference between my approach to memory and my approach to perception. With the exception of Chapter 3, pragmatism drops out of the picture in my discussions of memory in Chapters 4 and 5. The reader might reasonably ask why that is the case. The reasons I offer here are mostly practical, but there is also one important theoretical reason. To start with the practical reasons, pragmatism drops out of the picture because, as I said above, I thought it would be more effective to rely on a more established view in philosophy of perception to make the case for hybrid views of memory. A second practical reason is that, while Peirce talks a lot about perception, he does not have any detailed discussion of memory. Some Peircean scholars, most notably Joseph Ransdell, have indeed attempted to trace parallels between Peirce's view of perception and his view of memory, and in my bibliography research, I could only find one paper dealing with this topic explicitly, which happens to be currently under review. So, while I had a great number of resources to draw on when connecting Peirce's view of perception to the philosophy of perception, it was not clear to me how (or whether) this could be done in the case of memory.

This was aggravated by the worry, which I mention in the introduction to Chapter 3, that there did not seem to be a meaningful way to talk about pragmatic agreement in memory. This brings us to the main theoretical difference between my discussion of perception and memory. As I mentioned before, this worry was partly due to my inclination to interpret, early on, pragmatic agreement as being an *actual* situation rather than an *ideal* or *counter-*

factual one. While it is still not clear to me how to motivate talk of pragmatic agreement in the context of memory, I am now convinced that, if we interpret this notion in terms of ideal or counterfactual situations, the framework developed in Chapter 2 can be adapted to approach memory. Another important issue refers to whether there can be, in the case of memory, similar notions to the percept and the perceptual judgment, which were central in Chapter 2. If the similarities between perception and memory are as close as I picture them in this chapter, then, again, I think there could be memory equivalents of those notions. While these are all relevant topics to think about a pragmatic account of memory, I did not explore them in the thesis. But they are, indeed, the tasks of future projects. While it is understandable that some might find these reasons unconvincing, I hope that they suffice to at least give some context to the choices made in this chapter.

Outline

This chapter proposes a hybrid view of memory that reconciles insights from representationalist and relationalist views about the objects of memory. Section 5.1 introduces the problem of the objects of memory and analyzes the representationalist and the relationalist answers to it. Section 5.2 develops a hybrid view of memory, called *hybridism*, by relying on insights drawn from the philosophy of perception. I argue that, by adopting a hybrid view of memory inspired by hybrid views of perception, we can consistently preserve good elements from representationalism and relationalism about memory without inheriting their problematic aspects. Section 5.3 expands hybridism and applies it to two ongoing disputes in philosophy of memory: the co-temporality problem and the relationship between episodic memory and other forms of episodic thinking. Finally, Section 5.4 considers some objections to my project of drawing an analogy between memory and perception.

5.1 The objects of episodic memory

When we remember events from our personal past, our memories seem to refer, or to be about, things that long ceased to exist. For example, when I remember my tenth birthday

party, it seems to me that I stand in a relation to an event, my tenth birthday party, that occurred in the past, but that no longer exists. Recently, the psychologist Endel Tulving (1972; 1985a) called *episodic memories* the memories that refer to or are about events. Despite there being different kinds of memories (see, e.g., Squire 2009; Michaelian and Sutton 2017; Werning and Cheng 2017), I will be concerned exclusively with episodic memories, or what philosophers sometimes call “recollection”.

Although I said that episodic memories are memories about events, this characterization can be misleading. Because other kinds of memory, such as semantic memory, can also be about events — e.g., remembering *that* Uruguay won the 1930 Football World Cup — a more precise characterization would be that episodic memories are about events that subjects experienced previously in their lives.¹ So, remembering my tenth birthday party is an episodic memory because I was perceptually related to that event previously in my life. In contrast, remembering that Uruguay won the 1930 Football World Cup is not an episodic memory, for despite being about an event, I did not experience it.

One important question relating to episodic memories refers to the nature of their objects. Traditionally, two opposing accounts have been proposed. Representational theories, or simply *representationalism*, hold that when I remember a past event, I am *directly* related to a mental representation of the event, but only *indirectly* related to the event itself (Locke 1975; Hume 2011; Russell 1921). In contrast, relational theories, or simply *relationalism*, claim that, when I remember a past event, I stand in a direct relation to the event itself (Reid 2000; Russell 1912; Laird 1920; Debus 2008). In opposition to representationalism, relationalism denies the presence of any intermediaries between memories and the events remembered. Despite its historical prominence and importance, the question about the objects of memory has been largely neglected in the contemporary debate. It is not entirely clear why this is the case, though. As I will argue below, how one conceives of the objects of memory will change how one understands different aspects of the metaphysics and the epistemology of episodic memory. In this context, despite my main goal here being to provide a framework in which representationalism and relationalism can be reconciled, I also

¹In fact, even this characterization is problematic, as I can semantically remember events that I experienced previously in my life. For my purposes, however, these problems can be put aside.

expect to make clear the importance that this question has for contemporary philosophers of memory.

The opposition between representationalism and relationalism can be better visualized by looking at how they answer three different questions about memory. The first question, which I will call the *problem of error*, is how memory errors are possible. The second question, which I will call the *problem of indistinguishability*, is how and why successful and unsuccessful occurrences of remembering can be indistinguishable from the point of view of subjects. Finally, the third question, which I will call *problem of epistemic particularity*, is the question of how memory grounds our knowledge of particular past events.

Representationalism offers simple and intuitive answers to the two first questions. Because memory is said to represent past events, error can be explained by appealing to the notion of content. Content, as Rowlands (2017) points out, is normative, which makes it possible to assess memory representations for accuracy. In relation to indistinguishability, representationalism explains it by saying that successful and unsuccessful occurrences of remembering share a “common factor”, i.e., they all have representations as their objects. So, because their objects are of the same kind, successful and unsuccessful remembering can be phenomenologically indistinguishable. However, representationalism faces trouble explaining the problem of *epistemic particularity*. When I remember a past event, such as my tenth birthday party, it seems to me that I remember an event that actually happened (see Debus 2008, 2014; Perrin 2016). But, since qualitatively identical representations can be the objects of both successful and unsuccessful remembering, or even of other forms of episodic thinking (see De Brigard 2014a; Michaelian 2016c), the relation that the objects of successful remembering establish to the actual past events seems to be entirely contingent, which makes it hard to see how those mental states can ground our knowledge of the past.

In contrast to representationalism, relationalism gives the question of epistemic particularity a central place. By denying that there are intermediaries between remembering and the past events, subjects are now placed in direct contact with those events. Debus (2008), for example, argues that, due to the presence of an “experiential relation”, where such relation is understood as supervening on causal, spatial, and temporal relations holding between

subjects and the relevant events, remembering is capable of putting us in direct contact with the past events themselves. However, the attempt to explain epistemic particularity comes at a price, which is that of providing a counterintuitive and excessively complex account of error and indistinguishability. Because the objects of remembering are the past events, and because those events do not exist in cases of unsuccessful remembering, relationalism has to provide an account of the nature of the objects of unsuccessful remembering. While it is relatively easy to see how there can be a direct relation between a mental state and an event that exists or that existed at some point, it is unclear whether the same can be said of non-existing objects or events. Although some philosophers, most notoriously Brentano (1973) and Meinong (1960),² believe that such relation is possible, contemporary relationalists, such as Debus (2008), have favored a different strategy. This strategy consists in adopting a *disjunctivist* account of memory along the lines of disjunctivist accounts of perception (see Martin 2004; Fish 2009). The main claim of disjunctivism is that successful and unsuccessful remembering are only similar with respect to their phenomenology. Based on this, it is further claimed that phenomenological similarity is not enough to group those mental states under the same metaphysical kind.

Disjunctivism becomes appealing to relationalists because it helps them to explain the problems of error and indistinguishability. With respect to error, unsuccessful occurrences of remembering can be simply regarded as being different in kind from successful ones. This allows relationalists to consistently hold that, while successful remembering is essentially relational, unsuccessful remembering is not. As to the problem of indistinguishability, relationalists can simply deny that indistinguishability alone is sufficient to group mental states under the same kind (see Debus 2008, 414–5; Martin 2004, 37).

Although disjunctivism offers an alternative to relationalists in accounting for error and indistinguishability, it faces a number of problems that make it unappealing. Besides offering a counterintuitive account by underplaying the importance of phenomenology to understand the nature of memory, disjunctivism is also problematic when considered in relation to the neurocognitive mechanisms underlying episodic memory. Empirical work

²Although see Crane (2001; 2013) for a contemporary discussion on the topic.

suggests that episodic memory is an instance of a more general capacity to imagine events (Suddendorf and Corballis 1997, 2007; Schacter et al. 2012; Michaelian 2016c), or more generally to think counterfactually (De Brigard 2014a), which suggests that memory is, indeed, similar in important respects to those mental states. Moreover, disjunctivism leads to an extreme and unmotivated view of memory, in which successful occurrences of remembering become rare occurrences, for given the constructive character of episodic memory (see Bartlett 1932; Michaelian 2011), most of the memories occurring in ecological contexts should allow for some degree of inaccuracy (Conway and Loveday 2015; see also De Brigard 2014a).

In summary, a large part of the disagreement between representationalism and relationalism is due to different conceptions of what elements should be central in an account of memory. On the one hand, representationalism provides a simple and unified account of error and indistinguishability, but faces problems to explain epistemic particularity. On the other hand, relationalism provides an account of epistemic particularity, but commits to an implausible view, i.e., disjunctivism, to deal with error and indistinguishability. Despite the apparent incompatibility between them, the question of whether a reconciliatory view is possible has received little attention in the philosophy of memory. As Michaelian and Sutton (2017) have noted in a recent survey of the area, “the prospects for hybrid views of memory remain unexplored”. In contrast, hybrid views incorporating elements from relationalism and representationalism are becoming popular in the philosophy of perception (see Schellenberg 2010, 2014; Siegel 2010; McDowell 2013; Logue 2014; Sant’Anna 2018b) and, I shall argue, they provide insightful resources to think about the possibility of hybrid views of memory. In the remaining sections, I will focus on one prominent hybrid view of perception developed by Susanna Schellenberg (2010; 2011; 2016) and I will propose a similar hybrid view of memory based on it.

5.2 Towards a hybrid view of memory

In recent works, Schellenberg (2010; 2016) has suggested that the disagreement between representationalism and relationalism about perception is due to, at least in part, their focus on different elements of perception. According to her, representationalists are more concerned with explaining the *phenomenological particularity* of perception, where “a mental state manifests phenomenological particularity if and only if it seems to the subject that there is a particular present”, that is, “[...] if and only if the particularity is in the scope of how things seem to the subject” (2016, 28), while relationalists give more emphasis to its *relational particularity*, where “a mental state instantiates relational particularity if and only if the mental state is constituted by the particular perceived” (2016, 28).

The focus on different aspects of perception, Schellenberg adds, has motivated different strategies to individuate perceptual states. Representationalists, according to her, adopt the *mental state view*, where “experiences are individuated solely by the phenomenology that the subject experiences” (2010, 20). Alternatively, relationalists rely on the *environment-encompassing view*, where “experiences are individuated by the phenomenology and the material, mind-independent objects, properties, scenes, or events to which the subject is perceptually related” (2010, 21).

I think that drawing similar distinctions can be helpful to understand the disagreements between representationalism and relationalism about memory. To avoid confusion, let me start by defining the memory equivalents of those notions. I will say that a memorial state instantiates *phenomenological particularity* iff it seems to the subject that his mental state is about a past event that was previously experienced. Thus, remembering, misremembering, and confabulating alike can instantiate phenomenological particularity. Similarly, I will say that a memorial state instantiates *relational particularity* iff the memory is constituted by an event that took place in the subject’s personal past, which refers to the collection of events that the subject experienced in his life prior to the memory. Thus, successful occurrences of remembering instantiate relational particularity, but unsuccessful occurrences do not, for only the former are constituted by events of the subject’s personal past. Moreover, similar

to the perception case, the attempts to explain phenomenological particularity and relational particularity motivate different strategies to individuate memorial states. The first strategy, adopted by representationalists, is the mental state view described by Schellenberg (2010), according to which memories are individuated by their phenomenological character. The second strategy is what I will refer to as the *past-encompassing view*, according to which memories are individuated by their phenomenology and the events of the subject's personal past.

One clarification here refers to what I mean when I say that a memory is *constituted* by an event that took place in the past. I am using the term 'constitution' in a very general way, such that there are multiple ways in which a past event may constitute a present memory. One such way would be by means of a memory trace (Martin and Deutscher 1966; De Brigard 2014b; Robins 2016b). On this view, a past event is a constitutive part of a present memory because the latter preserves a causal connection to the former. The notion of constitution used in this case, however, is not that of material constitution. As I will discuss in more detail later (see section 5.3.1), an event can constitute the content of memory even when that event has ceased to exist. Another way would be to say, along with eternalists (e.g., Bernecker 2008), that the event itself is part of the memory because it never ceases to exist. In section 5.3.1, I suggest that favoring the first alternative can alleviate some metaphysical worries about the objects of memory, but for the purposes of my discussion of the reconciliation of representationalism and relationalism, I do not need to commit to any of these alternatives.

Let me now discuss how the distinctions introduced above relate to the discussion in section 5.1. Consider representationalism first. The claim that successful and unsuccessful occurrences of remembering share a "common factor" is supported by the mental state view because it seems to subjects that their memories relate to particular events, which explains why memory instantiates phenomenological particularity. On the representationalist account, successful and unsuccessful remembering both have representations as objects, thus it is not surprising that their phenomenologies might be indistinguishable. However, representationalism fails to explain relational particularity. Because representations are

decouplable (see Rowlands 2017), they can occur whether or not the things that are represented exist. So, the occurrence of a past event does not seem necessary for the occurrence of memory. This makes it hard to see how, on the representationalist view, memory can ground our knowledge of the past.

Consider now relationalism. By adopting the past-encompassing view, relationalists can explain epistemic particularity. Since successful remembering requires being directly related to past events themselves, it is easier to see how memory grounds our knowledge of the past. Moreover, because subjects are directly related to events, relationalism also explains why successful remembering has phenomenological particularity. However, because unsuccessful occurrences of remembering do not have objects, the past-encompassing view faces trouble to explain how they instantiate phenomenological particularity. That is why disjunctivism becomes appealing: because phenomenological particularity is arguably the only thing shared between successful and unsuccessful remembering, and because successful remembering is constituted by events, relationalists can postulate a fundamental separation between them based on the past-encompassing view. So, while it allows for a simple account of epistemic particularity, the past-encompassing view makes things complicated for relationalists in relation to error and indistinguishability.

The distinctions introduced here help us not only to understand the opposition between representationalism and relationalism, but also provide an initial framework to conceive of a possible reconciliation. The suggestion I want to put forward is that a hybrid or reconciliatory view needs to explain both the phenomenological particularity and the relational particularity of memory. However, this raises an important question, which is how this can be done without resulting in any inconsistencies. In the case of perception, Schellenberg argues that we need to preserve the idea that perception has content and adopt the environment-encompassing view as a general strategy to individuate perceptual states. I will follow her suggestion here and propose that, in the case of memory, preserving the idea that memory has content and adopting the past-encompassing view provides the path to reconcile representationalism and relationalism. But, before we move on, a few words on why it is important to preserve those things. On the one hand, the idea that memory

has content is important because it provides a way to distinguish between successful and unsuccessful remembering by allowing us to assess those mental states for accuracy. In doing this, we can explain, moreover, the nature of the “common factor” that those mental states share — i.e., both have contents — thus providing an account of indistinguishability. On the other hand, preserving the past-encompassing view is important because it allows us to distinguish successful remembering from unsuccessful remembering in terms of its relationship to past events. That is, in the same way that, according to Schellenberg’s environment-encompassing view, the particulars that are perceived constitute the contents of perception, the past-encompassing view requires that the particulars that are remembered — i.e., events, objects, etc. — constitute the contents of memory.³ This allows us to build into our theory the relational aspect that explains how memory grounds our knowledge of the past. More importantly, because we can now say that successful remembering and unsuccessful remembering have a common factor, the relational aspect introduced by the past-encompassing view need not lead us to disjunctivism. The differences between those mental states can be properly accounted for without positing a fundamental separation between them.

