

Time in Mind:
The Cognitive Science of Temporal Representation

by
Gerardo Viera

M.A., The University of Houston, 2009

B.A., New York University, 2006

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES
(PHILOSOPHY)

THE UNIVERSITY OF BRITISH COLUMBIA
(VANCOUVER)

December 2016

© Gerardo Viera, 2016

Abstract

Philosophers and cognitive scientists have always been interested in how people come to mentally represent time. Surprisingly though, contemporary philosophers have largely neglected the wealth of relevant empirical research coming from neuroscience, computational psychology, zoology and related fields. My dissertation is meant to remedy this neglect by bringing together major strands in the philosophical and empirical literatures on temporal representation in order to show how both fields can mutually benefit one another.

Chapter 1 describes what I call the *temporal coordination problem* and provides the needed philosophical background on mental representation that frames the majority of the thesis. Chapter 2 provides a taxonomy of the general approaches to explaining how animals coordinate their behaviors with the temporal structure of the world around them.

Chapter 3 argues that part of the explanation for how animals come to mentally represent time is through the operation of a genuine sense of time centered on the circadian systems that provides animals with information about the approximate time of day.

Chapters 4 and 5 argue for what I call the fragmentary model of temporal perception – temporal perception is not a unified capacity but is importantly fragmented. Chapter 4 argues that the fragmentary model undermines the central debate in the philosophical literature over *the mirroring constraint*. I conclude that there simply is no single story to be told about how the temporal structure of experience itself relates to the temporal content of experience.

While chapter 4 emphasizes the fragmentary nature of temporal perception, chapter 5 emphasizes the way in which time appears unified in perception and cognition and proposes an explanation of how this apparent unity comes about. Here I highlight how literature more commonly found in the history and philosophy of science on the unitization of measurement actually informs current understanding of the

mind. In particular, I argue that the brain comes to integrate the temporal information encoded in various time keeping devices by *unitizing* time in a manner that parallels how our cultural time keeping practices have unitized time.

Finally, chapter 6 concludes by recapping many of the major conclusions of the thesis.

Preface

All content is the original work of the author.

Versions of chapter 3 have been presented at the Canadian Philosophical Association, American Philosophical Association, and the Society for Philosophy and Psychology. Versions of chapter 4 have been presented at the Philosophy of Time Society and the Southern Society for Philosophy and Psychology.

While it was the aim of the author to reduce redundancies across chapters, the body of this thesis, chapters 2-5, was written as independent research articles, and therefore some redundancies exist. Furthermore, there will be some minor variations in terminology as the terminology chosen in each chapter was shaped by existing literature. However, as each chapter was designed as an independent research paper, any variations in terminology should remain clear within the individual chapters.

Table of Contents

Abstract	ii
Preface	iv
Table of Contents	v
List of Figures	viii
Acknowledgements.....	ix
Chapter 1: Introduction	1
1.1. The representational theory of mind	5
1.2. Theories of content	17
1.3. The temporal coordination problem.....	20
1.4. The mental representation of time	26
1.5. Moving forward	32
Chapter 2: Empirical Timing Models and the Explicit / Implicit Distinction.....	35
2.1. The explicit / implicit distinction in cognitive science	36
2.2. Minimal semantic interpretations and the implicit / explicit distinction.....	41
2.3. Explicit and dedicated timing mechanisms – internal clocks	49
2.4. No dedicating timing mechanism approaches.....	58
2.5. Summing up	70
Chapter 3: The Sense of Time.....	71

3.1. The senses in an information processing architecture.....	73
3.2. Against the sense of time	95
3.3. Towards a sense of time.....	100
3.4. Circadian rhythms and internal clocks.....	105
3.5. The semantics of (internal) clocks	119
3.6. An information-theoretic account	125
3.7. Conclusion	131
Chapter 4: The Temporal Structure of Experience: Against the Atomism / Extensionalism Debate	133
4.1. What is temporal experience?	136
4.2. The classic dialectic and the mirroring constraint.....	139
4.3. The case for extensionalism: introspection.....	145
4.4. The case for atomism	158
4.5. The fragmentary model of temporal perception.....	170
4.6. Conclusion	196
Chapter 5: The Units of Temporal Perception and the Unity of Time.....	198
5.1. The (purportedly) unit-free nature of temporal perception	200
5.2. The cultural unitization of time.....	216
5.3. Returning to Beck and Peacocke	226
5.4. The fragmentary model of temporal perception.....	227
5.5. Event perception	230
5.6. The units of temporal perception	236

5.7. Conclusion	242
Chapter 6: Conclusion.....	244
Bibliography	251

List of Figures

Figure 1: The simplified feedforward architecture.	80
Figure 2: The complex interactive architecture.	82
Figure 3: Common-cause information channel.....	103
Figure 4: Mirroring and non-mirroring experiences.	145

Acknowledgements

I would like to begin by thanking my supervisor, Eric Margolis, for the hours that he spent going over my work and advising me in how to tackle this dissertation and how to navigate the world of philosophy in general. I can't imagine having finished this dissertation without his constant support and guidance. I would also like to thank my committee members, Christopher Mole and Murat Aydede. Their comments, questions, and objections pushed me to significantly improve my work.

The graduate students at UBC deserve quite a bit of acknowledgment as well. Over the years, many of them have had a huge impact on my work. Tyler DesRoches and Jiwon Byun deserve special thanks for all of our conversations over the years. I also want to single out Max Weiss for always being available to talk about the most esoteric philosophical concerns and to help me make sense of my confusions.

A very special thanks goes to Emma Esmaili. Every word in this dissertation was improved by her influence. The dedication and high standards that she set for her own work, and the feedback that she gave me on mine, pushed me to try and do better at every step of the dissertation. Her support and friendship through the years got me through some of the toughest parts of this dissertation and I could not have done it without her.

Parts of chapters 3 and 4 were presented at a number of conferences. I would like to thank my audiences at the Eastern APA, Society for Philosophy and Psychology, Southern Society for Philosophy and Psychology, Canadian Philosophical Association, the Philosophy of Time Society, and the UBC Graduate Colloquia for your questions and comments. I would also like to thank David Sackris, Kathy Fazekas, and Daniel Booth for comments at the APA, Philosophy of Time Society, and CPA. I would also like to thank Raja Rosenhagen, Andre Sant'Anna, and Mark Fortney for comments on my work as part of the Virtual Dissertation Group.

Finally, I want to thank my family for supporting me throughout every step of this process.

Chapter 1: Introduction

Time has an important role in our mental lives. In fact, it has two important roles. The more general of these two roles is that our mental lives exist or unfold *in time*. We have a series of experiences that follow one another in time. Our beliefs and desires change as time passes. For the more reductively minded among us the point can be put at the level of the biological and neurological processes that underpin the mind. As time progresses, the various metabolic and chemical processes that drive the mind will themselves develop through time. This is all to say that our mental lives have their own complex temporal structure.

The second role that time has in our mental lives is as part of the content of a number of different mental states and processes. From navigating a busy sidewalk, interpreting speech, feeling the pain of an awkward pause in a conversation, listening to music, planning your day, wondering about the future, or even smoothly tying your own shoes, it is vital that we be able to keep track of the timing of events. So, it's not just the case that our mental lives have their own temporal structure, but much of our mental lives concern the temporal structure of the world around us. Our minds exist in time and they are also *directed at* or *represent* time.

That time plays these two roles in our mental lives hasn't gone unnoticed. In fact, throughout the history of philosophy, from antiquity through current discussions in philosophy and cognitive science, there has been a tendency to connect these two ways in which time is wrapped up with the mind. Something about how the mind exists in time is supposed to explain how the mind comes to be directed at or about time. Whether this is the case will be a recurring theme throughout the dissertation that we'll return to shortly.

Despite the fact that philosophers and cognitive scientists have been interested in understanding how time relates to the mind, it is surprising how little these two disciplines have interacted with one

another. Yes, philosophers will often discuss empirical results, particularly the existence of temporal illusions, when they are writing about temporal experience. But their engagement with the empirical literature tends to end there.¹ On the flip side, cognitive scientists may give a cursory acknowledgment of general philosophical issues that arise when they discuss experience or mental representation, but again, there simply isn't the deep engagement that you would expect given that both literatures are often investigating the very same phenomena.

This thesis aims to alleviate this lack of true engagement between these literatures by bringing together dominant strands in both the philosophical and empirical investigation of the role time has in the mind. The goal is that by bringing together these literatures we will gain a more complete understanding of the way in which the mind comes to acquire its temporal content and what role, if any, the temporal structure of our mental lives has in explaining the acquisition of that content. To put the point differently, it is by bringing together this literature that we'll gain a better understanding how the mind's being *in* time relates to the mind's being *about* time.

The main body of this thesis consists of four distinct chapters. Chapter 2 provides an overview of the different empirical approaches to understanding how minds might come to produce behavior that is coordinated with the temporal structure of the events in the environment. In laying out this taxonomy of possible explanatory approaches, I also articulate a form of the implicit / explicit representation distinction that is useful for distinguishing between these various explanatory approaches. While chapter 2 provides some of the background for the thesis, chapters 3, 4, and 5 deal more specifically with particular aspects of how animals come to represent the temporal structure of the world around them.

Chapter 3 focuses on the operation of the circadian system and its role in coordinating behavior with the daily pattern of events. In that chapter I discuss the challenges that people have raised for a

¹ Of course there are some notable exceptions to this general rule. For instance, see (Grush 2005, 2006, 2016, Lee 2014a, 2014b, Phillips 2010, 2012; Thompson 2007).

sensory based account of temporal representation and argue that those arguments fail. In fact, I further argue, that given a proper understanding of what sensory systems are we have good evidence for thinking that many animals have a genuine sense of time that is in many ways analogous to the classic senses like vision or olfaction, which provides them with information about the approximate time of day.

In chapters 4 and 5 I engage with what is typically picked out when philosophers and cognitive scientists talk about *temporal perception* or *temporal experience*. That is, in these chapters I focus on the way that animals come to acquire information from various sensory modalities about time at scales from around 10ms to several seconds. In chapter 4, I argue for what I call *the fragmentary model of temporal perception*. According to this model, temporal perception is not a single unified phenomenon, but is instead, a motley assortment of various capacities that allow animals to keep track of the temporal structure of their environment. With this model in hand I argue that the central debate in the philosophical literature on temporal perception, the debate over *the mirroring constraint*, rests on a faulty assumption. There is simply no single answer to whether the temporal content of experience mirrors the temporal structure of experience itself. By looking closely at the various sorts of mechanisms included in the fragmentary model we'll see that there is simply no single answer to be told as to how the temporal structure of perception relates to the temporal structure that is represented by perception. As a result, the central debate in the philosophical literature on temporal perception rests on a mistake. Temporal perception is not a single unified phenomenon and as a result will not admit of a single answer.

While the emphasis of chapter 4 is the fragmentary nature of temporal perception, chapter 5 focuses on the way in which *time is unified* in our perceptual and cognitive lives. We perceive the temporal structure of our world as though time is a single, seamless, and unified framework within which individual events occur, despite the fact that the temporal information we acquire from the world is given to us through a variety of sometimes radically distinct time keeping mechanisms and processes. Somehow, the temporal information encoded by these various *peripheral* time keeping mechanisms needs

to be integrated in order to produce this unified temporal order that we are aware of in perception and cognition and that is required to underpin our complex ways of physically interacting with the world. While the empirical work on this topic is still in its infancy, I argue that a fruitful way of understanding how the brain might overcome this integration problem is as a process of *unitizing time*. The integration problem that the perceptual system faces in many ways parallels the problem that society faced and that motivated the adoption of a system of temporal units for our cultural time keeping practices. By drawing out the parallels we see that the brain has all of the required resources to unitize time in a manner that parallels how we as a society unitized time, and in fact, this unitization would provide the framework for understanding how the brain overcomes its integration problem.

Chapter 3, 4, and 5 were written and designed as independent research papers. However, there are certain themes that run throughout each of the chapters. As mentioned already, one theme that runs throughout the thesis is the question of how the temporal structure of experience, or of our mental life more generally, relates to the temporal content of our mental life. Since, it will be argued, there is no single way in which animals, including humans, come to mentally represent time, there is no a priori reason for believing that once we figure out how we mentally represent some aspect of the temporal structure of our world that we will thereby have an explanation for how we mentally represent other temporal aspects of the world. We have to take on the various components of our capacity to represent time one at a time.

In trying to tease apart the role of the temporal structure of our mental lives in the explanation of the temporal content of mental representation, we will see over and over again appeals to the philosophical literature on mental representation. The question of how the physical stuff that we are made of comes to have states that have semantic properties – i.e. how the physical stuff we are made of comes to realize mental representations that are about the world – is arguably one of the most foundational questions in the philosophy of mind and philosophy of cognitive science. It is right up there with the

question of how consciousness can fit into the natural world. The importance of this literature on mental representation cuts two ways. First, by appealing to the literature on mental representation, and the developed theories that have appeared in that literature, we get the resources we need to understand how the mind's being in time comes to (or fails to) explain the mind's being about time. Second, by looking at the various sorts of mechanisms that seem to underpin our capacity for temporal representation, we will get a clearer understanding of how empirical, and in particularly neuroscientific, findings can inform long standing philosophical debates. That is, by looking at the relevant scientific literature, we see how scientific theories constrain both our attributions of representational states and our explanations for how representations come to have their contents.

In the remainder of this introduction, I'll do two things. First, provide some of the background from the philosophical literature on mental representation that will be appealed to throughout the thesis. Second, I will lay out a neutral way of characterizing how the mind can be about or is concerned with the temporal structure of its environment, in order to provide us with an adequate target for later discussion.

1.1. The representational theory of mind

As so many authors have put it, if we know anything, then we know our own minds.² Yet, despite the fact that we are supposedly so familiar with our own minds, it is mental phenomena that has posed the most difficulty in being integrated into the naturalistic understanding of the world. In particular, there are two aspects of the mind that Descartes pinpointed in his *Discourse on Methods* (1637/1960) that are thought to cause significant problems for a naturalistic understanding of the mind.

² Trying to give an exhaustive list of people who've said this would be impossible. It's simply that widespread. However, the classic statement of this sort of claim is typically attributed to Descartes in his *Meditations* (1641/1993), while the more recent revival of this sort of position can be found in Chalmers (1996).

First, there is the problem of *qualia* or *phenomenal consciousness*. We interact with the world, we gather information about our environment, and in many cases, there *is something that it is like for the subject* to undergo this sort of interaction with the environment. The problem of *qualia* asks for an explanation of why there is anything at all that it is like to go about the world as opposed to there simply being nothing. Why can't we all just go about our business like unconscious zombies?

The second problem for providing a naturalistic understanding of the mind, and the one that will be the primary focus of this dissertation, is the problem of providing a naturalistic account of *intentionality* or *the aboutness of mental states*. Our mental states exhibit intentionality when they are semantically evaluable (e.g. for accuracy, truth, or satisfaction) with regards to the way the world is. Most saliently our beliefs are readily evaluable for truth or falsity depending on whether they accurately depict the way the world is. Our desires are evaluable for whether or not they are satisfied depending on whether what we desire is the case or not.

But not only are our mental states often about the world, but our reasoning and thinking is such that it preserves the semantic relations between our mental states. I can go from the belief that it is raining outside, and the desire that I don't want to get wet, to the belief that I should grab my rain jacket. The question is then, how can a purely physical system operate in such a way that it respects these semantic / rational relations?

Now, many have ultimately come to see these two forms of the mind-body problem as being closely related. Some have seen these problems as so intertwined that solving one will solve (or put us within a short stone's throw from solving) the other.³ But we can focus on the approach to solving the

³ Two prime examples of this come from opposite directions. Strong representationalists (e.g. Drestke (1995), Lycan (1996), and Tye (1995, 2000)) argue that once we have an understanding of how the mind comes to have representational states, then we will have all we need to understand how a mind could be phenomenally conscious. From the other direction, phenomenal intentionality theorists (e.g. Mendelovici & Bourget (2014)) argue that intentionality will be borne out of an account of phenomenal consciousness.

problem of aboutness as the machinery involved in that debate will be of far more importance for our later discussions of temporal representation.

The standard theoretical framework that has dominated the philosophical and scientific literature since the 70's has been the *representational theory of mind* (RTM).⁴ According to RTM, the problem of aboutness and how a physical system could respect the appropriate semantic relations between ideas is to be solved by appealing to a system of internal representations. To avoid making the position sound trivial, it will take a few steps to outline the proposed theoretical framework.