The claims made above will be developed in more detail in the next sections. However, before we enter this discussion, I shall mention one important assumption that I will be making. In order to reconcile content and the past-encompassing view, I will be assuming an unconventional way to understand the relationship between the content and the phenomenology of mental states, which I refer to as *separatism*, and then explore how it allows us to coherently incorporate content and the past-encompassing view into a unified view. Separatism opposes *intentionalism*, which is the view that the phenomenology of our mental states can be adequately explained by their representational or intentional content (see, e.g., Dretske 1997; Tye 2000; Byrne 2001, 2009). So, as Fernández (2017) points out, separatism says that “the phenomenal and intentional [or representational] features of mental states are independent from each other” (97). Therefore, my claim is that *if* we take for granted that a form of separatism is true, we can have a proper hybrid view of memory. I

³See section 5.3.1 for a more detailed discussion of what it means to say that events constitute the contents of memory.

do not expect the reader to get on board with separatism at this point; in fact, a proper argument for a hybrid view of memory will require a proper argument for separatism. However, since my goal here is to explore whether a hybrid view of memory is possible, I shall set this question aside for the moment.

Another important thing to note here is that the hybrid view of perception offered by Schellenberg is not uncontroversial. One issue is that it leans too much on representationalism, which will make the view unappealing for those inclined to relationalism (see Sant'Anna 2018b; Chapter 2 for discussion). In particular, one central point of disagreement between representationalists and relationalists is about how to account for the phenomenology of perception. Relationalists, such as Martin (2004) and Fish (2009), tend to explain it in terms of particulars being constitutive parts of experiences. However, for the hybrid view, the phenomenology is explained by the modes of presentation, which are elements composing the representational content of experience (see 5.2.1). Another complaint, but now coming from representationalists, is that hybrid contents, or relational contents, build non-conceptual elements into the content of perception. This is problematic because it undermines an important additional motivation for representationalism, which is that it provides a neat account of how the content of perception can inform the content of other mental states, especially beliefs (see McDowell 1996).

Thus, it is not entirely clear whether hybridism has been successful in fully reconciling representationalism and relationalism. Whether there can be such a full reconciliation or what elements from each view should figure into a reconciliatory view are still open questions in the perception literature (see Locatelli and Wilson 2017, 209). The same challenges, I think, should be expected to arise in relation to memory. However, these challenges do not diminish the importance of developing hybrid views. Even if hybrid views fail to satisfy the demands of “pure” representationalists and “pure” relationalists, they still offer a promising alternative for those who are not convinced by either of the “pure” theories. Moreover, in historical terms, this is an important development, as hybrid views show that two apparently incompatible views are not incompatible after all.

5.2.1 Hybrid contents

Separatism allows for the formulation of a hybrid notion of content. *Hybrid contents*, as I will call them, refer to the satisfaction conditions of memory that are partly determined by its phenomenology and its relation to past events. Hybrid contents are, therefore, an alternative notion of content designed to preserve both the phenomenological particularity and the relational particularity of memory. I will now explore this notion in detail.

In line with my previous discussion, I will rely on Schellenberg's account of content here.⁴ In a recent paper, she has characterized hybrid contents as two-place relations holding between a mode of presentation of an object and a mode of presentation of a property (Schellenberg 2010). A mode of presentation here refers to how an object or a property appears or becomes cognitively available to the subject. So, on the classic example discussed by Frege (1980), the same object ("Venus") can have different modes of presentation ("morning star" and "evening star") in different thoughts. Similarly, the idea here is that, in the case of perception, objects and properties can be presented in different ways in perceptual experiences. In terms of their ontological status, modes of presentation can be viewed as parts of the representational content of mental states responsible for determining their phenomenology. Because they establish what it is like for subjects to undergo different mental states, it is possible for those mental states to refer to the same thing while not necessarily sharing the same phenomenology. With this in mind, Schellenberg suggests that the following characterization of the contents of perception can be given, where MOP refers to modes of presentation, o is an object, and p is a property:

$$\mathbf{Perception} = [\text{MOP}_1(o); \text{MOP}_2(p)]$$

I will adopt the same characterization to define the hybrid contents of memory, but instead of objects and properties of objects, I will replace these with events and properties of events. So, the hybrid contents of memory are characterized by the following, where MOP refers to modes of presentation, e is an event, and p is a property of e:

⁴I should note that the term "hybrid contents" is my own terminology. Despite not using the same term, I ascribe the core idea behind this notion to Schellenberg's (2010; 2011; 2016) account of perceptual content.

$$\mathbf{Memory} = [\text{MOP}_1(e); \text{MOP}_2(p)]$$

Let me clarify what these terms mean. I will not commit to any particular account of events here; the term is used straightforwardly to refer to situations such as visiting the Cologne cathedral or drinking beer at the pub. Properties of events are, accordingly, the particular constituents of those events. For example, when I remember visiting the Cologne cathedral on a cloudy day, being a “cloudy day” is a property of that event. Similarly, remembering “having pilsner” is a property of the event “drinking beer at the pub”. The semi-colon separating the two modes of presentation indicates that properties are presented as being instantiated by events despite the fact that, in the analysis of the content, they are related to different modes of presentation.

Another thing that needs to be clarified is what it means to say, in the case of memory, that modes of presentation present events and properties as being a certain way to subjects. Consider the case of visiting the Cologne cathedral again. When I remember this event, the mode of presentation presents it as being located in the past and as having occurred. In other words, modes of presentation of events are responsible for presenting events as being in a certain *temporal location* and as being *actual*, in the sense that they happened before, or as being *possible*.⁵ Contrast this with imagining visiting the Cologne cathedral. On such cases, the mode of presentation of the event places it in the future and identifies it as something that can possibly happen. Similarly, the modes of presentation of properties are responsible for presenting properties as being instantiated or not by a particular event.⁶ For example, when I remember visiting the Cologne cathedral on a cloudy day, that property is presented to me as being instantiated by the event in question. Likewise, when I think about how it would be to visit the cathedral on a sunny day, the mode of presentation presents the property of the event as being instantiated by the event too. The important thing to note here

⁵See section 5.3.2 for a more detailed account. I should note here that the temporal location specified by the modes of presentation is coarse-grained in the sense that it does not specify a particular day or time, but only whether the event is located in the past or in the future.

⁶While I distinguish between modes of presentation of instantiated and non-instantiated properties here, most occurrences of remembering and also of episodic thinking discussed in section 5.3.2 will contain only modes of presentation of instantiated properties. This is because, in most occurrences of those mental states, the properties are presented to subjects as being instantiated, even though they might not be in reality. Thus, while it might be possible for a subject to remember non-instantiated properties of events, such as remembering a cathedral and a sunny day, but not ascribing these to any particular event, I will focus, from now on, exclusively on cases where the properties are instantiated.

is that modes of presentation are the parts of the content responsible for the phenomenology of memory, so it is possible for modes of presentation to present events as being the case and properties as being instantiated without implying that memory refers to events that actually happened or to actual instantiated properties. Reference, as I will argue below, is established by another part of the content of memory, which does not necessarily influence what it is like for subjects to undergo memorial states.

Another important thing to note is that we do not need to restrict ourselves to only one of the properties of events. It is possible, in principle, to have a memory whose content has two or more modes of presentation referring to different properties of an event. For example, I can remember visiting the Cologne cathedral on a cloudy and hot day. In this case, the content of my memory has a mode of presentation relating to an event and two different modes of presentation relating to two different properties:

$$\mathbf{Memory} = [\text{MOP}_1(e); \text{MOP}_2(p_1); \text{MOP}_3(p_2)]$$

It might be argued that this account of the content of memory does not provide a clear way to distinguish between events and properties. Suppose that I remember visiting a building identical to the Cologne cathedral on a cloudy day, but for some reason I remember this event as taking place in Hamburg. There are two possible ways to characterize the constituents of the content of this memory. We can, on the one hand, say that it is composed by a mode of presentation of the event “Visiting a building identical to the Cologne cathedral in Hamburg” and a mode of presentation of the property “cloudy day”, or we can, on the other hand, say that it is composed by a mode of presentation of the event “Visiting a building identical to the Cologne cathedral” and two modes of presentation of the properties “in Hamburg” and “cloudy day”. If this is right, however, we will have different assessments of the accuracy of the same memory. In the first case, the memory might be said to be confabulatory because the event clearly did not occur. In the second case, it might be argued that the memory is an occurrence of misremembering because the event in question occurred, but one of the properties represented failed to be instantiated.⁷

⁷For more details on confabulation and misremembering, see section 5.2.2.

As I mentioned initially, I am not committing to any particular metaphysical view of events, so there is no principled way to say that one of the characterizations above is better. However, we can choose between them in the context of our analysis of the content of remembering. In other words, whether a property will be characterized alongside an event, as in the first case, or whether it will be assigned a separate mode of presentation, as in the second case, depends on the questions that we are trying to answer with our analysis of the content of remembering. If the relevant question is “Does the subject remember the event of visiting a building identical to the Cologne cathedral in Hamburg?”, then his mental state is clearly confabulatory. But, if the question is “Does the subject remember the event of visiting a building identical to the Cologne cathedral?”, then it would make more sense to say that he does remember such event, although he gets some of the details wrong.

While some might find this pragmatic answer unconvincing, it does not look completely absurd when we consider the fact that a large fraction of the memories that we usually consider “true” involve inaccurate elements. That is, because, as De Brigard (2014a) notes, “remembering is a particular operation of a cognitive system that permits the flexible recombination of different components of encoded traces into representations of possible past events [...] in the service of constructing mental simulations of possible future events” (158), it is not uncommon for the so-called “true” memories to have some inaccurate elements. Thus, if we want an account of the content of memory that avoids the conclusion that most of our memories are not true, a pragmatic strategy that focus on some but not all elements of the content according to the purposes of the analysis starts to make more sense.

Now, given this characterization of hybrid contents, it is possible to address the problems raised in section 5.1 in a reconciliatory framework. For the sake of simplicity, I will focus on cases that involve only one event and one property.

5.2.2 Error revisited

Let us start with the problem of error, which refers to the possibility of memory errors. For hybridism, because all occurrences of remembering have content, it is possible to assess them for accuracy. To clarify, consider successful occurrences of remembering first. Hy-

bridism says that the content of remembering consists in a two-place relation between a mode of presentation of a particular event e_1 , which took place in the past, and a mode of presentation of a property p_1 , which happens to be a property of e_1 :

$$\mathbf{Remembering} = [\text{MOP}_1(e_1); \text{MOP}_2(p_1)]$$

This characterization explains why remembering instantiates phenomenological particularity. The modes of presentation of events and properties, which are responsible for making them cognitively available to the subject, make it seem to the subject that he is remembering a particular event with a certain property. Moreover, because both modes of presentation are successful in establishing reference, the relevant event e_1 and the property p_1 instantiated by it become constitutive parts of remembering. This explains how remembering establishes a relation to past events, and consequently, how it instantiates relational particularity.

Consider now unsuccessful occurrences of remembering, which can be divided into two different kinds. The first kind is *misremembering* (see Robins 2016a; Michaelian 2016b), which refers the cases where we mistakenly remember some feature of a past event. For example, when I remember having chocolate cake at my tenth birthday party, I mistakenly remember a feature — that I had chocolate cake instead of strawberry cake — of an event that happened, that is, my tenth birthday party. The second kind is *confabulating*, which refers to cases where we remember events that did not happen, such as remembering that I went to the beach on my tenth birthday.

Hybridism says that the content of misremembering consists in a two-place relation between a mode of presentation of a particular event e_1 , which took place in the past, and a mode of presentation of a property.

$$\mathbf{Misremembering} = [\text{MOP}_1(e_1); \text{MOP}_2(__)]$$

However, because the property presented was not instantiated by e_1 , its mode of presentation fails to establish reference. In Schellenberg's (2010) account, when the modes of presentation fail to refer, the content becomes "gappy". Since, in the case of misremembering, only the second mode of presentation fails to refer, I shall say that its content is *partially gappy*. To further clarify this point, I should say a little bit more about what it means for

a content to be “gappy”. Although the most natural interpretation is to understand it as meaning that the content is somehow empty, in the sense that subjects will experience their memories as having “missing” parts, that is not what I have in mind here. *Gappiness* is a theoretical notion here and its meaning is simply that the mode of presentation failed to establish reference. As I discuss below in 5.2.3, because of separatism, the gappiness or non-gappiness of modes of presentation do not change the phenomenology of memory. For this reason, I should emphasize, again, that the memories whose contents are gappy will not be presented to subjects as being empty or as missing some part.

Now, turning back to the characterization of misremembering given above, note that it explains why misremembering instantiates phenomenological particularity. Despite one of the modes of presentation being gappy, it still seems to the subject that he is remembering an event with a property. Because of separatism, the presence or the absence of reference by the modes of presentation only needs to make a difference to content, but not to phenomenology. In contrast to remembering, however, misremembering only establishes *partial* relational particularity because only one of its modes of presentation is successful in establishing reference.

Finally, consider the case of confabulating. The content of confabulating consists in a two-place relation between a mode of presentation of a particular event and a mode of presentation of a property, but both fail to establish reference.

Confabulating = [MOP₁(__);MOP₂(__)]

That is, because the event that is made cognitively available to the subject did not happen, the first mode of presentation fails to refer to something. Consequently, the second mode of presentation also fails to refer, for there are no properties instantiated by events that did not happen. In contrast to remembering and misremembering, then, the content of confabulating is *fully* gappy.

Analogously to remembering and misremembering, this characterization of the content of confabulating explains why it instantiates phenomenological particularity. The gappiness of the modes of presentation, as I noted above, makes a difference to the representational content of the memory, but not to its phenomenology. However, unlike remembering and

misremembering, because neither the mode of presentation of the event nor the mode of presentation of the property are successful in establishing reference, confabulating does not instantiate relational particularity.

One question that might be asked here is whether it is possible to have a memory whose content contains a gappy mode of presentation of an event and a non-gappy mode of presentation of a property.⁸ In the terminology used here, the content would look like the following

$$\mathbf{M} = [\text{MOP}_1(_); \text{MOP}_2(p)]$$

One example would be remembering visiting the Coliseum on a cloudy day. Since I have never visited the Coliseum, the mode of presentation of the event would be gappy, but the mode of presentation of the property would not, for the property of being a cloudy day was instantiated by other events that I can remember — e.g., my visit to the Cologne cathedral. However, on the framework developed here, it is not possible to have a memory with such content. Whether or not a property is instantiated by an event is an objective fact about the event, so it cannot be the case that the property “cloudy day” was instantiated by the event “My visit to the Coliseum” because this event did not happen. Despite the fact that “cloudy day” is a property that was instantiated by another event that I can remember, the sense in which instantiation is used here does not imply that the property in the content above was instantiated by the relevant event. Because of this, the memory in question would have a fully gappy content and, therefore, it would be a confabulation.

To conclude the discussion of memory errors, I should point out that I have been concerned with only some forms of memory errors, namely, misremembering and confabulating. There are, however, other forms of errors that are of concern to philosophers of memory. In a recent discussion about how to provide a taxonomy of memory errors, cases of relearning and veridical confabulation have, in addition to misremembering and confabulating, played an important role in shaping the current theories (see Robins 2016a; Michaelian 2016b; Bernecker 2017).

⁸I'm grateful to Kirk Michaelian for pressing me on this point.

Cases of relearning refer to situations where we experience an event, forget about it, and then re-acquire information about the event from sources other than our own episodic memories. For example, I experienced my first day at school many years ago, and for some reason or another, I forgot about what happened that day. However, due to talking to my parents, I re-acquired or re-learned some relevant information about the event of my first day at school. Relearning is said to be a form of memory error because, despite the information that is re-acquired being accurate, it is acquired second-hand (e.g., from testimony), as opposed to the first-hand information that is acquired through remembering (from the past experience). Veridical confabulations are, in contrast, cases where subjects represent past events accurately, but the accuracy obtains accidentally (see Michaelian 2016b; Bernecker 2017). For example, a subject might describe accurately what he had for dinner yesterday and take himself to be remembering this event, but because the underlying processes producing his putative memory are not the usual processes that produce episodic memories, or perhaps because he is just guessing, he is said to be undergoing a veridical confabulation. Like cases of relearning, the information conveyed to the subject in veridical confabulations is not appropriately derived from his past experiences. While there might be room to dispute whether relearning and veridical confabulations are genuine forms of memory error (see, e.g., Bernecker 2017), both cases seem to be intuitively plausible. So, it seems reasonable to expect that a complete taxonomy of memory errors will need to account for them. The question is, however, whether the hybrid view can accommodate such errors.

I share the underlying motivation in the current literature that a complete taxonomy of memory errors needs to include relearning and veridical confabulations. It was not my task, however, to provide such a taxonomy in this paper. My discussion of the hybrid view focused instead on central cases of memory errors which have been discussed more extensively by philosophers of memory. Thus, the hybrid view should not be viewed, at least in this stage, as a complete account of memory errors. The discussion of relearning and veridical confabulation will, indeed, require further work from hybridists, but such work is complementary to the main task of this paper concerning memory errors, which was to discuss them in the context of the dispute about the objects of memory.

5.2.3 Indistinguishability revisited

Let us consider the problem of indistinguishability now. This problem refers to the fact that remembering, misremembering, and confabulating can be indistinguishable from the point of view of subjects. Hybridism shares with representationalism the idea that there is a “common factor” between remembering, misremembering, and confabulating, which explains why they can be indistinguishable. However, unlike traditional representationalist accounts, it claims that only *parts* of the contents of those mental states are shared. Those parts correspond to the modes of presentation. Because, again, modes of presentation are responsible for making events and properties cognitively available to subjects, they might be unable to distinguish, from their own points of view, between remembering, misremembering, and confabulating. But, the fact that modes of presentation are shared need not conflict with the fact that only remembering is relational. For separatism, representational content can be different even if the phenomenology is the same. This allows hybridism to incorporate the relationalist idea that, in remembering, the past event and its properties are constitutive parts of the content, which explains why remembering differs from misremembering and confabulating with respect to its relational particularity. On the hybrid framework, then, indistinguishability and relational particularity are properly integrated.

5.2.4 Epistemic particularity revisited

Finally, let us consider the problem of epistemic particularity, which refers to the question of how memory can ground our knowledge of particular past events. Hybridism explains the epistemic particularity of remembering in terms of the distinctive nature of its content; more specifically, in terms of the nature of its relational particularity. The content of remembering is the only one that is non-gappy, meaning that it is the only one that instantiates *full* relational particularity. It is because of the non-gappiness of its content that remembering, as opposed to misremembering and confabulating, can ground our knowledge of particular past events.

The hybrid account thus accepts that there is something distinctive about remembering,

but it does not, like relational accounts, interpret this distinctiveness as a reason to postulate a fundamental separation in relation to misremembering and confabulating. In this respect, it integrates this idea with the representationalist idea that remembering, misremembering, and confabulating share a “common factor”. This provides an account of memory that respects its phenomenology, which is central to representationalist accounts, and its epistemic particularity, which is central to relationalist accounts.

One important thing to note in the context of the discussion of epistemic particularity is that the hybrid view is not saying that successful occurrences of remembering are forms of *knowing* the past (see, e.g., James 2017). That is, despite providing us with true or accurate information about past events, it does not necessarily follow that we know anything about those events by simply remembering them. Instead, by saying that memory instantiates epistemic particularity, I want to say that the hybrid view explains how memory makes it *possible* for us to know anything about the past. This is a very important distinction in the context of the philosophy of memory. The reason is that whether memory is *factive*, i.e., whether memory implies knowledge, is a topic of controversy among philosophers. Since I do not want to take part in this debate, I shall say that instantiating relational particularity, and hence epistemic particularity, is a necessary but not sufficient condition to say that our memories allow us to form knowledge of the past.

5.3 Hybridism and the philosophy of memory

Hybridism provides an alternative solution to the longstanding dispute between representationalism and relationalism. However, it can also be insightful for contemporary debates in philosophy of memory. I will focus on two different topics here: the co-temporality problem and the dispute between continuists and discontinuists about the relationship between episodic memory and other forms of episodic thinking.

5.3.1 The co-temporality problem

The co-temporality problem, as Bernecker puts it, refers to the question of “[h]ow [...] can we be in direct touch with events which occurred and ended a long time ago?” or “[h]ow can the direct object of my present state of remembering be something that has ceased to exist?” (Bernecker 2008, 69). This problem is particularly pertinent for relationalism, for as Norman Malcolm (1975) notes, the claim that we are directly acquainted with past events implies that the past events exist now. But, as the Bernecker quote makes explicit, there is a tension, if not a contradiction, in saying that things that are past exist in the present. So, relationalists have to explain how events can be constitutive parts of remembering when they are not co-temporal with it.

One solution, proposed by Bernecker (2008), is to adopt *eternalism* about events. On this view, “[e]vents do not cease to exist when they cease to be present or when there ceases to be evidence for them. Once an event has happened, it exists eternally; the only thing that still happens to it is that it retreats into the more and more distant past.” (2008, 71). Essentially, eternalism goes against *presentism*, which is the view that only the present — and therefore only present events — exist. The co-temporality problem relies on presentism to make its case against relationalism, for it assumes, as a starting point, that past events are the kind of things that do not exist. So, in adopting eternalism, it is possible to say that past events can be constitutive parts of remembering and thus avoid the co-temporality problem.

It is not my goal to criticize eternalism here, but since some might see it as a high metaphysical price to pay in order to deal with the co-temporality problem (e.g., Michaelian 2016c, 63), hybridism might offer an alternative route for them. For hybridism, what must exist at the time of remembering are the hybrid contents of memory, and not the events themselves. To clarify why this provides a distinctive answer to the co-temporality problem, we should go back to the discussion of constitution earlier. In the beginning of section 5.2, I suggested that there are multiple ways in which an event can constitute the contents of memory. One such way is by means of the presence of a memory trace. A memory trace, as I understand it, is a *referential index* of the original event that figures into the

representational content of a memory.⁹ It is because memory traces work as referential indexes that we can say that events constitute the content of our memories. However, in order to make this clear, I should say more about what a referential index is. A referential index is a thing A that is responsible for *pointing* to the existence of another thing B. A can point to the existence of B because A is existentially dependent on B, that is, A would not have been the case if B had not been the case. For example, smoke is a referential index for fire because it points to the existence of fire and smoke is existentially dependent on fire — i.e., assuming that there is smoke iff there is fire. What is important to note about referential indexes is that they can exist even in the absence of the things that they existentially depend on. That is, even in the absence of fire, smoke can still function as a referential index for fire, in the sense that it points to the existence of fire at a *prior* moment. Similarly, a particular memory trace T is a referential index for a particular past event E because it points to the existence of the event E. Like the case of smoke and fire, T can exist in the absence of E, upon which it is existentially dependent. Thus, in the absence of E, T can still function as a referential index for E, in the sense that it points to the existence of E at a *prior* moment.

Now, the question is how does this help with the co-temporality problem? If one wants to avoid the problem, one needs to provide an account of how an event e can constitute the content of memory without requiring that e exists at each particular time that it is remembered. However, if e is to be a constitutive part of the content, this does not seem to be possible without appealing to eternalism. In logical terms, if A is a constitutive part of B, then, whenever B is the case, A is also the case. To avoid this worry, I will distinguish here between two senses of constituency. On the one hand, we can say that a *materially constitutes* B iff A is a material part of B. On the other hand, we can say that A *metaphysically constitutes* B iff the occurrence of B existentially depends on and is explained by the occurrence of A. Otherwise put, if A had not been the case, B would not have been the case.

For the hybrid view, events constitute contents in the second, but not the first, sense of constituency. That is, an event e is a (metaphysical) constitutive part of a content C iff the occurrence of C existentially depends on and is explained by the occurrence of e. So,

⁹The notion of a referential index is an adapted version of Charles Peirce's notion of an *index*, which plays a central role in his theory of representations.

when one successfully remembers an event e , it is not required that e materially exists at the time of remembering. What makes a given content C metaphysically constituted by e is the fact that, at the time of remembering, C is materially constituted by a memory trace. This is possible because memory traces are referential indexes, and as such, they establish the relationship of metaphysical constituency between contents and events by being material constituents of the contents. So, the hybrid view avoids the co-temporality problem because it allows for the claim that events (metaphysically) constitute the content of memories without requiring their material existence.

I do not expect the reader to get on board with this solution straightaway. One particular worry is that it relies heavily on a particular understanding of memory traces, so more is needed to motivate it properly. Another worry might be that this solution relies too much on the idea that there are memory traces, which might put off some readers. I have not argued for the existence of memory traces here, but since they are important in discussions about the neurobiology of memory (Thompson 2005; Poo et al. 2016), further exploring their implications for philosophical theorizing about memory can be potentially fruitful.

A third worry might be that more needs to be said in order to make explicit what it means for a thing A to be a metaphysical constituent of another thing B and whether referential indexes can be used to describe this relationship appropriately. One might argue, for example, that the relationship between hybrid contents and events established by memory traces is better understood in terms of dependence or entailment, rather than in terms of constitution. I want to acknowledge all these worries and say that they are worth developing in exploring the prospects of the hybrid view. However, due to the limitations of scope and space, I hope that the considerations above are sufficient to give at least an idea of how the hybrid answer to the co-temporality problem is supposed to work.

To conclude this discussion, I would like to address an objection that is likely to arise in the context of this discussion of constitution. This objection says that understanding constitution as metaphysical constitution undermines the main motivation for taking the hybrid view to be a form of hybrid-*relationalist* view. Relational particularity, one might argue, requires material constitution and not metaphysical constitution. In response, I do

not think this is right. The main motivation to account for relational particularity is that it grounds epistemic particularity, which is aptly explained by metaphysical constitution. Due to the referential indexical nature of memory traces, one will have a memory whose content C is about a particular event e iff e was the case. Thus, C necessarily and existentially depends on and is explained by e having been the case, which is what is required to account for epistemic particularity. This might not, of course, satisfy a “pure” relationalist, but since the hybrid view does not aim at being a combination of pure representationalism and pure relationalism, but rather a compromised reconciliation, this is not necessarily a problem for the view.

5.3.2 Continuism vs. discontinuism

As I discussed in section 5.1, the idea that episodic memory is closely related to other forms of episodic thinking, such as episodic future thinking (Michaelian 2016c; Szpunar 2010) and episodic counterfactual thought (De Brigard 2014a), is receiving increasing support from empirical research on memory (see Suddendorf and Corballis 1997, 2007; Schacter et al. 2007, 2012). This has raised the question of whether episodic memory differs from other forms of episodic thinking only in degree, or whether the difference is one of kind. While people involved in the debate often focus on different aspects of the relationship between memory and episodic thinking (see Perrin and Michaelian 2017 for discussion), I will focus here on the *metaphysical* question of whether episodic memory and other forms of episodic thinking belong to the same metaphysical kind.¹⁰

Continuists hold that the similarities between episodic memory and episodic thinking found by empirical research suggest that they are instances of the *same* kind (see, e.g., De Brigard 2014a; Michaelian 2016c). While continuists are not necessarily committed to the view that episodic memory and episodic thinking are the same thing, they do not think that the dissimilarities among them support a fundamental separation. In contrast, discontinuists hold that, while there might be striking similarities between episodic memory and episodic thinking, episodic memory possesses some features that makes it fundamentally

¹⁰See the Appendix A for a more detailed discussion of this.

distinct from other forms of episodic thinking. One common reason offered in favor of discontinuism is that episodic memory holds a causal relationship (Perrin 2016) or an “experiential” relationship (Debus 2008, 2014) to past events, while other forms of episodic thinking do not.

If a hybrid account of memory turns out to be correct, hybridism would support a broad continuist view according to which episodic memory and episodic thinking belong to the same metaphysical kind. A hybrid account would require one to take episodic thinking to possess hybrid contents too. Assuming that this is the case, we can explain the fact that episodic memory and episodic thinking belong to the same metaphysical kind by pointing out the fact that they possess hybrid contents. While hybridism sits on the continuist side, it can still incorporate the discontinuist idea that episodic memory is inherently relational. On this account, while (successful) episodic memory would be non-gappy because it is successful in referring to events and their properties, episodic thinking would be gappy, because it fails to do so.

Episodic thinking is considered to be gappy because it does not refer to existing events and to instantiated properties. In different forms of episodic thinking, whether we are thinking about events that can possibly happen or events that can no longer happen, our thoughts do not refer to particular events. Thus, their modes of presentation are gappy like the modes of presentation of confabulating. But, despite establishing that the contents of episodic thinking are gappy, hybridism can explain the phenomenological similarities between episodic memory and episodic thinking. Due to possessing modes of presentation in their contents, the phenomenological particularity of episodic thinking, or the fact that it seems to subjects that they are thinking about particular events is explained by hybridism. Moreover, it also explains why episodic memory is unique with respect to its relational particularity. Because of the gappiness of episodic thinking, episodic memory is the only one capable of instantiating relational particularity. So, hybridism would grant continuists that episodic memory and episodic thinking are continuous with respect to their phenomenological particularity, but would concede to discontinuists by saying that they are discontinuous in relation to their relational particularity.

One natural worry here would be in relation to the relationship between unsuccessful occurrences of remembering, more specifically confabulating, and episodic thinking. If hybridism is right, then it might be argued that there is no difference between episodic thinking and confabulating because both have fully gappy contents, which is an odd result. One strategy to address this worry would be to say that confabulating is a misnomer for what are actually forms of episodic thinking directed to the past. While this is a possible answer to the problem, I think this strategy should be avoided if other alternatives are available. The main reason is that hybridism would need to provide independent reasons to think that confabulating is simply a form of episodic thinking directed to the past, but it is not clear whether there are such reasons.

As an alternative, I think that a more detailed account of the modes of presentation present in the contents of remembering, misremembering, confabulating, and different forms of episodic thinking, can potentially provide a framework to distinguish between them. To clarify this point, let me first distinguish between different forms of episodic thinking. I will follow De Brigard (2014a) and distinguish not only between forms of thinking about possible past and future events, but also between forms of thinking about counterfactual past and future events. Consider, first, *episodic future thinking*. These refer to cases where we think about events that can possibly happen in the future. For example, when I think about my holidays at the end of the year, my thought presents me with an event instantiating certain properties. However, despite seeming so at the phenomenological level, the thought fails to refer to existing things, for the event in question does not exist and the properties are, consequently, non-instantiated.

Consider, second, *future-oriented counterfactual thinking*. These refer to cases where we think about events that could have been the case in the future, but are no longer possible. For example, when I think about how my holidays would be if I had saved money to pay for them, I am thinking about an event that would happen in the future if I had done something differently in the past. However, at the present, it is no longer possible for me to save money to pay for the holidays, which makes this an impossible event. In such cases, the events in question also do not exist and their properties are non-instantiated. Thus, despite presenting

me with a particular event and some seemingly instantiated properties, the thought fails to refer to existing things.