The primary domain for which RTM was developed was for our propositional attitudes. According to RTM, states like *beliefs*, *desires*, *hopes*, etc, the so called 'propositional attitudes', amount to a subject bearing a relation to some semantic entity, *a proposition*. The *attitudinal* component of the phrase 'propositional attitude' simply picks out the different ways that someone might be related to a proposition (e.g. that attitude one takes to a proposition will differ between belief and desire). The insight of RTM was how they account for someone's being related to a proposition.

According to RTM, we are related to a proposition through some internal representational state that bears the proposition as its semantic content, and the appropriate attitude is fixed through the functional profile of the particular instantiation of the internal representation. For instance, to have the *belief that it's raining right now*, is to have an internal representation whose semantic content is the proposition <it's raining right now> and for this representation to be functionally related to other states in a manner that is characteristic of belief (for instance, it might lead you to infer other beliefs such as *it's probably cloudy*). To *desire that it be raining right now*⁵ is to also have an internal representational state whose semantic content is again the proposition <it's raining right now>, but this time the representation

⁴ For classic statements of RTM see Field (1978) and Fodor (1975).

⁵ The grammar of the propositional clause needs to be changed slightly to fit English grammar. However, to be exact I should have said *desire that it's raining right now*.

will have a distinct functional profile that is typical of desire (for instance, it might drive me to curse the bright, blue, sunny sky I see out my window).

What proponents of RTM need to provide is an account of what it is to have an internal representational state that could take on a proposition as its semantic content and what it is for a physical system to instantiate a representational state of this sort with the appropriate functional profile. Let's take each of these in turn.

As Fodor (1975, 1987, 2008) argued, we cannot account for the full richness of the propositional attitudes if we were to take each proposition that can be entertained by a person as involving its own unique and unanalyzable representational state. The primary reason being that we are finite physical beings, so do not have an unlimited storage space to house a unique individual symbol for each proposition we can entertain. Furthermore, our propositional attitudes exhibit two properties, *systematicity* and *productivity*. *Systematicity* is the property of our propositional attitudes whereby the ability to grasp one proposition is systematically related to the ability to grasp other propositions. For example, being able to entertain the belief *that Sally loves dogs* has a systematic relation to the ability to entertain other beliefs, such as *dogs love Sally*. *Productivity* is the property of our propositional attitudes to entertain an *in principle* unlimited number of propositions. Once I can entertain the thought *it's not the case that it's raining* I can also entertain the thought *it's not the case that it's not the case that it's raining*. Simply by reiterating *it's not the case that* I can produce an endless number of novel thoughts. However, beyond these trivial cases of novel thoughts produced through reiteration, I can also entertain novel thoughts that I have never confronted before. With the right resources I can perfectly well understand the sentence "the piebald scorpion rode the L-train to Eastern lands" even though this sentence is one that I have never confronted before in my life.

Fodor argued that in order to account for the systematicity and productivity of thought it must be the case that our propositional attitudes involve complex structured representations that can be combined

in a way that accords with a compositional semantics. All that we need, according to Fodor's version of RTM, *The Language of Thought* (LOT) (1975, 1998, 2008), is that we possess concepts (i.e. the mental representations that are the building blocks of propositional representations) that individually have semantic content, and can be combined in a manner that allows for their individual semantic values to contribute to the semantic value of the larger complex representation.

The major explanatory burden that's then placed on this form of RTM is to provide an explanation for how some neurally implemented *symbol*, i.e. state of a neural system, could come to have its semantic content. In other words, if the mental representations of propositions are composed of concepts, and the concepts are the primary bearers of meaning, then we need an account of what determines the content of the individual concepts.

In the literature, the theory that would provide this sort of explanation goes under the title of a *theory of content* or *psychosemantics*. Since the aim is to provide a *naturalistic* understanding of the mind, one that is broadly in line with our other scientific commitments, then whatever theory of content we provide must be one that accords with what we know about science. And specifically, it cannot appeal to other semantically significant psychological states, as any appeals to other semantically significant psychological states would send us down the path towards a problematic regress. Any semantically significant states that are appealed to by a theory of content would themselves require some sort of theory of content that explains how *those* representational states acquire their content. The very same problem would simply arise again. At some point we need an account of how our mental states come to be about the world that does not rely on there already being mental states that are about the world. So, once we have an account of how the individual symbols of RTM have their content, then it's a matter of designing a physical system whose transitions respect these semantic contents.

The guiding analogy in the construction of RTM has always been a linguistic one. In particular, RTM is nicely characterized by analogy to formal languages. Formal languages possess a distinct

semantics and syntax. The semantics provides a mapping of the individual basic terms of the language to a domain of objects that are the semantic values of the interpreted formal language. This specifies “the meaning” of the symbols. The syntax operates over the non-semantic aspects of the symbols and dictates what inferential rules are allowed within the system and what combinations of the individual symbols compose a well-formed formula or well-formed complex symbol. If it can be shown that the syntax is such that it preserves the semantic properties of the system, e.g. preserves truth, then we will have shown the *soundness of the system*. Taking the analogy on board, then for RTM to explain how we can not only have semantically evaluable thought but reason in a way that respects these semantic relations, what we need is an account of a physical system’s mechanical operations, i.e. its syntax, for which soundness can be shown. We know from the advent of computer design that systems of this sort are possible, just look at how your calculator’s operations respect rules like *addition*, but what we need is an account of how the brain (or whatever the relevant physical system is) can implement a system of this sort. Similarly, the attitudes will be given the same gloss as the *syntax* of the system. We simply need to find a way in which the brain implements a functional architecture that preserves the relevant *attitudes*.

Now, RTM and specifically Fodor’s LOT are specific proposal for providing a naturalistic account of our propositional attitudes. And they were designed with the propositional attitudes specifically in mind. However, researchers in various fields have extended the general theoretical framework of RTM to explain other aspects of the mind. Perhaps the most successful (or at least most influential) extensions of RTM to other aspects of the mind can be found in the theories of visual perception given by Biederman (1987), Marr (1982), and Treisman (Treisman & Gelade, 1980). According to them, vision itself exhibits something that parallels the productivity of thought in that we can in principle perceive an incredibly large range of novel visual scenes and recognize an incredibly large number of objects in the world.⁶ They argue, that in order to account for this sort of productivity in

⁶ I describe this form of productivity in terms of *an incredibly large* range of visual scenes and not *an infinite number* since depending on how we individuate visual scenes there may in fact no be an infinite number of possible

vision we need to posit the existence of a complex system of internal representations that can be combined in various ways to construct complex visual representations.

However, in both the case of the propositional attitudes and in the case of perception, appeals to representational theories have not gone unchallenged. While the various criticisms differ depending on which critic's argument we are considering and whether we are considering the propositional attitudes or perception, there is nevertheless a common core to their arguments. The basic strategy is to provide another explanation of the phenomena in question that does not appeal to any representational states in the production of the phenomena in question. Since positing representations places you in a position where you acquire particular explanatory burdens, i.e. needing to explain how those representations acquire their content, then a theory that does not posit representations will be more parsimonious. Further, many attacks on RTM come from theories that are claimed to provide a *deeper* understanding or more widespread predictive power. For instance, the attack on RTM from *dynamical systems theory* argues that without any appeal to representational states, and by only using the mathematical tools of dynamical systems theory that are used to understand the operation of a number of physical phenomena, one can provide a *deeper* understanding of the origin of behavior (e.g. see Chemero (2000), van Gelder & Port (1995)).

Whether or not dynamical systems theory is able to live up to its own hype is still a matter of debate.⁷ What is important for our discussion is the general approach that the anti-representationalist argument takes. Take some psychological phenomenon and then see if you can provide a non-representationalist explanation of that psychological phenomenon. If you can, then, the inference goes,

scenes. For instance, if we falsely understood visual scenes on analogy to computer displays where there is an array of pixel locations that can take on one of a fixed number of values, then the visual field containing a finite number of pixel locations would only be capable of producing a finite number of distinct visual representations. However, the pull of productivity is still there. The incredibly large number of different visual states is enough to motivate the need for some system of structured representations.

⁷ Although see Eliasmith (1996) for a criticism that dynamical systems theory fails to provide the sort of revolutionary alternative to RTM that proponents of the view suggest.

representationalism is shown to be false. However, again following this sort of reasoning, if, and only if, you cannot provide a non-representationalist explanation, then representationalism is vindicated. Now, in some cases, we may be able to nicely specify some psychological phenomenon and clearly determine whether a non-representationalist theory provides us with an adequate explanation of that theory (i.e. allows us to make reliable predictions). However, there are cases in which it simply isn't clear which theory provides us with a *better* (or the *true*) understanding of some phenomenon as different theories might allow us to make different types of inferences, predictions, or generalizations that all are empirically useful.

Let's take an example from Fodor that he takes as vindicating RTM over a non-representational neuroscientific theory for conceptual thought (1987). Take my behavior that upon speaking with you several weeks ago, I empty my car's trunk and go to the airport today at a specific time. Now consider the non-representational neuroscientific story of my behavior. As Fodor puts it, the neuroscientific explanation would appeal to various neural mechanisms that cause my muscles to contract in just the right way that I turn on the car, drive it across town, at just the right time of day. Since the *specific* neural mechanisms that underpin my movement would be different than the neural mechanisms that underpin your movements in a similar situation, given that our brains likely have significant differences between them, then a neuroscientific explanation of why I went to the airport and why you went to the airport would be different. We would require different explanations for each instance of this sort of behavior, and as Fodor puts it, the similarity in the reasons for why I performed an action and why you performed a similar action would be lost.

However, Fodor continues, a representational story makes both your and my behavior intelligible and allows us to make certain generalizations that the purely non-representational neuroscientific explanation does not allow for. If we allow for representational states, then we can explain our airport directed movement as the result of our having a belief that we should go to the airport at such and such a

time ready to pick you and your luggage up. While the particular neural implementation of these representational states might differ between individuals, it is nevertheless the case that we may have internal representations that have the same content (in much the same way that ‘the sky is blue’ and ‘*the sky is blue*’ differ in their implementation, that is, differ in the sorts of marks on the page, but still share a single content). And this content will have a particular relation to actions, i.e. we will *believe* these propositions. Further, we can generalize from these representational states to conclude that *for any individual if they believe that they should be at the airport at such and such a time, then without any outside hindrances they will go to the airport at that time.*

This sort of generalization, that appeals to mental states that are individuated or characterized by appeals to their semantic content, allows us to understand something about human behavior that simply could not be captured by the non-representational explanation. If we characterize mental states in non-representational terms, then we will invariably appeal to either the means by which those mental states are realized (i.e. the neural machinery involved) or by their functional profiles (i.e. how those mental states, characterized in non-representational terms, relate to other mental states again characterized in non-representational terms). However, it’s unclear that everyone who has the belief *that they should be at the airport at such and such a time* will have the very same neural machinery realizing that state, since not only are there interpersonal differences in the neural wiring found in brains at gross anatomical levels but at the finer grain level the same occurs (Friston et al., 2004). Similarly, as Fodor (1987) has famously argued, there will also be functional differences in how these beliefs are realized in individual people. In one person the belief that one should go to the airport will have the functional properties that it cues that person to constantly check the time. In another person, the same representational state will have a distinct functional profile that disposes them to call a taxi cab. While the functional profiles differ, and thus the mental states characterized at the level of functional roles will differ, we nevertheless treat both token mental states as being *beliefs that one should be at the airport at such and such a time*. However, this

characterization of the mental state is one that appeals to the representational properties of the belief and not merely the local physical or functional properties of the belief.

Now Fodor makes the same inference that his critics do. He finds that representationalism allows for certain generalization that he privileges, and which the opposing theory does not capture, therefore, he takes the representational theory as being *the account of the mind*. But, if our goal in theory choice is to adopt those theories that provide us with the best predictive power, then why should we presume that only one theory is correct?⁸ Insisting on the idea that *one theory is getting at the world correctly* requires that we have some means of evaluating these claims. But any evidence we acquire to test whether a theory is getting at the *real* structure of the world would simply be further empirical data that is either predicted by or not predicted by our theories. Ultimately, the only epistemic grounds by which we have to evaluate our theories is the predictive power of these theories.

The representational story allows for certain predictions, but the non-representational neuroscientific explanation allows for predictions that occur at a smaller scale (and even in some cases that will manifest at behavioral levels).⁹ Of course, if the theories in competition are in direct conflict, then we cannot adopt both theories. However, if the various theories can be reconciled, then there should be no barring adopting both sorts of theories. The inference from the existence of a non-representational theory of some psychological capacity, to the conclusion that representationalism is false, simply doesn't

⁸ The spectre of Carnap (1937) looms large here. Carnap argued that we have a freedom in choosing our theories that is not restricted *simply* by the data. The *principle of tolerance* that he argued for was developed primarily in response to the choice of a formal / mathematical system for describing the world given to us in modern relativistic physics. There is no *true* formalization in describing any aspect of the world. Instead, we can adopt any number of distinct formalizations to describe the world (i.e. we can describe the physical structure of the world given either Euclidean or Non-Euclidean geometry). Instead, the only constraint in our choice of theory or formalization is due to the pragmatic factors in theory choice. Some formalization will simply be more useful than others in that they allow us to formalize empirical data.

⁹ For example, explanations given in terms of the neural implementation of some capacity can make novel predictions about how some capacity would degenerate due to particular sorts of neurological intervention.

hold. Multiple types of explanations are acceptable, provided that each one brings with it some additional explanatory powers and that the theories be compatible with each other.

Now, the situation becomes more complicated when one of the theories is thought to apply to a level of explanation that is supposedly implemented by what the other theory is talking about. In this case, the representational properties of the mind are not thought to be free-floating but are thought to be implemented by physical stuff that makes up the body (presumably, what makes up the brain).

If the representational theory is supposed to be implemented by the mechanisms described by the neuroscientific theory, then it better be the case that the representational capacities we attribute to the organism be realizable in the neural machinery. This much should be obvious. However, for some, namely Fodor (1991), the idea that a representational or psychological theory should be beholden to the underlying neuroscientific theory is simply mistaken. Our psychological theories, or more generally, the special sciences (i.e. anything other than basic physics), are supposed to be *independent* of the theories of the underlying physical mechanisms. Instead of arguing directly against his position, we can take some of the content of the later chapters as direct counterexamples to Fodor's conclusion. The representational states that we attribute in the explanation of behavior are not independent of our theories of the underlying neural machinery. One reason to notice this, and that will be appealed to in the later chapters, is that independent of any commitments as to what the underlying neural machinery involved in some process might be, we can actually attribute radically different representational states to explain the very same behavior. However, by looking at the underlying machinery, we are given further data points that constrain the attribution of semantic content.¹⁰ The representational story might still be underdetermined,

¹⁰ It should be emphasized, however, that nothing in these data points *logically require* a different attribution of representational content. As was argued by Quine (1960) in his famous *gavagai* example, our attributions of semantic content are logically underdetermined by any observable behavior. As Quine argued, by merely noting a foreign language speaker pointing to what we call 'a rabbit' and hearing the person say 'gavagai' as they point, we cannot know for sure if the person intends to pick out with their exclamation the rabbit, a collection of undetached rabbit parts, a particular time-slice of a rabbit, etc. Instead, the point is that the information we learn about the neural implementation of representational states provides us with pragmatic grounds for adopting a more parsimonious attribution of representational states.

like all theories are, but taking into account the mechanisms that do the implementation, we place further restrictions on the viable representational stories.

So, the upshot of this discussion so far is that when we are evaluating whether any given psychological phenomena should be given a representational or non-representational explanation we need to look to whether the representational explanation provides us with explanatory power that is over and above the neuroscientific or biological explanations that must be admitted due to the biological facts that we all must accept. And further, whatever representational story we do provide must be such that it fits with the constraints imposed by the underlying neural machinery.

One more qualification needs to be made. There's a way of putting the debate over representational theories that overly simplifies the situation. And that way of putting the debate is as follows: take some phenomenon in question, provide the characterization of that phenomenon, and then see what explains it. However, as Rey (1997) points out, a significant amount of begging the question can be done by how one characterizes the relevant phenomena. For instance, if I specify some psychological phenomenon as *the ability to express the structured thought about the rains in Vancouver being too frequent*, then this sort of phenomena might seem to require a representational explanation, since the very characterization of the relevant phenomena described the phenomena in representational terms. However, if I specify the phenomena as *the ability to produce the utterance 'the rains in Vancouver are too frequent'*, then it's possible to provide a non-representational account of this phenomena, since the characterization does not beg the question in favor of a representational account. This later characterization of what needs to be explained simply describes the production of certain sounds, but doesn't say that those sounds bear any semantic content.