Consider, now, thoughts that are directed at the past. Assuming that episodic memory is a form of episodic thinking, as De Brigard (2014a) does, we have, on the one hand, successful and unsuccessful occurrences of remembering, which we have discussed already, and, on the other hand, *past-oriented counterfactual thinking*. The latter refers to cases where we think about events that could have been the case in the past. For example, when I think about how my holidays would have been in 2010 had I saved money that year, I am entertaining a thought about an event that could have happened in the past if I had done something differently. Similarly to episodic future thinking and future-oriented counterfactual thinking, the event in question does not exist and the properties are non-instantiated. Thus, despite presenting me with a particular event and some seemingly instantiated properties, the thought fails to refer to existing things.

Now, to see how we can distinguish between those occurrences of episodic thinking, including successful and unsuccessful occurrences of remembering, remember that, for hybridism, the content of a given thought is partly determined by its phenomenology *and* partly determined by how it relates to the things it is about. These are, respectively, the phenomenological particularity and the relational particularity of mental states. The claim hybridism puts forward with respect to the relationship between episodic memory and different forms of episodic thinking is that they are continuous in terms of their phenomenological particularity, but discontinuous in terms of their relational particularity. So, we need an account of the relationship between episodic memory and episodic thinking that explains (1) how all those forms of thinking instantiate phenomenological particularity; (2) how occurrences of remembering instantiate full or partial relational particularity; and (3) how forms of episodic thinking differ from each other in terms of content.

While the difference of content of successful and unsuccessful occurrences of remembering is due to their relational particularity, or whether and how the modes of presentation refer, the difference of content of different forms of episodic thinking is to be found in how phenomenological particularity is instantiated, or in the kinds of modes of presentation that

compose their contents. The motivation for this is quite simple: because different forms of episodic thinking are considered to be fully gappy on the hybrid account, the differences in their contents should come from the parts responsible for their phenomenologies. To make sense of these differences, however, we need an account of what kinds of modes of presentation there are and which ones are present in different occurrences of episodic thinking.

When we look at the phenomenology of the various forms of episodic thinking described above, it is possible to identify two important differences among them. First, they differ with respect to whether the events are presented as being about the past or the future and whether they are presented as being actual, possible, or not possible. And second, they differ in relation to whether the properties are presented as instantiated or not. This allows us to distinguish between four kinds of modes of presentation for events and two kinds of modes of presentation for properties to explain the differences in the phenomenology of those thoughts.

Events	
Kinds of MOPs	Temporal location / Actuality or possibility
$MOP_{past/actual}$	Past / Actual
$MOP_{past/n-possible}$	Past / Not-possible
$MOP_{future/possible}$	Future / Possible
$MOP_{future/n-possible}$	Future / Not-possible

Properties	
Kinds of MOPs	Instantiation
$MOP_{instantiated}$	Instantiated
$MOP_{n-instantiated}$	Non-instantiated

Let us now consider this in relation to the content of episodic thinking. To start, consider remembering, misremembering, and confabulating. In all of them, the modes of presentation present subjects with events that seem to be past and actual and with properties that

seem to be instantiated by those events. While only remembering is successful in referring to actual past events and to instantiated properties, it is still the case that it *seems* to subjects that, in remembering, misremembering, and confabulating, their mental states are about such things. So, we can say that the modes of presentation composing the content of remembering, misremembering, and confabulating are modes of presentation presenting events as being past and actual and modes of presentation presenting properties as being instantiated. This leads us to a more fine-grained characterization of the content of those mental states:

Remembering: $[MOP_{past/actual}(e); MOP_{instantiated}(p)]$

Misremembering: $[MOP_{past/actual}(e); MOP_{instantiated}(_)]$

Confabulating: $[MOP_{past/actual}(_); MOP_{instantiated}(_)]$

Note that this characterization is compatible with the account given in section 5.2. Remembering, misremembering, and confabulating still differ with respect to their contents, but they have the same phenomenology because their contents are composed by the same kinds of modes of presentation.¹¹

In comparison, consider now episodic future thinking, future-oriented counterfactual thinking, and past-oriented counterfactual thinking. In cases of episodic future thinking (EFT), the modes of presentation present subjects with events that seem to be future and possible and with properties that seem to be instantiated by those events. In cases of future-oriented counterfactual thinking (FOCT), the modes of presentation present subjects with events that seem to be future and non-possible and with properties that seem to be instantiated by those events. In cases of past-oriented counterfactual thinking (POCT), the modes of presentation present subjects with events that seem to be past and non-possible and with properties that seem to be instantiated by those events. In all those cases, because the events

¹¹It might be argued here that the non-gappy modes of presentation in remembering and misremembering are different from the gappy modes of presentation in confabulating. In response, I want to clarify that when I say that the modes of presentation in remembering, misremembering, and confabulating are of the same kind, I mean that *how* they present events to the subject is the same, regardless of whether or not they are successful in referring. Thus, because modes of presentation are responsible for the phenomenology of memory and because the phenomenology does not necessarily change when modes of presentation are fulfilled, it is not incoherent to say that remembering, misremembering, and confabulating can have the same kind of modes of presentation.

do not exist, the contents are fully gappy. Thus, we have the following characterization of the contents of those mental states:

EFT: [MOP_{future/possible}(__);MOP_{instantiated}(__)]

FOCT: [MOP_{future/n-possible}(__);MOP_{instantiated}(__)]

POCT: [MOP_{past/n-possible}(__);MOP_{instantiated}(__)]

Now that we have a characterization of the content of different forms of episodic thinking, we can see whether hybridism satisfies (1)–(3) above. (1) establishes that a hybrid account of the relationship between episodic memory and episodic thinking must explain how they instantiate phenomenological particularity. In the characterization of the content of episodic thinking given above, this requirement is satisfied. Because the contents of different forms of episodic thinking contain modes of presentation, we can explain why, like in occurrences of remembering, it seems to subjects that their thoughts are about particular events with particular properties. Condition (2) requires, in contrast, an account of why some occurrences of remembering instantiate full or partial relational particularity, while other forms of episodic thinking do not. The hybrid account provided also satisfies (2): remembering instantiates full relational particularity because both modes of presentation are successful in referring; misremembering instantiates partial relational particularity because only the mode of presentation of events establishes reference; and confabulating and other forms of episodic thinking are fully gappy because none of their modes of presentation are successful in referring. Finally, (3) requires an explanation of how different forms of episodic thinking and confabulating differ in terms of content, given that they are all fully gappy. This is also accommodated in the framework above, for despite being fully gappy, the content of confabulating and other forms of episodic thinking differ because they are constituted by different modes of presentation. This, again, is fully consistent with the notion of hybrid content developed in section 5.2, for phenomenology partly determines the content of mental states. Thus, taking different forms of episodic thinking to have hybrid contents preserves important continuist and discontinuist intuitions.

5.4 Objections

One assumption that underlies the proposal of this paper is that the dispute between representationalism and relationalism in perception and memory have sufficiently similar elements, such that the enterprise to analyze hybrid views of perception to start thinking of hybrid views of memory is justified. In this section, I will briefly consider some differences between memory and perception that could serve as bases for objections against my proposal. I also consider one objection which says that hybridism does not succeed in avoiding disjunctivism.

5.4.1 The character of memory vs. the objects of memory

It might be argued that there is an important difference between the contemporary dispute between representationalism and relationalism about perception and the same dispute in memory. Unlike in the philosophy of memory, contemporary philosophers of perception are concerned with the *character* of perception, as opposed to the nature of its objects. Although how exactly to characterize the objects of perception across different modalities is a matter of dispute (see, e.g., O’Callaghan 2011, 2016), it is widely agreed that they are mind-independent things or events in the environment. What representationalists and relationalists ultimately disagree about is whether perception is, fundamentally, a matter of representing the world or of being directly related to it. Representationalists, as Tim Crane (2006) points out, hold that “a perceptual representation need not essentially involve a relation to what it represents” (133), meaning that there can be instances of perceptual experiences that do not relate to anything, such as hallucinating seeing a unicorn (see, e.g., Tye 2000; Byrne 2001; Dretske 2003). In contrast, relationalists insist that there cannot be perception without a relation (see, e.g., Martin 2004; Brewer 2007; Fish 2009), which motivates, in part, their appeal to disjunctivism to deal with the occurrence of hallucinations.

While it is true that discussions about the character and the objects of perception and memory are two different things, it is also true that they are closely related. It is not possible to give a proper account of what the objects of memory are if we do not have a proper

understanding of the character of memory. Hybridism is, in this perspective, an attempt to show that the dispute between representationalism and relationalism about the nature of the objects of memory can be resolved by adopting an appropriate (i.e., a hybrid) view of the character of memory. In other words, if the character of memory is hybrid, as hybridism suggests, the objects of memory are indeed the past events themselves, but we only become aware of those events by undergoing representational states whose contents are inherently relational. Recently, Bernecker (2008) seems to hint at a similar view when he says that

Though remembering something may require the having of memory-data, there is no reason to suppose we are aware of these memory-data themselves. I am aware of a past event by internally representing the event, not by being aware of the internal representation of the event. Memory-data do not function as the primary objects of awareness, but are merely the vehicles of the remembered information. Memory is indirect in the sense that it involves a series of causal intermediaries between the past event and the memory experience (memory-datum). But from this it does not follow that memory is indirect in the sense of involving a prior awareness of something other than the past event. (75)

Memory-data, or, as I have been using the term, representations, are the vehicles by means of which we become aware of past events. Despite his account being compatible with hybridism, Bernecker does not say much about what those memory-data should be in order to make it possible for one to be aware of past events themselves by means of representing them. Hybridism, in contrast, deals with this question directly. In doing so, it provides an account of what the character of memory should be in order to reconcile representationalist and relationalist intuitions, such that a common ground about what the objects of memory are can be found.

5.4.2 The temporality of the objects of perception and memory

Another disanalogy between perception and memory relates to the *temporality* of their objects. The objects of perception are co-temporal with perceptual experiences, which makes

it easy to see how they can be constitutive parts of their contents. However, the objects of memory are not co-temporal with memorial states. Past events no longer exist when we remember them, so it is hard to see how they can be constitutive parts of the content of remembering. Thus, perhaps it is simply misleading to say that memory has hybrid contents in the same way that perception does.

As I pointed out in section 5.3.1, even if it is true that this disanalogy is genuine, it does not threaten hybridism. Alternatively, it is possible to question whether the disanalogy is actually the case. As Bernecker (2008) notes, “also in the case of perception we have to allow that what is directly perceived is not contemporary with the act of perceiving it”, for “perceiving a physical object is a causal process that takes time” (69, see also Russell 1912, 17–8). Thus, strictly speaking, the objects of perception are not co-temporal with perceptual experiences, for the causal processes leading up to perceptual experiences require time to happen. Although the time separating remembering and past events is significantly longer than the time separating perceiving and its objects, the difference here is one of degree and not one of kind. So, if there is a disanalogy between perception and memory, it is not as dramatic as it initially appeared.

5.4.3 Auto-noetic consciousness

A third difference between perception and memory is that the latter seems to involve a unique kind of consciousness, namely, what Tulving (1985b; 2002) called *auto-noetic consciousness*, or simply *auto-noesis* (see also Klein 2015). Broadly speaking, one might understand auto-noesis in two ways. On the first understanding, auto-noesis refers to what philosophers usually call the “feeling of pastness” (Russell 1921, 161–2) associated with episodic memories. On the second understanding, it refers to the sense of self or “ownership” that episodic memories carry with them (see Klein and Nichols 2012). Although Tulving distinguished these two understandings later in his works, the initial characterization provided by him, according to which auto-noesis refers to the sense of self in subjective time, seems to suggest that auto-noesis involves both the elements above. Since this is perhaps the most common definition of the term, I will stick to it here.

By relying on a characterization of content inspired by perceptual experiences, it might be argued that hybrid contents fail to account for this essential feature of memory. Although some, such as Fernández (2016), argue for a characterization of auto-noesis in terms of content, it is commonly accepted that auto-noesis belongs to the phenomenological dimension of memory. Thus, when we take into account hybridism's commitment to separatism, it is not required that auto-noesis be an essential element of the content of memory. It is consistent with hybridism that differences in the phenomenology of a mental state need not imply differences in its content, as long as those differences are explained by something external to the content. Another alternative would be to suggest that auto-noesis is not a feature of episodic memories themselves, but of doxastic states accompanying them, which are "recruited" by particular elements of their contents (Sant'Anna and Michaelian 2018). For example, we can say that your memory of your tenth birthday party has the feelings of pastness and of ownership not because they are built into the content or into the phenomenology of memory, but rather because you hold certain beliefs that accompany your memory, such as that the child represented in the thought is you, that events in which you are a child are in the past, and so on. Another alternative would be to explain auto-noesis as arising out of the functioning of metacognitive processes responsible for detecting specific cues present in the content of memory (see, e.g., Dokic 2014).

Given that there is no agreement as to what auto-noetic consciousness is, and given that this is still a controversial topic in the literature, I do not want to commit to any particular alternative here. A full hybrid account of episodic memory will, of course, be required to provide an account of auto-noesis, and some of the alternatives described above are potentially compatible with the hybrid view. However, showing that hybridism can in fact rely on them, or whether a new alternative is needed to explain auto-noesis, is the task of a future project.

5.4.4 Disjunctivism again?

To conclude the paper, I would like to reply to one objection that can be raised to hybridism. This objection says that hybridism fails to avoid disjunctivism, for the framework I pro-

vided is logically compatible with the disjunctivist claim that successful and unsuccessful occurrences of remembering are only similar in relation to their phenomenology.¹²

This objection overlooks two important points of hybridism. The first point refers to the fact that hybridism is not committed to the claim that the only thing shared by successful and unsuccessful occurrences of remembering is the phenomenology. As it became clear in the discussion of the content of episodic thinking, phenomenology is only a part of the content of remembering and episodic thinking. So, successful and unsuccessful occurrences of remembering are similar with respect to their phenomenology because they share a more basic “common factor”, which is that they are mental states with representational content. So, those mental states are similar in a more fundamental way, which actually explains *why* they share the same phenomenology, thus making hybridism incompatible with the kind of metaphysical conclusions advanced by disjunctivists.

The second point is that the objection overlooks the fact that hybridism is not concerned with the logical coherence of disjunctivism, but with the metaphysical conclusions drawn from the claim that the only thing shared by successful and unsuccessful occurrences of remembering is the phenomenology. Even if, for the sake of the argument, we grant that hybridism is compatible with this claim, it does not follow that it is committed to disjunctivism, or at least to disjunctivism in the way described in section 5.1. The problematic aspect of disjunctivism, from the hybridist’s point of view, is that the disjunctive claim supports the conclusion that successful and unsuccessful occurrences of remembering belong to two different metaphysical kinds. Since this claim is not implied by hybridism, it is not the case that it is committed to disjunctivism even if we accept that the only thing shared by successful and unsuccessful occurrences of remembering is the phenomenology.

¹²I would like to thank Markus Werning for raising this objection.

Conclusion

This thesis consisted of five self-standing articles dealing with issues surrounding the dispute between representationalism and relationalism about perception and memory. The overall argument of the thesis was that a form of hybridism, where elements from each of these views are integrated, is to be preferred in both domains. However, because of the self-standing nature of the chapters, the argument offered for hybridism was not linear. As a conclusion, I would like to say a bit more on how, despite the non-linearity of the argument, the thesis still offers a consistent defense of hybridism.

Let me start with perception. Part I, which dealt exclusively with perception, consisted of two chapters. Chapter 1 built on pragmatist ideas to develop an unconventional relational theory of perception. This theory, which I called *perceptual contextualism*, is an unconventional form of relationalism because it advocates for a different understanding of the objects of perception, where they are viewed as the impositional elements of perceptual experiences, that attempts to avoid disjunctivism. While I did not recognize this at the time Chapter 1 was written, and despite my resistance to acknowledging an objection that made this point (see Sect. 1.3.3), perceptual contextualism was too permissive to be an “austere” relational view, that is, a view that explains all the fundamental aspects of perception only in terms of its relational aspect (see Schellenberg 2010). This only became clear, however, when I attempted to provide a more solid grounding of my approach to perception in the classical pragmatist proposal of Charles Peirce.