1.2. Theories of content

We just outlined one way in which the debate between representational and non-representational theories proceeds. However, in this section we turn to a slightly different tactic that is often appealed to in tandem with the approach we just discussed. In the last section we saw that the debate is often evaluated in terms of the explanatory power of the competing theories. However, in addition to that approach, some anti-representationalists attempt to show that adopting representationalism introduces theoretical burdens that cannot be maintained. Specifically, a number of anti-representationalists have argued for their non-representationalist theories by showing that representationalists cannot explain how the individual representations they posit come to have their semantic content.¹¹

The way in which these arguments go is that the critic chooses some particular theory of content, typically the theory that the anti-representationalist thinks is strongest¹², and shows that that chosen theory cannot account for the representations posited by a representational theory.¹³ If the theory of content cannot account for the particular representational content that is in question, then representationalism is discarded. Now, this sort of argument actually has similarities to arguments that we find within the representationalist camp over what is *the* correct theory of content. Amongst representationalists we find theorists proposing counterexamples to particular theories of content by providing cases where it appears that some theory of content *gets things wrong* in that the theory of content would predict some content that is not what the particular representational theory requires. As a result of these sorts of counterexamples, the particular theory of content in question is generally taken to be false *across the board*. What the representationalist infighting and the anti-representationalist arguments have in common is the idea that there will be *a* theory of content, and if *some* theory fails to explain how *some*

¹¹ Examples of this can be found in Chemero (2000), Hutto & Myin (2013), Ramsey (2007), and Stich (1982).

¹² Chemero (2000) focuses on the teleological theory of Millikan (1989). While Hutto and Myin (2013) focus on the information theoretic approach of Dretske (1981, 1995).

¹³ We'll see this sort of argument in many of the later chapters. I'll leave off the explanation of the various types of theory of content until then

representation gets its content, then that theory will simply be false. Perhaps a concrete example should help.

Consider the theory of content given by Jesse Prinz in his book *Furnishing the Mind* (2002). In that book Prinz proposes a theory of content according to which a representation's content is determined by the class of objects (or states) in the world that the representation is causally sensitive to *and* whatever actual object originally caused the production of the initial representational state. According to Prinz, I may have a representational state that gets triggered by both ducks and geese, however, this representation does not have the disjunctive content of *duck or geese* but rather represents *ducks* just in case the original cause of the representation being tokened, or activated, in my mind was a duck and not a goose.

Now, clear counterexamples for this theory are easy to come up with. For instance, I may acquire the concept DUCK through exposure to wooden ducks in a hunting store or at a museum. If Prinz's theory of content were correct, then my concept DUCK would not actually have ducks (i.e. the water loving birds that quack) as its content, but instead would represent replica dummy ducks. This seems to clearly be getting the content of this sort of representational state wrong. But I would caution us from thereby inferring that the particular theory of content should be abandoned. Instead, it may simply be the case that in order to understand how the various representational states of the mind come to have their content that we must appeal to a variety of different types of theories of content. That is, to make sense of our attributions of representational content it may simply be the case that we require different theories of content. However, this is not to conclude that there are distinct types of content. Instead, the point is just that there may simply be multiple explanations for how some representation comes to have its content. Let me explain through an analogy.

Consider the property of *being a parent*. We can give a characterization of what it is to be a parent that fits across the various instances of being a parent. Namely, it is to have a parental relationship

to a child, where *parental relationship* can be cashed out in terms of a long list (likely disjunctive) of the sorts of responsibilities that a parent has. However, there may be various ways in which someone comes to be a parent. You might become a parent through childbirth, through adoption, or even through some other means whereby you simply take on the parenting role. In all of these cases we have rather different stories or explanations for how someone comes to be a parent, but there's no sense in which any of these ways of becoming a parent are in conflict with one another. The fact that one person became a parent through adoption doesn't thereby bar another person from becoming a parent through childbirth.

Now, let's turn back to the matter of representation. What is it to be a mental representation? Well, it takes two things. First, there must be a *representational vehicle*. That is, some physical state of the brain (or some other physical system) that is the bearer of the semantic content. In this way, the representational vehicle is like the linguistic symbol that is just a physical mark on the page. Then, in addition to the representational vehicle, a semantic value must be attached to the vehicle. Take this characterization of what it is to be a mental representation as being analogous to the characterization of being a parent. Now, the distinct theories of content are proposals for how something may come to have the properties that are attributed to mental representations. However, in just the same way that there needn't be a single story about how someone might become a parent, there needn't be any single story about how representations come to have their content.

It should be kept in mind though that we have to be careful here as not any proposed theory of content should be taken as a viable one. Once again, as will become clear in later chapters, what theory of content we adopt is in part constrained by what we know about the underlying neural machinery and the way this underlying neural machinery is related to the world via the behaviors and other mechanisms possessed by the animal. A theory of content should be able to explain how the relevant components of the underlying neural machinery may come to bear certain contents that are attributed to that machinery in the service of providing a predictive and explanatory theory.

1.3. The temporal coordination problem

With all of that ground clearing out of the way, let's turn to the main event. My initial characterization of the topic of interest has been put, and will be put, in terms of how animals come to mentally represent time. But from the previous sections, we saw that this is a question begging way of putting the issue. This isn't an accident but a choice since it provides us with a much clearer initial characterization of what we want to explain. However, at the end of this section I will provide a more neutral reformulation of the problem in terms of how animals coordinate their behaviors with a temporally structured environment.

Most discussions of temporal perception and cognition, at least within philosophy, begin with an appeal to introspection. Simply by reflecting on your experience of the world it should be clear that we are conscious of and / or perceive the temporal structure of the world around us. As you sit there at a sidewalk table, you hear and see the cars passing by at different speeds, some waiting longer at stop signs than others. Pedestrians walk by; some even jog. You hear the music coming from the nearby café. First one song, and then as that song fades into the past, another song comes on. You might even manage to react quickly enough to catch a pen as it rolls off the table. In all of these ways, you simply need to reflect on your experience to notice that you are being presented with a dynamically rich world of temporally structured events.

Moving beyond these perceptual cases, our cognitive lives are also replete with temporal content. You sit there and wonder whether you'll be able to leave work before rush hour begins or whether it will be quicker if you take one route home or another. You might rethink your weekend plans once you see that the weatherman says to expect rain. You might even start reliving old experiences recalling past situations. Again, in all of these cases, reflection reveals the temporal content involved in your mental life.

But notice that if we were to limit our discussion of temporal representation to just those cases in which introspective reflection on our own mental lives reveals temporal content, then we would be doing a great disservice to the role that the mental representation of time plays. As we'll see even in humans, let alone other animals, there is an incredible range of time keeping capacities that simply aren't available to us through introspection.

Let's begin by sticking to two human cases. First, take speech perception. When someone produces an utterance we break down the speech pattern into discrete phonemes. In many cases we take the particular pitches and stops as what would delimit the distinct phonemes that we perceive. However, much of our phoneme perception is the result of a fine-grained sensitivity to the temporal characteristics of the sounds being produced. Buonomano and Karmarkar (2002) have a wonderful example to show this. Take the famous Hendrix lyric "kiss the sky". By simply modifying the temporal pattern of the speech sounds, and without changing any of the other characteristics of the sound, we can change the perceived lyric from "kiss *the sky*" to what would have made for a very different song, "kiss *this guy*"! On reflection, at least for the naïve amongst us, we probably would have explained the perceived difference as there simply being a phoneme difference. However, in order to make this difference our perceptual system must be keeping track of the timing of the speech sound at the scale of milliseconds. The temporal content of speech perception simply isn't worn on the sleeve of speech perception itself.

As another example consider auditory localization. Humans have a fairly decent capacity to judge the location of a sound source based purely on auditory cues. Chief amongst the auditory cues used in spatial localization is a sensitivity to the interaural timing difference brought about by the delay between a sound wave impacting upon one ear versus the other. By keeping track of these very slight temporal differences (along with a host of other differences that are in part due to the shape of the ears) we are able to determine the location of a sound source (Wightman & Kistler 1992). Now, when you reflect on this sort of capacity, it is spatial content that is made available – you *localize* the sound. However, it is only

through empirical investigation that we find that the ability to localize sound in this manner requires the fine grained temporal sensitivity *at the scale of microseconds*.