This realization resulted in the production of Chapter 2, where Peirce’s theory of perception is put to work. By developing what I called *perceptual pragmatism* — which can reasonably be seen as an updated version of perceptual contextualism — the hybrid aspect

of my approach finally became central. The Hybrid Triad, which was introduced in Chapter 2, guided my quest to find a suitable hybrid account of perception based on pragmatism. While, at least initially, the option to focus on Peirce's pragmatism was motivated primarily by my previous interests in his philosophy, during the development of the chapter, it became clear to me that pragmatism could also provide the starting point to explore a new territory in the perception literature. In other words, perceptual pragmatism is one of the first attempts in contemporary analytic philosophy of perception, if not the first one, to tackle more systematically the question of whether there can be a hybrid account of perception that is not biased towards representationalism.¹³ In this sense, my pragmatist-inspired hybrid approach offers an important contribution to the hybrid literature in that it shows that, at least in principle, there can be hybrid views of perception that are more sympathetic to austere versions of relationalism. So, while only Chapter 2 explicitly advocates for a hybrid account of perception, the ideas first developed in Chapter 1 were fundamental to shape the hybrid account of Chapter 2.

Let me now turn to memory. Part II consisted of the work that I did in the philosophy of memory. Despite the change in the subject, the ideas advanced in this part could not have been developed if it were not for the work done in Part I. Chapter 3 was directly motivated by my interests in an alternative account of the objects of perception. Because there was also a need for an alternative account of the objects of memory in the context of episodic hypothetical thought, elements of the pragmatist framework developed in Part I were introduced in this new context. While, again, there was no explicit mention of hybridism in this chapter, the resulting pragmatist account of the objects of memory, and, more generally, of the objects of episodic hypothetical thought, contained elements from both representational and relational accounts of memory. Chapter 4, too, did not address hybridism, but it was a required step towards a proper motivation of hybridism about memory. As I discussed above, relational views are very unpopular in the memory literature. This is because re-

¹³While my hybrid account provides a novel perspective for the contemporary analytic philosophy of perception literature, it is very close to hybrid views developed in the pragmatist literature, especially that of Peirce scholars. As it is clear from my discussion in Chapter 2, my hybrid account relies heavily on Sandra Rosenthal's (2001; 2004) interpretation of Peirce's theory of perception. Similar interpretations that highlight the hybrid dimension of Peirce's view have also been offered by Haack (1994); Bergman (2007); Legg (2014b); Wilson (2016).

lationalists have failed to engage effectively with the current discussions in the area. So, despite Chapter 4 being critical of relationalism, its main goal was not to contribute to its dismissal, but rather to provoke relationalists to engage more explicitly with their critics. Again, the work done in Chapter 4 would also not have been possible if it were not for the work done in Part I on perceptual errors, as the problems that relationalism about perception and relationalism about memory face in relation to error are very similar.

Given the proximity identified between problems of concern to philosophers of perception and philosophers of memory, and given my hybrid take on perception, the question of whether a hybrid approach to memory was possible quickly became central. However, as I discuss in the introduction to Chapter 5, due to the fact that no hybrid views had so far been offered in the memory literature, I ended up deciding for a more conservative hybrid approach to memory. This is why, among other things, the framework used in Chapter 5 was the one provided by Schellenberg's hybrid account of perception, and not the one provided by my own hybrid account in Chapter 2. While this could create a potential theoretical conflict between Chapter 2 and Chapter 5, strategically speaking, the more conservative approach of Chapter 5 was more likely to provide a better support for the overall argument of the thesis. That is, because talking about hybrid views is already a novel topic in the memory literature, relying on a more established hybrid approach to perception, such as Schellenberg's, was more likely to make philosophers of memory sympathetic to entertaining the possibility of hybrid views in memory. Moreover, because hybridism requires taking relationalism seriously, Chapter 5 offers a potentially fresh start for relationalist-minded philosophers of memory, as it offers a framework that does not require abandoning representationalism altogether.

In summary, while not all chapters dealt explicitly with hybridism, they played a fundamental role in building up the required elements for the hybrid accounts of perception and memory offered in Chapters 2 and 5. Thus, although the support that those chapters offer to Chapters 2 and 5 cannot be straightforwardly constructed in a linear way, they still consist in important steps towards reaching the conclusion that hybridism is to be preferred in both perception and memory.

Appendix

Appendix A

Mental time travel and the philosophy of memory

Introduction

The idea that episodic memory is a form of mental time travel has played an important role in the development of memory research in the last couple of decades. Despite its growing importance in psychology, philosophers have only begun to develop an interest in philosophical questions pertaining to the relationship between memory and mental time travel. Thus, this chapter proposes a more systematic discussion of the relationship between memory and mental time travel from the point of view of philosophy. I start by discussing some of the motivations to take memory to be a form of mental time travel. I call the resulting view of memory the mental time travel view. I then proceed to consider important philosophical questions pertaining to memory and develop them in the context of the mental time travel view. I conclude by suggesting that the intersection of the philosophy of memory and research on mental time travel not only provides new perspectives to think about traditional philosophical questions, but also new questions that have not been explored before.

This appendix chapter serves as a complement to various discussions pertaining to memory that appeared in different parts of the thesis. This short paper was written as an invited contribution to a special issue on the topic for the *Unisinos Journal of Philoso-*

phy. While it does not contribute positively to the overall argument of the thesis, the paper might prove useful as an additional resource, as it provides more accessible discussions to important issues that appear in the thesis.

Outline

This chapters considers the implications of the idea that episodic memory is a form of mental time travel to the philosophy of memory. Section A.1 provides a brief background of the emergence of this idea. Section A.2 discusses this idea in more detail focusing on issues of interest to philosophers. Finally, Section A.3 discusses the implications of mental time travel research to the philosophy of memory. I discuss issues relating to the causal theory of memory (Sect. A.3.1), mental time travel and our knowledge of the past (Sect. A.3.2), and the objects and metaphysics of mental time travel (Sects. A.3.3 and A.3.4).

A.1 Introduction (paper)

The idea that episodic memory is a form of mental time travel has played an important role in the development of memory research in the last couple decades. Mental time travel, according to Suddendorf and Corballis (1997), “comprises the mental reconstruction of personal events from the past (episodic memory) and the mental construction of possible events in the future” (133). “The real importance of mental time travel”, they add, “applies to travel into the future rather than into the past; that is, we predominantly stand in the present facing the future rather than looking back at the past” (Suddendorf and Corballis 1997, 147).

Traditionally, memory has been taken to be primarily about the past, in the sense that it allows us to recall things that happened. However, the suggestion that episodic memory is just a form of mental time travel challenges this idea, for “the primary role of mental time travel into the past is to provide raw material from which to construct and imagine possible futures” (Suddendorf and Corballis 2007, 302). These considerations raise a number of important philosophical questions. A first relevant question refers to whether memory requires

an appropriate causal connection to past experiences or events. Since Martin and Deutscher (1966), it has been standard to assume that remembering requires such connection (see, e.g., Bernecker 2008; Debus 2008; Michaelian 2011; Robins 2016a). A second relevant question is whether episodic memory can be a source of knowledge of the past (see Debus 2014; Michaelian 2016c). Since mental time travel into the past, or episodic memory, is in the service of providing raw material to simulate future scenarios, it is not clear whether or under what conditions it can provide us with reliable information about past happenings. A third and more general question refers to the relationship between memory and other forms of mental time travel, such as imagining future events. Because both are a result of similar cognitive capacities, the question of whether they belong to the same metaphysical kind becomes central (see Perrin and Michaelian 2017)

These and other questions have attracted attention from philosophers concerned with memory (see, e.g., De Brigard 2014a; Debus 2014; Michaelian 2016c; Perrin 2016). In this paper, I will explore some of the implications that the *mental time travel view of memory*, as I will refer to it, has to the philosophy of memory. I will start by discussing some motivations to consider episodic memory as a form of mental time travel. Subsequently, I will explore the implications of this idea to the philosophy of memory.

A.2 Episodic memory and mental time travel

Before we discuss the relationship between episodic memory and mental time travel, it will be helpful to first clarify what *episodic memory* is. The term was initially introduced by Endel Tulving (1972), and roughly speaking, it corresponds to the memory system responsible for receiving and storing “information about temporally dated episodes or events, and temporal-spatial relations among these events” (385).¹ So, when you episodically remember an event, your memory contains information about the *what*, the *where*, and the *when*

¹The term *memory* is ambiguous and it might refer to different things, such as one’s *capacity* to remember (e.g., “John has a good memory”), the *cognitive system* responsible for producing memories (e.g., “Your memory is not working well”), or the outputs of that cognitive system, namely the *mental states* that we call “memories” (e.g., “I have a memory of my tenth birthday party”). For my purposes, I use the term to refer both to the cognitive system responsible for producing memories as well as to the individual mental states produced by it.

associated with that event. That is the so-called what-when-where view of episodic memory, or simply the *www view*. Episodic memories, on Tulving's initial formulation, contrast with semantic memories. Those refer to memories about general facts that were not necessarily experienced. For example, when I remember that the Second World War ended in 1945, I am semantically remembering a fact by using language. The semantic memory system, Tulving says, refers to the "organized knowledge a person possesses about words and other verbal symbols, their meanings and referents, about relations among them, about rules, formulas, and algorithms for the manipulation of these symbols, concepts, and relations" (1972, 386). Thus, in contrast to episodic memories, semantic memories do not require the previous experience of the relevant events.

The important thing to note about this definition of episodic memory is that it is primarily based on the kind of information that is processed and stored. And, because of this, it faces some important problems. One such problem refers to the fact that some semantic memories possess the relevant "www" information; for example, my memory that the Waterloo battle was fought in 1815. Thus, it is not entirely clear whether episodic memories and semantic memories can be distinguished solely on the basis of the information possessed by them. Another problem refers to the phenomenological dimension of episodic memories. Remembering a particular event that was previously experienced seems to involve more than the retrieval of information. Episodically remembering seems to have a distinctive phenomenology, involving a "feeling of pastness" (Russell 1921, 161–62) and a "feeling of warmth and intimacy" (James 1890). In other words, besides the information carried, episodic memories seem to make reference to the past ("feeling of pastness") and to belong to subjects in a unique way ("feeling of warmth and intimacy"). For example, when I remember my tenth birthday party, the memory not only presents the event as having occurred in the past, but also as being "mine", in the sense that I seem to *own* the memory.

These and other difficulties have led Tulving to reformulate his first characterization of episodic memory. Later on, he proposed a definition that took into account the phenomenological aspects described above. According to him, besides carrying "www" information, episodic memories involve a unique kind of consciousness, which he called *autonoetic*

consciousness or simply *autonoesis* (see Tulving 1985a, 2005). Autonoesis, Tulving says, “refers to the kind of conscious awareness that characterizes conscious recollection of personal happenings”; that is, it is what makes subjects “aware that the present experience is related to the past experience in a way that no other kind of experience is” (Tulving 2005, 15).²

The definition of episodic memory as involving autonoesis is very important. Because “[t]he act of remembering [...] is characterized by a distinctive, unique awareness of reexperiencing here and now something that happened before, at another time and in another place” (Tulving 1993, 68), remembering makes subjects “capable of *mental time travel*: [...] [a] person can transport at will into the personal past, as well as into the future” (67, my emphasis). So, besides being responsible for the unique feeling associated with episodic memories, autonoesis gives subjects a more general capacity to “travel” in subjective time. This is not difficult to motivate on phenomenological grounds. As Klein (2015) notes, there is a “perceived temporal symmetry between movements toward (future) and away (past) from the present” (21). To illustrate, imagine that you are thinking about your holidays at the beach next year. Similarly to episodic memories, you have the feeling that the thought is owned by you, in the sense that the holidays are yours and not someone’s else. However, because the event is something that *can* happen, it is presented to you as being “future” to your current thought. Thus, it looks like we can “relocate” ourselves to the future in the same way that we can do it in relation to the past.

The capacity endowed to us by autonoesis to travel both to past subjective time and to future subjective time consists in an important motivation to take episodic memory to be just one form — among others — of *mental time travel*. Despite giving emphasis to phenomenological considerations above, there are also good empirical reasons to endorse

²Although initially characterized in phenomenological terms, there is no agreement as to what autonoesis is exactly. Some have argued, for instance, that autonoesis has an important epistemic value. For example, Dokic (2001; 2014) holds that episodic memory carries a “feeling of knowing”, in the sense that it tells subjects that it originates in their past experiences. Fernández (2016) defends a similar view, but he builds autonoesis into the content of memory rather than in its phenomenology. Quite recently, Mahr and Csibra (2018) have proposed a “communicative” account of the function of episodic memory, in which autonoesis is viewed as being responsible for “[delineating] which of our claims about the past we can assert epistemic authority”. Despite these important developments, I shall take for granted the more standard idea that autonoesis is mainly a phenomenological feature of episodic memory.

this view. In a recent survey, Perrin and Michaelian (2017) discuss similarities between episodic memory and future mental time travel found in different domains. In developmental studies, for example, it has been shown that the children's capacity to remember the past and imagine the future arise at approximately the same time (Suddendorf and Busby 2005; Atance 2008; Fivush 2011). In studies with patients with memory impairments, it has been found that deficits in memory incur similar deficits in the ability to think about future scenarios (Klein et al. 2002; Rosenbaum et al. 2005; Hassabis et al. 2007). Moreover, imaging studies also show that there is a strong overlap in the brain regions associated with episodic memory and future mental time travel (Addis et al. 2007; Schacter et al. 2007, 2012).

I will not attempt to review the relevant literature here.³ I shall, instead, point to an important development of the *mental time travel view of memory*. More recently, some researchers have suggested that the primary function of mental time travel is not to allow us to remember the past. Suddendorf and Corballis (1997), for example, argue that “[t]he real importance of mental time travel applies to travel into the future rather than into the past; that is, we predominantly stand in the present facing the future rather than looking back at the past” (147). In a similar spirit, De Brigard (2014a) says that “remembering is a particular operation of a cognitive system that permits the flexible recombination of different components of encoded traces into representations of possible past events [...] *in the service of* constructing mental simulations of possible future events” (158, my emphasis). And, more recently, Michaelian (2016c) says that “remembering is not different in kind from other episodic constructive processes” (103), thus “[w]hat it is for a subject to remember [...] is for him to imagine an episode belonging to his personal past” (111).

The idea that the primary function of mental time travel is not to remember the past, but to imagine the future, has important consequences. One such consequence is that our common sense conception of memory, according to which memory's function is to store information of what happened, seems to be threatened. It is compatible with the mental time travel view that our representations of the past be inaccurate as long as they are beneficial for future actions. So, as De Brigard (2014a) notes, “many ordinary cases of misremem-

³For a detailed review concerned with philosophical questions, see Perrin and Michaelian (2017).

bering *should not* be seen as instances of memory's malfunction" (158, his emphasis). This raises a further question, which is of particular interest to philosophers, about whether, and if so, how, memory provides knowledge of the past. Because the primary function of remembering is not to recover information about the past, we need a proper account of how knowledge can be formed on the basis of memory. Similarly, the mental time travel view poses important questions pertaining to the relationship between memories and the past events. The causal theory of memory, which has been predominant in philosophy for the past four decades, stipulates that remembering requires the preservation of an appropriate causal connection to past events. However, if memory is a form of mental time travel in the same way that imagination is, and "if imagining need not draw on stored information ultimately originating in experience of the relevant episode" (Michaelian 2016c, 111), there is no principled reason to say that such requirement holds for memory.