We can continue this process of empirically uncovering cases where individuals must keep track of time to cases involving non-human animals. It's when we start to investigate how animals coordinate their behaviors with their environment that we uncover a genuine treasure trove of temporal sensitivities. Chapter 3 will discuss the ability of animals to keep track of the time of day in the service of a discussion of the sense of time and the circadian systems, but here we can discuss other cases of time keeping capacities.

When we look at the animal kingdom we find that all animals that need to interact with and successfully navigate their environments must be capable of coordinating their behaviors with the temporal structure of the events around them.¹⁴ Merely being able to act in accordance with *the what* and *the where* of the world simply isn't enough. For animals to move about a dynamic world, to forage properly, to hunt properly, to migrate appropriately, to do a number of different tasks, animals cannot leave it up to chance that they will be at the right place *at the right time*. They must have some sensitivity to the temporal structure of the world.

Now, the truly stunning fact about animal time keeping capacities is that in the animal kingdom we find animals that exhibit a sensitivity to the temporal structure of their environment over a time scale that spans 10 orders of magnitude (Buhusi & Meck, 2005). To give you an idea of the sheer range of this sensitivity consider what it would mean to have a spatial sensitivity over the same scale. That would require that an individual creature be able to discriminate distances as short as the length of an individual

¹⁴ This point extends to all living things. Even plants and bacteria must be able to coordinate their activities with the temporal structure of the world around them. For example, plants must open and close their stoma to coordinate their gas exchange with environmental conditions. If the stomata are opened at the incorrect time, whether this ill-timed opening is due to temporary moisture levels or temporary carbon dioxide levels, then the plant will not survive for long. In this way, the plant itself must have some mechanism that coordinates its activities with the temporal occurrence of events in the environment. Temporal coordination, in this sense, is ubiquitous throughout not just the animal kingdom but throughout all of biological life.

flea to distances as long as the circumference of the planet Earth! All without the aid of any technological tools or written mathematics.

The sensori-motor sensitivities that rely on time keeping capacities should be fairly clear by this point. Animals, just like humans, move about in a dynamic and changing world where the changes that occur are entirely contingent¹⁵. In order to not constantly be bumping into things, animals, again like humans, need some way of keeping track of or coordinating their activities with these goings on. But, animal temporal sensitivities do not end with sensori-motor sensitivities.

Take for instance the feeding and caching behavior of the Western Scrub-Jay (Dally et al., 2006; Gallistel, 1990; Raby et al., 2007; Thom & Clayton, 2013). Scrub-jays exhibit not only the sensori-motor temporal sensitivities that we expect of any creature that physically navigates its environment, but they also exhibit rather surprising capacities for temporal reasoning and prospective planning. Scrub-jays eat a variety of different foods. However, these foods are typically available in cycles. At certain points some food sources, nuts or worm larvae, might be more abundant than other food sources. What the scrub-jay will do in these situations is harvest more than they can eat in one sitting and store the remainder in a food cache somewhere in its territory, thereby exhibiting some future directed concerns. If withholding food from their competitors were their main concern, then they could accomplish this task by simply destroying the food that they do not eat. The scrub-jay seems to genuinely plan for its future well-being.

But, their abilities begin to appear even more impressive when we consider how scrub-jays *harvest* their food caches. They do not simply feed at whichever food cache is closest or according to which is their favorite source of food, as would be the case if they simply stored the information concerning the spatial location of the caches or what type of food was stored. Instead, the scrub-jay will harvest the food that is most likely to spoil quickest (Clayton & Dickinson 1998). Clayton & Dickinson

¹⁵ Whether the world is a deterministic one or not doesn't matter for this point. The events in the world are *epistemically contingent* in that no animal has Laplacian knowledge of their environment, therefore the events going on in the environment are typically not ones that individual creatures could predict.

(1998) showed that scrub-jays will avoid harvesting worm larvae (their preferred food source in the experiment) from their caches, and instead go for the less desirable, but longer lasting, nuts, provided that the larvae were cached *too long ago* (between several hours and a day). Importantly, it was shown that scrub-jays were not responding to the *smell of rotten larvae* but were tracking the length of time different foods were stored. Scrub-jays not only keep track of how long certain food caches will remain usable, but they will also keep track of *when* the food was placed in those caches. We're lucky in that our cell phones and computers tell us the date and our jugs of milk have their expiration dates stamped on their label. Imagine trying to keep track of these things without a calendar and without a best by date. Put in those terms, the temporal sensitivity of the scrub jay is genuinely stunning.

It's here where we run up against a precaution that was raised in the previous section. In all of these cases, the human and the animal cases, we see that animals exhibit a sensitivity to the temporal structure of the world around them. What then is the phenomena that we need to explain? Well, if we want to leave it open as to whether these phenomena require a representational story or not, then we cannot characterize the question we want to answer as *how do animals come to represent the temporal structure of the world around them*. That would be to give a representationally loaded characterization of the phenomena that would likely beg the question in favor of a representationalist story.

However, a more general characterization of the phenomena can be given (and in fact, we just saw this briefly). The world we live in is a temporally dynamic one. At any given moment certain states of affairs obtain at that time and at other times different states of affairs obtain.¹⁶ In order for an animal to successfully navigate its environment, it must have some way of coordinating its behaviors with the temporal structure of the world around it. Call this *the temporal coordination task*. All animals must have some way of overcoming this task. This coordination is what we want to explain. And notice, there is no

¹⁶ This is all I mean by dynamic. Surely this wouldn't satisfy someone with sympathies for McTaggart (1908), however, for the purposes of understanding something about the mind, this form of dynamicism is enough to set up the problem that I am laying out.

begging the question in favor of a representationalist story. Instead, it is left open whether positing mental representations of time for the explanation of this sort of coordination is required.

Now, the interest in noticing the extreme time range over which animals are able to coordinate their behavior with their environment was not just to produce a *wow!* response. Instead, the added purpose of showing that wide sensitivity was to make it plausible that the temporal coordination task, even within a single organism, is not a single task but instead breaks down into a number of distinct capacities to coordinate behavior with different aspects of the temporal world. From a pure engineering standpoint, it's unlikely that any neural mechanism, made of noisy biological materials, would be precise enough to explain how animals can track microsecond timing differences as well as differences that span several days or weeks. Furthermore, there are a number of different empirical cases that show that different timing capacities come apart. At very coarse grained levels, Buhusi and Meck (2005) provide an overview of how sensori-motor timing comes apart from circadian timing. Destruction of the system responsible for many circadian behaviors in mammals, *the suprachiasmatic nucleus*, leaves animals unable to produce regular circadian behaviors but leaves more fine-grained sensori-motor capabilities undisturbed (Lewis et al., 2003). Craver et al (2014) provide evidence that sensori-motor capacities also come apart from the timing capacities involved in memory. The famous case of *Mr. B*, a patient who suffered almost complete loss of the ability to store new memories could successfully navigate his environment, and even engage in complex sensori-motor tasks of drinking a beer (Craver et al., 2014), despite the fact that he could not retain any information for more than approximately 2 seconds. And in chapter 4 and 5 I provide arguments that even at the scale of milliseconds to seconds, the timing capacities of humans are due to a motley assortment of different types of time keeping devices.

The temporal coordination task is many, not one.

1.4. The mental representation of time

We're leaving it open that the explanations for how animals overcome the various temporal coordination tasks they face should involve an appeal to the mental representation of time. In fact, the majority of the next chapter is devoted to discussing the various sorts of approaches that can be found in the literature to explaining how animals overcome the temporal coordination task. It is there that we'll lay out the various representationalist and non-representationalist approaches. However, in this section we'll lay out why some authors have thought that the mental representation of time poses particularly damning problems. Recall that from previous sections it was mentioned that one of the significant hurdles that a representationalist theory of some psychological capacity faces is an explanation for *how the purported representations come to have their content*. That is, the problem comes up as providing a theory of content for the representations in question.

It is in trying to provide a theory of content for temporal representation that *the temporal coordination task* becomes *the temporal coordination problem*. The worry can be put bluntly. Time is an odd aspect of our environment, so we'll need a correspondingly peculiar account of how we come to represent time. But, for a statement of the problem, put with somewhat more finesse, we can turn to a quote from the psychologist Lera Boroditsky. She says,

All of our experience of the world is physical, accomplished through sensory perception and motor action. And yet our internal mental lives go far beyond those things observable through physical experience: we invent sophisticated notions of number and time [...] so how is it possible that physical organisms who collect photons through their eyes, respond to physical pressure in their ears, and bend their knees and flex their toes in just the right amount to defy gravity are able to invent and reason about the unperceivable and abstract?

(Boroditsky 2011, p. 333)

Time, it would appear, is simply an odd aspect of our world. Unlike the medium sized objects that we interact with and their spatial locations, time isn't something that we can readily point to. Yes, we use hand gestures when we speak about time. An impatient parent might point down into their palm when they want something done *right now*. Or a Westerner might point behind them when they talk about the past being *way back then* while an Easterner might point in front of them when they speak about the *past being known*. However, in none of these cases where physical gestures are used alongside speaking about time do we take the gestures to literally be pointing to the temporal features of the world we have in mind. The worry, also, isn't just one where time is abstract and objects and locations are not. We can readily point to concrete instantiations of abstract categories. I can point to a pair of people and take myself to be pointing to an instance of friendship. But again, we simply can't point to time in this way. The temporally extended properties that events possess in our world simply aren't made available to us *at a time* to be able to point to them.¹⁷

Time is also something that we cannot readily manipulate. We can grasp objects and move them around thereby changing their spatial relations to other objects. We can figure out where to hang a picture by putting it *here*, and then *there*, and then back to the first place. We can move ourselves around the room to get different perspective on the layout. We can even control the speed at which we move through space, moving slowly when we want and faster when there's the need. In so far as we seem to be able to control anything, we seem to be able to control at least some of the spatial properties of the world around us.

However, our ability to manipulate the temporal relations of events is much more limited. In some cases, we can delay events or bring them about quicker. I can choose to send an email now or later.

¹⁷ This is in part one of the problems that arises with temporal representation. Many other cases of representation, including pointing, involve a representational vehicle and a represented content that are both entirely present at the moment of representation. However, with temporally extended phenomena, the thing being represented does not exist at the moment of being represented. Rather, we only get access to one part of what we are pointing to.

But this manipulation of the temporal order of events only occurs in the prospective case – we can change the expected or planned temporal order of events. But once an action is performed, we cannot undo it and change its temporal relations to other events. Once events are in the past, there is nothing that we can do to change their location in the past. The analogy to space would require that once we locate an object somewhere, then we cannot move that object to another location. However, nothing like that happens with our ability to manipulate the spatial location of objects. Insofar as an object can be moved, there is no limitation on moving an object back and forth in space, but with time this back and forth simply isn't possible. We also can't change the speed at which time passes.¹⁸

Finally, time is also something that seems to lack any causal influence on the world around us. Time isn't some form of energy that can be transmitted and impact our sensory transducers. Time isn't the sort of thing that we might bump up against as we're walking down the stairs. Time simply doesn't seem to be the right sort of thing to have any causal influence on us. Now, we do sometimes speak as though time is causing certain effects in the world. We might say that the milk went bad *because of how long* it was out on the counter. Or, we might say that the joke failed *because of its being delivered too quickly*. However, a number of philosophers working on the metaphysics of time would caution us from taking these expressions as genuine evidence that time exerts causal influence on the world (Lewis, 1973; Maudlin, 2002; Newton-Smith, 1980). According to these philosophers we should avoid attributing to time genuine causal influences and instead understand time as being a requirement for causal influences to occur at all. In the case of the milk, the genuine causal story, they would suggest, is one that appeals to the metabolic processes of the bacteria found in the milk and the warm temperatures. While these processes unfold in time that does not make time a *cause* of the milk's going bad. Whether these sorts of

¹⁸ An initial reaction to all of this might be to say, well, we can do these things, modern physics allows for all sorts of temporal distortion and varieties of time travel. However, once we look at that literature we find that the ability to change the past is not like how it's depicted in movies like *Back to the Future* or *The Terminator*. Instead, discussions of time travel still involve the acceptance of the idea that *what's done is done*, however, in cases of time travel the stories just become more complicated. For an example of this see Lewis (1976).

arguments that time is causally inefficacious go through, it's nevertheless the case that if time is to be causally efficacious thing in our world, then it does so in a way that is rather different than how other physical things in our world cause things to happen. There would seem to be no medium or vehicle, over and above the physical things in the world that we already believe possess causal powers, for time to influence anything. Temporal causation would seem to be entirely superfluous. Again, time is simply an odd feature of our world.

How then do we explain how animals may come to have mental representation that track such an odd aspect of our world? As we'll see in subsequent chapters, I do not believe that the peculiarity of time poses any genuine problems for providing a theory of content that makes sense of our attributions of temporal content. In fact, in what follows I argue that the resources of information theoretic accounts of content (e.g. Dretske (1981)) are sufficient. However, a number of researchers have doubts that standard stories can be given for temporal representation.

In particular, if we look at the theories that are proposed by psychologists working on the temporal reasoning capacities of humans, we find that they often endorse theories that make use of some rather sophisticated machinery.

Perhaps the most salient of these approaches has its origins in the work of Whorf (1944) and the later work of Lakoff and Johnson (1980). According to the *spatio-cultural* approach to explaining temporal representation our ability to reason about time involves the scaffolding of temporal representations on an existing system for the mental representation of space.¹⁹ Time, they argue, is *spatialized*. The details of the accounts vary from author to author but the general evidence for the view remains the same across writers. Perhaps the most salient of the data points that they take to support their claim is the widespread use of terms with a primarily spatial meaning to speak about temporal relations.

¹⁹ Instances of this sort of approach can be found in Lakoff and Johnson (1980), Lera Boroditsky (2000, 2001, 2011; Boroditsky & Ramscar 2002), Daniel Casasanto and colleagues (Casasanto 2008; Casasanto & Boroditsky 2008; Casasanto et al 2010; Merritt et al 2010), and Prinz (Boroditsky & Prinz 2008).

We speak of *long* or *short* times. We speak of *the past being behind us* and your *future being ahead of you*. Also, the specific ways in which time is spatialized will vary from culture to culture. For instance, native speakers of languages that are written from *left-to-right* will typically associate the past with the left and the future to the right. While native speakers of language that are written from *right-to-left* will typically associate the past with the right and the future to the left (Boroditsky 2011).²⁰

Now given the widespread use of spatial terms for speaking about time and the associations that can be found between spatial and temporal properties, there does seem to be something correct about the spatio-cultural approach to explaining temporal reasoning. However, it's unclear how this sort of approach actually solves the problem that Boroditsky described in the quote above (this fact is worth noticing as Boroditsky herself is one of the major figures behind the spatio-cultural approach to temporal reasoning).

The original question that we were tackling in this section was how mental representations of time come to have their temporal content. The story that we are considering here would explain this by saying that temporal representations involve distinct uses of the underlying machinery used in spatial representation. Presumably, since these approaches take space to be *more concrete*, it should be the case that the underlying representational system has the ability to represent space. That much should be granted. However, the question then becomes how do we explain how a system used for representing space comes to represent time. The question of how the temporal representations we attribute in the explanation of some capacity come to have their temporal content simply remains unanswered. Furthermore, if we require a sophisticated explanation for how *these temporal representations* acquire their content, then what should we say of the vast number of other temporal representations that cannot be explained via these sophisticated means? Surely prior to the acquisition of language and prior to a

²⁰ Additional evidence for the view comes from the existence of SNARC effects. See chapter 5 for a discussion of this evidence.

significant amount of cultural influence, children are capable of navigating their environment, and this likely involves some temporal representation. Therefore, it's likely that not all temporal representations should be given such a complicated explanation.

Another of these sophisticated explanations for how humans come to acquire their temporal reasoning capacities bases temporal reasoning on an understanding of *causation* (Hoerl, 2009; Hoerl & McCormack, 2011, 2016). The theoretical ancestor for this view comes from a failed metaphysical project of Hans Reichenbach's (1956, 1957). Reichenbach attempted to give a metaphysical analysis of time in terms of causation. Time, or the temporal structure of our world, he argued is not fundamental but is based in the causal order of the world. It's widely accepted in the literature that this analysis fails (see Hoerl and McCormack (2011) for a discussion), however, several researchers have taken this analysis as a plausible one for understanding the development of temporal reasoning capacities in children. There is some evidence to show that infants have difficulties in reasoning about time, especially for novel temporal sequences as opposed to highly familiar scripted encounters. However, this ability to reason about time seems to be come online as more and more sophisticated capacities for *causal reasoning* emerge. The proposal is then that the concepts of EARLIER THAN and LATER THAN emerge from the concepts of CAUSE and EFFECT.