In summary, the mental time travel view of memory raises a lot of important questions for philosophers concerned with memory. In an attempt to motivate those problems, I will consider, in the next section, some implications that the mental time travel view of memory has to the philosophy of memory.

A.3 Mental time travel and the philosophy of memory

The mental time travel view of memory not only challenges important traditional conceptions about memory, but also offer prospects for future research on the subject. In this section, I will consider some topics that are of potential interest to philosophers of memory concerning the mental time travel view of memory. However, because the interest of philosophers on these topics is still very recent, there is not a lot of work dealing systematically with the questions that I discuss below. For this reason, rather than attempting to survey the debate, I will try to motivate some problems of potential interest.

A.3.1 The causal theory of memory

After the publication of Martin and Deutscher's seminal paper "Remembering" (1966), philosophers in the analytic tradition started to develop an increasing interest in philosophical questions pertaining to memory. Martin and Deutscher proposed what is now known as the *causal theory of memory* (CTM). The CTM has been very influential and it still shapes, to a large extent, how philosophers think about memory today.⁴ However, if correct, the mental time travel view raises important concerns about the CTM.

The CTM provides us with a set of criteria to determine whether a given mental state counts as remembering or not. For the CTM, a subject *S* counts as remembering an event *e* iff:

1. *S* represented *e* in the past; (*Past representation condition*)
2. *S* has a current mental representation of *e*; (*Current representation condition*)
3. The content of the current mental representation of *e* is sufficiently similar to the content of the past representation of *e*; (*Content condition*)
4. There is an appropriate causal connection between the current representation of *e* and the past representation of *e*. (*Causal connection condition*)⁵

To clarify these points, consider my putative memory of my tenth birthday party. In order for me to count as remembering this event, I need to have experienced it previously. That is the *past representation condition*. Additionally, I need to be able to represent the same event in the present. That is the *current representation condition*. But my past and current representations can only be representations of the same event if their contents are sufficiently similar (the *content condition*); for example, if the contents of both representations contain members of my family and friends, a chocolate cake, etc.⁶ Finally, remembering

⁴For a recent and comprehensive assessment of the CTM in relation to recent developments in the philosophy of memory, see Michaelian and Robins (2018).

⁵This discussion is adapted from Bernecker (2010, ch. 1). See also Bernecker (2015, 302).

⁶Martin and Deutscher (1966) conceive of this similarity in terms of a structural analogy holding between the past representation and the current representation. They say that "the past experience must constitute a *structural analogue* of the thing remembered, to the extent to which he can accurately represent the thing"

requires that my current representation of my tenth birthday be caused, in an appropriate way, by my past representation of the same event (the *causal connection condition*). The requirement for such causal connection constitutes the main novelty of the CTM. Moreover, since it is also the source of the problems that arise in the context of the mental time travel view of memory, I will focus on it more closely.

The causal condition is supposed to rule out cases that, intuitively, we do not count as remembering, but that are allowed by (1)–(3). To see this, consider the case of Kent described by Martin and Deutscher (1966):

A man whom we shall call Kent is in a car accident and sees particular details of it, because of his special position. Later on, Kent is involved in another accident in which he gets a severe blow on the head as a result of which he forgets a certain section of his own history, including the first accident. He can no longer fulfill the first criterion for memory of the first accident. Some time after this second accident, a popular and rather irresponsible hypnotist gives a show. He hypnotizes a large number of people, and suggests to them that they will believe that they had been in a car accident at a certain time and place. The hypnotist has never heard a thing about Kent nor the details of Kent's accident, and it is by sheer coincidence that the time, place, and details which he provides are just as they were in Kent's first accident. Kent is one of the group which is hypnotized. The suggestion works and [...] [Kent] believes firmly that he has been in an accident. The accident as he believes it to be is just like the first one in which he was really involved. (174)

Kent's case satisfies (1) and (2) above, as he had a past representation of the car accident and has a current representation of the same event. Moreover, it satisfies (3) too, for

(191, my emphasis). It is not entirely clear, however, where the structural analogy is to be found. The most natural interpretation seems to be that the content of the past representation must have the same kind of structure as the content of the current representation, but they do not say anything as to what the structure of those contents are supposed to be. Another issue is that it is not clear how much "structural analogy" is required for S to count as remembering. While we do not want to require the content of the past representation to be the same as the content of the current representation, it is hard to find a principled way to determine how much similarity is required. For my purposes, I shall put these worries aside. For a related discussion, see Michaelian (2011) and Michaelian (2016c, 90).

Kent's current representation is sufficiently similar to his past representation. Nevertheless, it seems wrong to say that Kent is genuinely remembering. The reason is that his current representation does not preserve the right kind of causal connection to his past representation. To use Martin and Deutscher's (1966) term, the past representation is not "operative" in producing the current representation. In Kent's case, the operative cause, so to speak, is the hypnotist. For the CTM, then, remembering is not only a matter of getting the details of a past experience of an event right, but also of standing in an appropriate causal relation to that experience.

Besides offering a way to rule out cases not contemplated by (1)–(3), the causal connection condition has also been used to provide a taxonomy of memory. As it stands, the CTM is an answer to the general question of what it takes for a subject to remember. However, there is more than one way in which one can successfully or unsuccessfully remember something, which requires an account of those differences. For example, it is consistent with remembering my tenth birthday party that I get some of its details wrong.⁷ I can correctly remember that my whole family was there and that the party took place at a certain location, but I can simultaneously remember, incorrectly, that I had strawberry cake. In this case, we can say that I am *misremembering* my tenth birthday party. Thus, Sarah Robins (2016a) has recently argued that, given the constructive character of memory (see Bartlett 1932; Schacter et al. 2007, 2012; Michaelian 2011; De Brigard 2014a), we need to appeal to a causal connection between past and current representations to distinguish remembering from misremembering.⁸ In a similar spirit, Bernecker (2017) has suggested that one can only distinguish successful remembering from confabulations (see Hirstein 2005) if one requires that the former, but not the latter, preserves a causal connection to past experiences (see also Robins 2016a, 2017a). The causal connection, therefore, is not only important to provide an adequate analysis of remembering, but also of the different kinds of successful and unsuccessful remembering.

The mental time travel of view of memory challenges the central status given to the

⁷Although, again, how much inaccuracy is consistent with remembering is not entirely clear. See Michaelian (2011) and note 6.

⁸See Michaelian (2016b) for a critique of Robins's proposal and an attempt to provide a taxonomy of memory that abandons the causal connection altogether.

causal connection condition in a theory of memory. As I discussed above, on the mental time travel view, the primary function of memory is not to remember the past (see Suddendorf and Corballis 1997; De Brigard 2014a; Michaelian 2016c). But, if that is the case, then it is hard to see why we should endorse the CTM. There are multiple reasons to think this. One reason is that, as Michaelian (2016c, 111) notes, because other forms of mental time travel need not have such causal connection to past experiences, there is no principled way to require it in the case of memory. This does not mean, of course, that there cannot be such connection, but only that it is not necessary.

Another reason is that, from the perspective of the mental time travel view, straightforward occurrences of remembering would be ruled out by the CTM. The causal connection allows us to preserve the intuition that, in cases such as Kent's, subjects do not count as remembering. However, intuitively we do not seem to require that *all* occurrences of remembering preserve an appropriate causal connection to past events. Consider the following case. Imagine that I experienced my tenth birthday party in the past and that I now have a putative memory of it. I remember my friends and family being there and I remember having chocolate cake. However, suppose that my current representation is not being caused by my previous representation of my tenth birthday party, but rather by two different experiences that involved the relevant elements of my current representation. In this case, the content of my current representation is partly derived from, say, my experience of my ninth birthday party, which was attended by the same individuals, and partly derived from my experience of another party that I attended, where there was a chocolate cake. In this case, there is no causal connection of the sort required by the CTM, but it seems too stringent to say that the subject is not remembering the relevant event only because the content of his current representation is not derived from the content of the original experience.⁹

A third reason why the mental time travel view challenges the CTM is that the latter is incompatible with the constructive character of mental time travel. Because mental time

⁹One might argue here that, intuitively, the case above does not count as a straightforward occurrence of remembering rightly because there is no causal connection. I do not mean to dispute people's intuitions about this and other similar cases, but, as long as we want our intuitions to be compatible with what empirical research tells us about memory, this seems the most plausible way to describe them. In other words, given the *constructive* character of memory (see, e.g., Bartlett 1932; Schacter et al. 2007, 2012; Michaelian 2011; De Brigard 2014a), it is not unlikely that cases as the one described above can happen.

travel is in the service of simulating events to assist subjects in future interactions with the environment, it seems too restrictive to require that our representations of the past have to draw content from only one singular source. For example, in thinking about how I should act in my job interview next week, my current representation of the past will benefit more from drawing on different past experiences of job interviews than drawing on only one singular experience.¹⁰

In sum, the CTM has occupied a central position in philosophical theorizing about memory for the past fifty years. Besides providing an analysis of remembering that accounts for a wide range of cases, it provides a useful principle to conceive of a taxonomy of remembering. However, if the mental time travel view of memory is right, the centrality of the CTM might not be warranted.

A.3.2 Mental time travel and our knowledge of the past

One direct consequence of abandoning the causal condition can be seen in the epistemology of memory. Because the causal condition is no longer necessary to remember, there is no guarantee that the content of our current representations derive from the content of our past representations. That being the case, the question that poses itself is whether, and if so, how, we can form knowledge of what happened in the past on the basis of our current representations. Is mental time travel capable of providing such knowledge? Before I turn to this question, it is important to distinguish between two senses in which it can be asked. On the one hand, we can ask the pragmatic question of whether memory provides us with information that, in practical contexts, allows for useful inferences about how things were in the past. Call this the *pragmatic epistemic question*. On the other hand, we can ask whether memory actually provides knowledge of the past, in the sense that it serves as grounds for our justified beliefs about it. Call this the *strict epistemic question*.

This distinction is important because a positive answer to the pragmatic epistemic question does not necessarily give us a positive answer to the strict epistemic question. It might be the case that the content of my memory of my tenth birthday party is the same or very

¹⁰See, however, Sutton (1998) and Michaelian (2011) for different attempts to provide a causal view compatible with the constructive character of memory. For a related discussion, see Robins (2016b; 2016a).

similar to the content of the memories that other people have of this event, such that I can make useful inferences about the event in relevant contexts, but it does not follow from this that my memory allows me to know anything about this event. An answer to the strict epistemic question, in contrast, requires identifying what makes it possible that our current memories serve as grounds for our justified beliefs about the past.

The causal condition provides an answer to the strict epistemic question. Because the content of my current representation of an event is caused by my past representation of it, the causal connection makes it possible for memory to ground our knowledge of the past. Otherwise put, the beliefs that we form on the basis of memory are justified because there is an appropriate causal connection between memories and past events. However, if, as the mental time travel view suggests, this condition is not necessary for remembering, how can we explain the relationship between the content of our past and current representations?

It is not entirely clear what the alternative for defenders of the mental time travel view are here. In fact, because he is the most systematic critic of the causal condition, Michaelian (2016c) has been the only one so far to provide an explicit treatment of the question. His approach consists in adopting a broad reliabilist framework in epistemology, according to which “the epistemic status of a belief is determined by the reliability of the process that produced it” (Michaelian 2016c, 39, see also Goldman 2012). Roughly, the idea is that one is justified in holding a certain belief if that belief was produced by a reliable process. On Michaelian’s proposal, then, we can explain why memory serves as grounds for forming knowledge of the past in terms of the reliability of its underlying processes. This solution, however, will not be appealing if one is not already inclined to a form of reliabilism. The reason is that, as Michaelian (2016c, 40) recognizes, it takes reliabilism as a starting point and then proceeds to explain *how* memory is reliable. However, if one is skeptical of the idea that reliability itself can provide an account of epistemic justification, an account of how memory is reliable will not suffice to address the strict epistemic question.

The question of whether reliabilism is a good account of epistemic justification is beyond my scope here. However, given the question at hand about how memory can form knowledge about the past, it might be useful to explore other alternatives. One possible

approach might be to adopt an *eternalist* view of events (e.g., Bernecker 2008). According to eternalism, past events do not cease to exist when they become past. Eternalism is promising because it allows one to say that past events are constitutive parts of memories. To see this, consider an analogy with perception. Relationalists about perception claim that mid-sized objects are constitutive parts of perception, in the sense that I could not have a visual experience of the chair in my office if this object were not there (see, e.g., Campbell 2002; Martin 2004; Brewer 2007; Fish 2009). An important motivation for acknowledging the constitutive role played by objects in perception is that it allows one to explain how they ground our knowledge of the world (see Schellenberg 2016 for a recent discussion). Similarly, it might be argued, acknowledging the constitutive role played by past events in memories allows one to explain how they ground our knowledge of the past (see Sant’Anna 2018a; Chapter 5).¹¹

Eternalism faces important problems. It is not obvious, for example, how our memories can be constituted by events located in a different spatiotemporal location. While it makes room, at least in principle, for that relation to take place by recognizing the existence of past events, an account of how they relate to our current mental representations is still required. The problem is that it is hard to see how such an account would look like. Another problem for eternalism is that it requires us to pay a high metaphysical price to account for how remembering grounds our knowledge of the past. Because we are required to postulate the existence of past events, some might view this solution with skepticism (e.g., Michaelian 2016c, 63).

Another alternative, which I shall call the *pragmatist solution*, is to deny that the pragmatic epistemic question is different from the strict epistemic question. On such view, having knowledge about the past is simply a matter of making useful inferences about how things were back then. Whether or not we have knowledge of the past, the pragmatist will say, depends on how our memories can inform our future behavior. If memories allow for

¹¹Debus (2008) makes the exact same point when she claims that “the Relational Account [of memory] must be true if we accept (as we should) that people can sometimes gain knowledge about the past on the basis of their [memories]” (406–7). However, her account of memory requires postulating a fundamental separation between memory and other forms of mental time travel view, which makes her view unpromising here (see Debus 2014 and section A.3.4).

behaviors that lead to coordinated action with other individuals in relevant settings, such as discussing who attended your birthday party, or more primitively, discussing where food can be found, then that is all that is required to say that we have knowledge of the past. The pragmatist will deny, therefore, that there needs to be, necessarily, a causal connection to past representations, as long as the current representations allow for useful inferences about the past.

The pragmatist solution also faces important problems. The first problem is similar to the one raised above to reliabilism. In other words, it will only look appealing for those who are already inclined to a pragmatist view in epistemology. The second problem is that the pragmatist solution seems arbitrary, in the sense that it seems to imply that our knowledge of the past depends on what certain individuals “agree” to be the case. However, it is not clear who the relevant individuals are in each situation, or even if there is a principled way to identify them. Moreover, the focus on usefulness might lead to counterintuitive results, for a memory might be useful to guide the current behavior of different individuals without being true of the past. In other words, it is completely plausible that subjects might misremember some or all details of an event in a similar way, such that their memory reports agree with each other, but nonetheless fail to effectively describe what happened.

To conclude this part, it seems that an account of how we form knowledge of the past according to the mental time travel view might require some controversial commitments. While these commitments might take place at different domains — e.g., in metaphysics, as in the eternalist solution, or in epistemology, as in the reliabilist and the pragmatist solutions — a convincing answer to this question will inevitably require a proper motivation of those commitments.

A.3.3 The objects of mental time travel

The mental time travel view of memory also raises important questions about the objects of mental time travel. If memory is only one form of mental time travel, then an account of the objects of memory will inevitably depend on a more general account of the objects of mental time travel. Traditionally, philosophers have addressed the question of the objects of

memory in quite some detail. Inspired by Hume (2011) and Locke (1975), *representational* or *indirect realist* views hold that the objects of memory are internal representations of events (see, e.g., Russell 1921; Byrne 2010). *Relational* or *direct realist* views, in contrast, say that the objects of memory are the past events themselves (see, e.g., Reid 2000; Laird 1920; Russell 1912; Debus 2008). Given this framework, one natural suggestion here to address the question of the objects of mental time travel would be to take one's preferred account of the objects of memory and apply it to mental time travel. However, this seems to get things backwards. On the mental time travel view of memory, the mental time travel category is more basic than the category of memory, so we first need an account of the objects of mental time travel, which will only then inform our account of the objects of memory.

The question of the objects of mental time travel has not been addressed in the literature so far. So, there are no established views about it. However, this should not prevent us from thinking about what an answer to the question might look like. One way to start addressing it is to distinguish between different forms of mental time travel. Although this is not always made explicit in discussions on the subject, there is more than one way in which mental time travel into the past and into the future can happen. Besides episodic memory, which refers to mental time travel to past events that occurred, and episodic future thinking, which refers to mental time travel to events that might occur, we also think about counterfactual events located in subjective time (see De Brigard 2014a). For example, I can think about how my life would be right now if I had not gone to college. In this case, I am thinking about an event that could have happened in the past, and that would influence the present, but that is no longer possible. Similarly, I can think about how my life will be in ten years if I had not gone to college. In this case, I am thinking about an event that would be the case in the future if some other event in my past had been different. In both cases, then, I am entertaining thoughts about counterfactual situations oriented to the past and to the future.

The above suggests that an account of the objects of mental time travel needs to take into account not only episodic memory and episodic future thinking, but also forms of episodic

counterfactual thought (see De Brigard 2014a) directed to the past and to the future. This makes the initial question significantly harder, for now we have to explain how things that can no longer be the case can somehow be the objects of our thoughts. One promising line of investigation might be to appeal to the notion of *intentional objects*. Intentional objects, as originally introduced by Brentano (1973), are non-existent objects which are the direct objects of awareness of the mind. Although this is a promising line, no one has pursued it systematically as of yet.¹²

Another alternative might be to look at the traditional accounts of the objects of memory as starting points. While relational views have been defended more consistently in the context of memory, they do not seem to offer promising prospects for a more general account of the objects of mental time travel. The reason is that the objects of mental time travel, except for arguably the objects of memory, do not exist, which makes it impossible for us to be related to them. So, unless one is willing to commit to more controversial metaphysical views, such as the view that there are intentional objects (e.g., Crane 2001, 2013) or some form of modal realism (Lewis 1986), it is not clear whether relational views can be coherently sustained. In contrast, representational views might be more promising. Because the objects that are represented by mental time travel need not exist to be represented, there is no need to worry about the metaphysical status of those events. What is relevant to explain how we are aware of the relevant events are the existence of the representations, which would serve as proxies for the events. It is not clear, however, what the problems for a representational account of the objects of mental time travel would be. Since this question has not been explored in enough detail, it remains to be seen whether representationalism can stand up to a more detailed analysis.

A.3.4 The metaphysics of mental time travel

The consideration of the questions above finally leaves us in a position to consider a more general question about the metaphysics of mental time travel. As we saw, the mental time travel view of memory raises a lot of different issues regarding the epistemology and the

¹²See, however, Crane (2001; 2013) for potentially helpful discussions about intentional objects in philosophy of mind.

metaphysics of memory. But how pressing those questions are will depend on how one sees the category of memory in relation to the broader category of mental time travel. Until now, I have taken for granted that there are good reasons to accept that memory is just another occurrence of mental time travel. However, some philosophers have resisted this view. Debus (2014), for example, argues that memory and future future-oriented mental time travel — or what she calls sensory imagination — are occurrences of different kinds because there are important metaphysical dissimilarities between them.

The debate about the metaphysics of mental time travel is still very recent and, as with some of the other questions above, there are not well-established views in the literature. Despite this fact, I will follow Perrin and Michaelian (2017) here and distinguish between *continuist* and *discontinuist* metaphysical views of mental time travel. Continuists accept that the similarities between memory and other forms of mental time travel support the more general view that they are occurrences of the same kind. Discontinuists, in contrast, say that those similarities are not enough to say that memory and other forms of mental time travel are occurrences of the same kind.

Reasons for endorsing continuism vary. The general motivation, though, seems to stem from different strands of research in the empirical sciences. As I discussed in section A.2, there is a great variety of empirical work that highlight important similarities between episodic memory and mental time travel. Perhaps the most distinctive motivation comes from the fact that mental time travel into the past and mental time travel into the future draw on very similar cognitive resources, which suggests that a common or “core” cognitive mechanism responsible for mental time travel will be eventually identified (Addis et al. 2007; Schacter et al. 2007, 2012). In more philosophical terms, then, we can see continuism as relying on a more naturalistic stance towards the question of the relationship between episodic memory and mental time travel. In other words, for continuists views, because there is a lot of different empirical evidence suggesting that episodic memory is just another occurrence of mental time travel, we should take this evidence seriously when thinking about the metaphysics of mental time travel.

Discontinuist views, in contrast, seem to be motivated by more general a priori con-

siderations about the metaphysics of mental time travel. This is not to say, of course, that discontinuists simply ignore the empirical evidence on which continuism relies.¹³ Instead, they believe that other considerations, such as whether mental time travel establishes an appropriate causal connection to the events in question, are also important to provide an appropriate picture of the metaphysics of mental time travel. Debus (2014), for example, argues that episodic memory and other forms of mental time travel are occurrences of two fundamentally distinct kinds.¹⁴ To support this claim, she says that, unlike episodic memory, other forms of mental time travel fail to put subjects in an experiential relationship with the relevant events. The notion of an experiential relationship is a technical one, which refers to the causal and spatiotemporal relationship that subjects have to the events that their thoughts are about. In episodic memory, this relationship obtains because the relevant events occurred and we can, at least potentially, draw the causal connection between the current memory and the past event. In other forms of mental time travel, in contrast, the relationship does not obtain because the relevant events do not exist.

Besides reflecting different metaphilosophical attitudes towards the same question, the dispute between continuism and discontinuism reflect different commitments taken in relation to the questions discussed in previous sections. Consider the question of whether episodic memory requires an appropriate causal connection to past events. While continuism is compatible with the CTM, it does not give the causal connection condition a central place in its metaphysical theorizing of mental time travel. For continuism, the presence (or the absence of) a causal connection reflects, at best, only a difference of degree between episodic memory and other occurrences of mental time travel. For discontinuists, however, this question is central for the metaphysics of mental time travel. The presence (or the absence of) a causal connection is sufficient to separate two mental occurrences as being of two different kinds.

The same applies to the question of our knowledge of the past and the objects of mental

¹³See, for example, Perrin (2016) for a more modest discontinuist view that takes into account the similarities highlighted by empirical research.

¹⁴Debus discusses only the relationship between episodic memory and future-oriented thinking, or sensory imagination as she calls it. However, since her argument seems to suggest that other forms of mental time travel are equally distinct from episodic memory, I shall not make this distinction here.

time travel. For continuists, like Michaelian (2016c), a proper account of how episodic memory provides us with knowledge of the past can be given by looking at the reliability of the mechanisms that produce memories, which, in turn, do not require causal connections to the past. Thus, the things that make us aware of the past events are the internal representations, which are detachable from those events. In this sense, continuists might be more inclined to adopt a representational view of the objects of mental time travel. For discontinuists, in contrast, episodic memory is capable of providing subjects with knowledge in a way that other forms of mental time travel cannot. This is because it puts us in a relationship to past events, which necessarily involves a causal connection to them, that is not possible by means of other forms of mental time travel. Thus, discontinuists might not be satisfied with a representational view of the objects of mental time travel, as representations of events can occur in the absence of causal connections to the relevant events. A direct realist or relational view of memory (see Debus 2008) will, therefore, seem more appealing for discontinuists, which Debus (2014) recognizes to be central to her discontinuist account.

A.4 Conclusion

The view that memory is a form of mental time travel offers exciting prospects for new research in the emerging sub-field of the philosophy of memory. Traditional views of memory, such as the causal theory of memory, and traditional questions about memory, such as how it provides knowledge of the past and what is the nature of its objects, need to be reconsidered in the broader framework of mental time travel. These questions, however, are inter-related with more general and new questions that arise only in the context of the research on mental time travel, i.e., what the objects of mental time travel are and what is the metaphysical status of those mental states. Thus, the intersection of the philosophy of memory and research on mental time travel not only provides new perspectives to think about traditional questions, but also new questions that have not been explored before.

Bibliography

- Addis, D. R., Wong, A. T., and Schacter, D. L. (2007). Remembering the past and imagining the future: common and distinct neural substrates during event construction and elaboration. *Neuropsychologia*, 45(7):1363–1377.
- Antony, L. (2011). The openness of illusions. *Philosophical Issues*, 21(1):25–44.
- Atance, C. M. (2008). Future thinking in young children. *Current Directions in Psychological Science*, 17(4):295–298.
- Austin, J. L. (1962). *Sense and Sensibilia*. Oxford University Press.
- Ayer, A. J. (1956). *The Problems of Knowledge*. Macmillan.
- Bartlett, F. C. (1932). *Remembering: A Study in Experimental and Social Psychology*. Cambridge University Press.
- Bergman, M. (2007). Representationism and presentationism. *Transactions of the Charles S. Peirce Society*, 43(1):53–89.
- Bernecker, S. (2008). *The Metaphysics of Memory*. Springer.
- Bernecker, S. (2010). *Memory: A Philosophical Study*. Oxford University Press.
- Bernecker, S. (2015). Memory in analytic philosophy. In Nikulin, D., editor, *Memory: A History*, pages 298–315. Oxford University Press.
- Bernecker, S. (2017). A causal theory of mnemonic confabulation. *Frontiers in Psychology*, 8:1207.

- Brentano, F. (1973). *Psychology from an Empirical Standpoint*. Routledge. Translated by A. C. Rancurello, D. B. Terrell, and L. McAlister.
- Brewer, B. (2007). Perception and its objects. *Philosophical Studies*, 132(1):87–97.
- Brewer, B. (2011). *Perception and Its Objects*. Oxford University Press.
- Brewer, B. (2013). Attention and direct realism. *Analytic Philosophy*, 54(4):421–35.
- Broad, C. D. (1925). *The Mind and Its Place in Nature*. Routledge & Kegan Paul.
- Brogaard, B., editor (2014). *Does Perception Have Content?* Oxford University Press.
- Byrne, A. (2001). Intentionalism defended. *Philosophical Review*, 110(2):199–240.
- Byrne, A. (2009). Experience and content. *The Philosophical Quarterly*, 59(236):429–451.
- Byrne, A. (2010). Recollection, perception, imagination. *Philosophical Studies*, 148(1):15–26.
- Byrne, A. and Logue, H. (2008). *Disjunctivism: Contemporary Readings*. MIT Press.
- Campbell, J. (2002). *Reference and Consciousness*. Oxford University Press.
- Chalmers, D. (2004). The representational character of experience. In Leiter, B., editor, *The Future for Philosophy*, pages 153–181. Oxford University Press.
- Chemero, A. (2011). *Radical Embodied Cognitive Science*. MIT Press.
- Cheng, S. and Werning, M. (2016). What is episodic memory if it is a natural kind? *Synthese*, 193(5):1345–1385.
- Cheng, S., Werning, M., and Suddendorf, T. (2016). Dissociating memory traces and scenario construction in mental time travel. *Neuroscience & Biobehavioral Reviews*, 60:82–89.
- Chirimuuta, M. (2017). Perceptual pragmatism and the naturalized ontology of color. *Topics in Cognitive Science*, 9(1):151–171.

- Clark, A. (2000). *A Theory of Sentience*. Clarendon Press.
- Conway, M. A. and Loveday, C. (2015). Remembering, imagining, false memories & personal meanings. *Consciousness and Cognition*, 33:574–581.
- Cook Wilson, J. (1926). *Statement And Inference: With Philosophical Papers*, volume 1 and 2. Clarendon Press.
- Crane, T. (2001). Intentional objects. *Ratio*, 14(4):336–349.
- Crane, T. (2006). Is there a perceptual relation? In Gendler, T. and Hawthorne, J., editors, *Perceptual Experience*, pages 126–146. Oxford University Press.
- Crane, T. (2013). *The Objects of Thought*. Oxford University.
- Crane, T. and French, C. (2017). The problem of perception. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2017 edition.
- D'Argembeau, A., Renaud, O., and Van der Linden, M. (2011). Frequency, characteristics and functions of future-oriented thoughts in daily life. *Applied Cognitive Psychology*, 25(1):96–103.
- De Brigard, F. (2014a). Is memory for remembering? Recollection as a form of episodic hypothetical thinking. *Synthese*, 191(2):155–185.
- De Brigard, F. (2014b). The nature of memory traces. *Philosophy Compass*, 9(6):402–414.
- 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. (2017). Memory causation. In Bernecker, S. and Michaelian, K., editors, *The Routledge Handbook of Philosophy of Memory*, pages 63–75. Routledge.

- Dewey, J. (1958). *Experience and Nature*. Courier Corporation.
- Dokic, J. (2001). Is memory purely preservative? In Hoerl, C. and McCormack, T., editors, *Time and Memory: Issues in Philosophy and Psychology*, pages 213–232. Oxford University Press.
- Dokic, J. (2014). Feeling the past: A two-tiered account of episodic memory. *Review of Philosophy and Psychology*, 5(3):413–426.
- Dretske, F. (1997). *Naturalizing the Mind*. MIT Press.
- Dretske, F. (2003). Experience as Representation. *Philosophical Issues*, 13(1):67–82.
- Fernández, J. (2016). Epistemic generation in memory. *Philosophy and Phenomenological Research*, 92(3):620–644.
- Fernández, J. (2017). The intentional objects of memory. In Bernecker, S. and Michaelian, K., editors, *The Routledge Handbook of Philosophy of Memory*, pages 88–99. Routledge.
- Fish, W. (2008). Disjunctivism, indistinguishability, and the nature of hallucinations. In Haddock, A. and Macpherson, F., editors, *Disjunctivism: Perception, Action, Knowledge*, pages 144–167. Oxford University Press.
- Fish, W. (2009). *Perception, Hallucination, and Illusion*. Oxford University Press.
- Fish, W. (2010). *Philosophy of Perception: A Contemporary Introduction*. Routledge.
- Fivush, R. (2011). The development of autobiographical memory. *Annual Review of Psychology*, 62:559–582.
- Frege, G. (1980). On sense and reference. In Geach, P. and Black, M., editors, *Translations from the Philosophical Writings of Gottlob Frege*. Blackwell.
- Gallagher, S. and Zahavi, D. (2013). *The Phenomenological Mind*. Routledge.
- Gallo, D. A. (2010). False memories and fantastic beliefs: 15 years of the DRM illusion. *Memory & Cognition*, 38(7):833–848.

- Genone, J. (2014). Appearance and illusion. *Mind*, 123(490):339–376.
- Genone, J. (2016). Recent work on naive realism. *American Philosophical Quarterly*, 53(1):1–25.
- Godfrey-Smith, P. (1998). *Complexity and the Function of Mind in Nature*. Cambridge University Press.
- Godfrey-Smith, P. (2010). Dewey, continuity, and McDowell. In MacArthur, D. and de Caro, M., editors, *Naturalism and Normativity*. Columbia University Press.
- Goldman, A. (2012). *Reliabilism and Contemporary Epistemology: Essays*. Oxford University Press.
- Haack, S. (1994). How the critical common-sensist sees things. *Histoire Épistémologie Langage*, 16(1):9–34.
- Haack, S. (2009). The meaning of pragmatism. *Teorema*, 28(3):9–29.
- Hacking, I. (1983). *Representing and Intervening*. Cambridge University Press.
- Hanna, R. (2015). *Cognition, Content, and the A Priori: A Study in the Philosophy of Mind and Knowledge*. Oxford University Press.
- Harman, G. (1990). The intrinsic quality of experience. *Philosophical Perspectives*, 4:31–52.
- Hassabis, D., Kumaran, D., Vann, S. D., and Maguire, E. A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. *Proceedings of the National Academy of Sciences*, 104(5):1726–1731.
- Hausman, C. (2006). Peirce’s semeiotic applied to perception – The role of dynamic objects and percepts in perceptual interpretation. *Cognitio*, 7(2):231–246.
- Hellie, B. (2007). Factive phenomenal characters. *Philosophical Perspectives*, 21(1):259–306.

- Hinton, J. M. (1967a). Experiences. *Philosophical Quarterly*, 17(66):1–13.
- Hinton, J. M. (1967b). Visual Experiences. *Mind*, 120(302):217–227.
- Hirstein, W. (2005). *Brain Fiction: Self-Deception and the Riddle of Confabulation*. MIT Press.
- Holland, R. (1954). The empiricist theory of memory. *Mind*, 63(252):464–486.
- Hookway, C. (2012). *The Pragmatic Maxim: Essays on Peirce and Pragmatism*. Oxford University Press.
- Hume, D. (2011). *A Treatise of Human Nature*. Clarendon Press. Original publication 1739.
- Husserl, E. (1982). *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy — First Book: General Introduction to a Pure Phenomenology*. Nijhoff. Translated by F. Kersten.
- Husserl, E. (1988). *Cartesian Meditations*. Kluwer. Translated by D. Cairns.
- Hutto, D. D. and Myin, E. (2012). *Radicalizing Enactivism: Basic Minds Without Content*. MIT Press.
- James, S. (2017). Epistemic and non-epistemic theories of remembering. *Pacific Philosophical Quarterly*, 98(S1):109–127.
- James, W. (1890). *The Principles of Psychology*. Macmillan & Co.
- Klein, S. B. (2015). What memory is. *Wiley Interdisciplinary Reviews: Cognitive Science*, 6(1):1–38.
- Klein, S. B., Loftus, J., and Kihlstrom, J. F. (2002). Memory and temporal experience: The effects of episodic memory loss on an amnesic patient's ability to remember the past and imagine the future. *Social Cognition*, 20(5):353–379.
- Klein, S. B. and Nichols, S. (2012). Memory and the sense of personal identity. *Mind*, 121(483):677–702.

- Laird, J. (1920). *A Study in Realism*. Cambridge University Press.
- Legg, C. (2014a). Charles Peirce's limit concept of truth. *Philosophy Compass*, 9(3):204–213.
- Legg, C. (2014b). "Things unreasonably compulsory": A Peircean challenge to a Humean theory of perception, particularly with respect to perceiving necessary truths. *Cognitio*, 15(1):89–112.
- Lewis, D. (1980). Veridical hallucination and prosthetic vision. *Australasian Journal of Philosophy*, 58(3):239–249.
- Lewis, D. (1986). *On the Plurality of Worlds*. Cambridge University Press.
- Locatelli, R. and Wilson, K. A. (2017). Perception without representation. *Topoi*, 36(2):197–212.
- Locke, J. (1975). An essay concerning human understanding. In Nidditch, P., editor, *The Clarendon Edition of the Works of John Locke: An Essay Concerning Human Understanding*. Oxford University Press.
- Loftus, E. F. (2005). Planting misinformation in the human mind: A 30-year investigation of the malleability of memory. *Learning & Memory*, 12(4):361–366.
- Logue, H. (2013). Good news for the disjunctivist about (one of) the bad cases. *Philosophy and Phenomenological Research*, 86(1):211–237.
- Logue, H. (2014). Experiential content and naïve realism: A reconciliation. In Brogaard, B., editor, *Does Perception Have Content?*, pages 220–241. Oxford University Press.
- Mahr, J. and Csibra, G. (2018). Why do we remember? The communicative function of episodic memory. *Behavioral and Brain Sciences*, 41, E1.
- Malcolm, N. (1975). Memory as direct awareness of the past. *Royal Institute of Philosophy Lectures*, 9:1–22.

- Martin, C. B. and Deutscher, M. (1966). Remembering. *Philosophical Review*, 75(2):161–96.
- Martin, M. (2001). Out of the past: Episodic recall as retained acquaintance. In Hoerl, C. and McCormack, T., editors, *Time and Memory: Issues in Philosophy and Psychology*, pages 257–284. Oxford University Press.
- Martin, M. (2004). The limits of self-awareness. *Philosophical Studies*, 120(1):37–89.
- Martin, M. (2017). Elusive objects. *Topoi*, 36(2):247–271.
- McCormack, T. and Hoerl, C. (1999). Memory and temporal perspective: The role of temporal frameworks in memory development. *Developmental Review*, 19(1):154–182.
- McDowell, J. (1996). *Mind and World*. Harvard University Press.
- McDowell, J. (2013). Perceptual Experience: Both Relational and Contentful. *European Journal of Philosophy*, 21(1):144–157.
- Meinong, A. (1960). The theory of objects. In Chisholm, R. M., editor, *Realism and the Background of Phenomenology*. Free Press.
- Menary, R. (2009). Intentionality, cognitive integration and the continuity thesis. *Topoi*, 28(1):31–43.
- Michaelian, K. (2011). Generative memory. *Philosophical Psychology*, 24(3):323–342.
- Michaelian, K. (2016a). Against discontinuism: Mental time travel and our knowledge of past and future events. In Michaelian, K., Klein, S. B., and Szpunar, K. K., editors, *Seeing the Future: Theoretical Perspectives on Future-Oriented Mental Time Travel*, pages 62–92. Oxford University Press.
- Michaelian, K. (2016b). Confabulating, misremembering, relearning: The simulation theory of memory and unsuccessful remembering. *Frontiers in Psychology*, 7:1857.
- Michaelian, K. (2016c). *Mental Time Travel: Episodic Memory and Our Knowledge of the Personal Past*. MIT Press.

- Michaelian, K. and Robins, S. K. (2018). Beyond the causal theory? Fifty years after Martin and Deutscher. In Michaelian, K., Debus, D., and Perrin, D., editors, *New Directions in the Philosophy of Memory*, pages 13–32. Routledge.
- Michaelian, K. and Sutton, J. (2017). Memory. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2017 edition.
- Millar, B. (2014). The phenomenological directness of perceptual experience. *Philosophical Studies*, 170(2):235–253.
- Misak, C. J. (2004). *Truth and the End of Inquiry: A Peircean Account of Truth*. Clarendon Press.
- Nagel, T. (1973). What is it like to be a bat? *Philosophical Review*, 83(4):435–456.
- Nanay, B. (2011). Do we see apples as edible? *Pacific Philosophical Quarterly*, 92(3):305–322.
- Nanay, B. (2013). *Between Perception and Action*. Oxford University Press.
- Newman, E. J. and Lindsay, D. S. (2009). False memories: What the hell are they for? *Applied Cognitive Psychology*, 23(8):1105–1121.
- Noë, A. (2004). *Action in Perception*. MIT press.
- O’Callaghan, C. (2011). Lessons from beyond vision (sounds and audition). *Philosophical Studies*, 153(1):143–160.
- O’Callaghan, C. (2016). Objects for multisensory perception. *Philosophical Studies*, 173(5):1269–1289.
- Peacocke, C. (1992). *A Study of Concepts*. The MIT Press.
- Peirce, C. S. (1931–36). *The Collected Papers of Charles Sanders Peirce*, volume 1–6. Harvard University Press. Charles Hartshorne and Paul Weiss, editors.

- Peirce, C. S. (1958). *The Collected Papers of Charles Sanders Peirce*, volume 7–8. Harvard University Press, Cambridge, MA. Arthur Burks, editor.
- Perrin, D. (2016). Asymmetries in subjective time. In Michaelian, K., Klein, S. B., and Szpunar, K. K., editors, *Seeing the Future: Theoretical Perspectives on Future-Oriented Mental Time Travel*, pages 39–61. Oxford University Press.
- Perrin, D. and Michaelian, K. (2017). Memory as mental time travel. In Bernecker, S. and Michaelian, K., editors, *The Routledge Handbook of Philosophy of Memory*, pages 228–239. Routledge.
- Perrin, D. and Rousset, S. (2014). The episodicity of memory. *Review of Philosophy and Psychology*, 5(3):291–312.
- Poo, M.-m., Pignatelli, M., Ryan, T. J., Tonegawa, S., Bonhoeffer, T., Martin, K. C., Rudenko, A., et al. (2016). What is memory? The present state of the engram. *BMC Biology*, 14(1):40.
- Price, H. H. (1936). Memory-knowledge. *Proceedings of the Aristotelian Society*, Supplementary volume 15:16–33.
- Reid, T. (2000). *An Inquiry into the Human Mind on the Principles of Common Sense*. Pennsylvania State University Press. Original publication 1764.
- Robins, S. K. (2016a). Misremembering. *Philosophical Psychology*, 29(3):432–447.
- Robins, S. K. (2016b). Representing the past: Memory traces and the causal theory of memory. *Philosophical Studies*, 173(11):2993–3013.
- Robins, S. K. (2017a). Confabulation and constructive memory. *Synthese*. Online ahead of print.
- Robins, S. K. (2017b). Memory traces. In Bernecker, S. and Michaelian, K., editors, *The Routledge Handbook of Philosophy of Memory*, pages 76–87. Routledge.

- Rosenbaum, R. S., Köhler, S., Schacter, D. L., Moscovitch, M., Westmacott, R., Black, S. E., Gao, F., and Tulving, E. (2005). The case of KC: Contributions of a memory-impaired person to memory theory. *Neuropsychologia*, 43(7):989–1021.
- Rosenthal, S. (2001). The percipuum and the issue of foundations. In Bergman, M. and Queiroz, J., editors, *The Commens Encyclopedia: The Digital Encyclopedia of Peirce Studies*.
- Rosenthal, S. (2004). Peirce's pragmatic account of perception: issues and implications. In Misak, C., editor, *The Cambridge Companion to Peirce*. Cambridge University Press.
- Rowlands, M. (2010). *The New Science of the Mind: From Extended Mind to Embodied Phenomenology*. MIT Press.
- Rowlands, M. (2017). Arguing about representation. *Synthese*, 194(11):4215–4232.
- Russell, B. (1912). *The Problems of Philosophy*. Home University Library.
- Russell, B. (1921). *The Analysis of Mind*. George Allen & Unwin, London.
- Russell, B. (1992). *Human Knowledge: Its Scope and Limits*. Routledge. Original publication 1948.
- Sant'Anna, A. (2018a). The hybrid contents of memory. *Synthese*. Online first.
- Sant'Anna, A. (2018b). Perception pragmatized: A pragmatic reconciliation of representationalism and relationalism. *Philosophia*, 46(2):411–432.
- Sant'Anna, A. and Michaelian, K. (2018). Thinking about events: A pragmatist account of the objects of episodic hypothetical thought. *Review of Philosophy and Psychology*. Online first.
- Schacter, D. L., Addis, D. R., and Buckner, R. L. (2007). Remembering the past to imagine the future: the prospective brain. *Nature Reviews Neuroscience*, 8(9):657–661.

- 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.
- Schellenberg, S. (2010). The particularity and phenomenology of perceptual experience. *Philosophical Studies*, 149(1):19–48.
- Schellenberg, S. (2011). Perceptual content defended. *Noûs*, 45(4):714–750.
- Schellenberg, S. (2014). The relational and representational character of perceptual experience. In Brogaard, B., editor, *Does Perception Have Content?*, pages 199–217. Oxford University Press.
- Schellenberg, S. (2016). Perceptual particularity. *Philosophy and Phenomenological Research*, 93(1):25–54.
- Searle, J. R. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press.
- Senor, T. D. (2014). Epistemological problems of memory. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2014 edition.
- Siegel, S. (2006). Which properties are represented in perception? In Gendler, T. and Hawthorne, J., editors, *Perceptual Experience*, pages 481–503. Oxford University Press.
- Siegel, S. (2010). *The Contents of Visual Experiences*. Oxford University Press.
- Smith, A. D. (2002). *The problem of perception*. Harvard University Press.
- Snowdon, P. (1980). Perception, vision and causation. *Proceedings of the Aristotelian Society*, 81:175–192.
- Snowdon, P. (1992). How to interpret direct perception. In Crane, T., editor, *The Contents of Experience*, pages 48–78. Cambridge University Press.

- Snowdon, P. (2008). Hinton and the origins of disjunctivism. In Haddock, A. and Macpherson, F., editors, *Disjunctivism: Perception, Action, Knowledge*, pages 35–56. Oxford University Press.
- Solymosi, T. (2013). Cooking up consciousness. *Contemporary Pragmatism*, 10(2):173–191.
- Soteriou, M. (2016). The disjunctive theory of perception. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2016 edition.
- Speaks, J. (2009). Transparency, intentionalism, and the nature of perceptual content. *Philosophy and Phenomenological Research*, 79(3):539–573.
- Squire, L. R. (2009). Memory and brain systems: 1969–2009. *Journal of Neuroscience*, 29(41):12711–12716.
- Stokes, D. (2012). Perceiving and desiring: A new look at the cognitive penetrability of experience. *Philosophical Studies*, 158(3):477–492.
- Stokes, D. (2013). Cognitive penetrability of perception. *Philosophy Compass*, 8(7):646–663.
- Suddendorf, T. and Busby, J. (2005). Making decisions with the future in mind: Developmental and comparative identification of mental time travel. *Learning and Motivation*, 36(2):110–125.
- Suddendorf, T. and Corballis, M. C. (1997). Mental time travel and the evolution of the human mind. *Genetic, Social, and General Psychology Monographs*, 123(2):133–167.
- 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(03):299–313.
- Sutton, J. (1998). *Philosophy and Memory Traces: Descartes to Connectionism*. Cambridge University Press.

- Szpunar, K. K. (2010). Episodic future thought: An emerging concept. *Perspectives on Psychological Science*, 5(2):142–162.
- Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press.
- Thompson, R. F. (2005). In search of memory traces. *Annual Review of Psychology*, 56:1–23.
- Travis, C. (2004). The silence of the senses. *Mind*, 113(449):57–94.
- Travis, C. (2017). Deliverances (indirection). *Topoi*, 36(2):229—246.
- Tulving, E. (1972). Episodic and semantic memory. In E. Tulving and Donaldson, W., editors, *Organization of Memory*, pages 381–402. Academic Press.
- Tulving, E. (1985a). *Elements of Episodic Memory*. Oxford University Press.
- Tulving, E. (1985b). Memory and consciousness. *Canadian Psychology/Psychologie canadienne*, 26(1):1–12.
- Tulving, E. (1993). What is episodic memory? *Current Directions in Psychological Science*, 2(3):67–70.
- Tulving, E. (2002). Episodic memory: from mind to brain. *Annual Review of Psychology*, 53(1):1–25.
- Tulving, E. (2005). Episodic memory and autoevidence: Uniquely human? In Terrace, H. S. and Metcalfe, J., editors, *The Missing Link in Cognition: Origins of Self-Reflective Consciousness*, pages 3–56. Oxford University Press.
- Tye, M. (1997). *Ten Problems of Consciousness*. MIT Press.
- Tye, M. (2000). *Consciousness, Color, and Content*. MIT Press.
- Varela, F., Thompson, E., and Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.

Werning, M. and Cheng, S. (2017). Taxonomy and unity of memory. In Bernecker, S. and Michaelian, K., editors, *The Routledge Handbook of Philosophy of Memory*, pages 7–20. Routledge.

Wilson, A. B. (2016). *Peirce's Empiricism: Its Roots and Its Originality*. Lexington Books.

Zahidi, K. (2014). Non-representationalist cognitive science and realism. *Phenomenology and the Cognitive Sciences*, 13(3):461–475.