Whether or not this causal explanation of the origins of temporal concepts succeeds, the same explanatory burden remains for this theory as it did for the spatio-cultural approach. That is, the question was one of how the temporal representations we attribute in the explanation of some capacity come to have their temporal content. Saying that these temporal concepts are born out of representations of causal relations does not tell us that story. We still lack the explanation that we require.²¹

²¹ This isn't to say that these sophisticated accounts *couldn't* provide the needed explanation. Rather, it's just to point out that the important question of how we come to acquire representations with a particular type of content remains unanswered.

The point to take from all of this is that even granting that time itself is a somewhat peculiar aspect of our environment, the difficulty of explaining how we come to represent time will not be exhaustively solved by proposing complex and sophisticated explanations of those temporal capacities. Furthermore, we simply have no a priori reasons for assuming that there will be a single explanation for how we represent time as opposed to there being a multiplicity of ways in which animals across the animal kingdom come to represent time. And, as I will argue throughout the thesis, by looking at the specific neural machinery involved in the production of time sensitive behavior we will get a better sense as to what theories of content are viable for our explanatory purposes, and in doing so we will see how the *temporal coordination problem* is overcome.

1.5. Moving forward

The interest in the mental representation of time has been around for an incredibly long time and throughout this history the relation between the temporal structure of the mind and the temporal content of mental states has always been up for debate. As Locke put it,

It is evidence to anyone who will but observe what passes in his own mind, that there is a train of ideas which constantly succeed one another in his understanding, as long as he is awake.

Reflection on these appearances of several ideas one after another in our minds, is that which furnishes us with the idea of succession...

(Locke, 1689 Chapter XIV, 3)

In order to explain how we come to have *ideas*, or in modern parlance *concepts*, of time we simply needed to appeal to the actual temporal structure of our experiences themselves. It was the very temporal sequence of non-temporal mental states that would furnish us with mental states that were *about time*.

This story of Locke's quickly met resistance. Both Reid (1855) and Kant (1781/1998) argued that a mere succession of experiences could not amount to an experience of succession. That is, the mere temporal structure of experience could not account for the content of experience. Instead, it was argued that something further is needed to explain how mental states come to be directed at or about time.

Since the discussion of temporal representation in the early modern period quite a lot has occurred. In particular, perhaps the most widespread and influential discussions of the mental representation of time in philosophy came in the form of William James' (1890) and Edmund Husserl's (1917/2008) discussions of temporal experience. Both authors proposed accounts of how we could come to experience the temporally dynamic world around us, and both authors have been interpreted in a number of distinct ways.²²

Unfortunately, much of the philosophical discussion of the mental representation of time remained focused on the narrowly defined *experience* of time. That is, philosophers have focused on how we come to consciously experience the temporal structure of the world around us and the *flow* of time's passage. Furthermore, while there has been an impressive explosion of empirical work on temporal representation coming from computer science, cognitive psychology, neuroscience, zoology, linguistics, and anthropology, there has been only limited engagement between that empirical literature and philosophical discussions of temporal representation. Perhaps the reason for why this additional relevant literature hasn't been appealed to by philosophers is due to the fact that philosophers have concerned themselves primarily with the conscious experience of time, and this other work doesn't directly bear on the conscious experience of time.

²² For interesting historical work on these authors see (Andersen & Grush, 2009) and (Andersen, 2014). As we go through the thesis, one might notice that I do not speak about the views of either James or Husserl in any detail. There is good reason for this. Both historical figures produced views that have been interpreted in a great variety of different ways. As a result, in order to fruitfully discuss either view, either as a topic in and of themselves or in the service of some further aim, would require a significant amount of historical analysis. And that is something that would simply take us too far from the main themes of this thesis.

In what's to come we'll be looking at temporal representation in a way that goes beyond the mere conscious experience of time. By beginning with the *temporal coordination problem* more generally we are able to consider the various ways that animals come to keep track of and coordinate their behaviors with the temporal structure of the world around them. Once we consider temporal experience as being just one aspect of this more general temporal coordination problem the wealth of empirical work from various fields instantly becomes relevant to our research. In discussing this point, we'll see that positing internal representations of time is often required, especially when it comes to discussing how we experience time, and I will make strides towards explaining how it is that those internal representations of time come to have their temporal contents. Ultimately, we'll not only understand something about how the mind comes to be directed at or about time, but also, we will make some headway in seeing that the mind can be given a naturalistic understanding through seeing how the mental representation of time can itself be naturalized.

Chapter 2: Empirical Timing Models and the Explicit / Implicit

Distinction

As described in the introduction, all animals face a form of the temporal coordination problem in that all animals need to coordinate their actions with the temporal structure of the world around them. Rather than explain how any aspect of the temporal coordination problem is solved, the goal of this chapter is to lay out some of the general ways in which the problem *could be solved* by giving a taxonomy of the explanatory approaches found in the empirical and philosophical literature.

As we'll see, the major axis according to which the taxonomy in this chapter will be structured concerns the degree to which these explanations appeal to explicit representations of time. At one extreme, the temporal coordination problem is overcome through the use of dedicated mechanisms for the explicit representation of time. Towards the other extreme are positions that simply do away with any representational states whatsoever. Between these extremes are a number of different positions that preserve some role for representational states and even some role for representations of time in explaining temporal coordination.

In order to get a handle on these different explanatory approaches it will be useful to distinguish between two ways that a system might be said to represent some content, that is, to distinguish between *explicit* and *implicit* representation. The chapter will begin by tackling this task. In section 1, I will go over a number of different means by which people have attempted to draw the explicit / implicit representation distinction and argue that many of them, while tracking interesting psychological divides, do not track genuine differences in *how* some content is represented. In section 2, I do the work of articulating a sense of the explicit / implicit distinction that does in fact draw a line between two ways of representing some content. With that distinction in hand, I'll turn to the main task of this chapter and provide a taxonomy of the various explanatory strategies for understanding how animals overcome the

temporal coordination problem. Since, as we'll see throughout the thesis, the temporal coordination problem that any animal faces actually breaks down into a number of distinct coordination problems, there is unlikely to be any single explanatory strategy that will explain the entirety of the coordination problem. Therefore, instead of arguing that any of these strategies is correct, the goal here is to note the variation in explanatory approaches and to simply lay them out so that we may appeal to them later on when needed.

2.1. The explicit / implicit distinction in cognitive science

In this section, we'll look at how some in the cognitive sciences have attempted to draw the implicit / explicit representation distinction and argue that none of these attempts track genuine differences in how representations encode their content. In the next section, we'll turn to providing an account of the implicit / explicit distinction that does track a difference in the manner of encoding some content.

Before we lay out an account of the explicit / implicit distinction that tracks a genuine difference in ways of representing some content, it will be helpful to survey some of the ways that people have attempted to draw the explicit / implicit distinction in the cognitive sciences. In each case, the distinction drawn, while tracking some distinction amongst representational states, fails to distinguish between two ways of encoding some information in a representation. Instead, these existing approaches mark the distinction not in how the representation encodes its content, but how the representations fit within a larger functional system.

How then has the explicit / implicit distinction been made? Perhaps the most common way of marking the distinction is in terms of the conscious availability of some representational content (Schacter et al, 1993). According to this way of marking the boundary, only information that is available to

conscious report is explicitly represented whereas implicitly represented information is not so available. What this distinction rests on, however, is not a distinction in how representations encode their content, but rather in *how content encoded by a representation is used*. Just a quick survey of different theories of consciousness shows us that many theories of consciousness distinguish between conscious and unconscious mental states (or processes) not in terms of how they encode information, but rather, in terms of the functional role played by particular representations.²³ For instance, on the popular Global Workspace Model of Consciousness (Baars 2005; Dehaene et al 1998) what distinguishes conscious from unconscious representations is whether the content encoded in the representation is being made available to a wide range of consumer systems as a result of entering into the global workspace. Notice that on this sort of model the very same representation can be conscious at one moment and unconscious at another, and as a result, the same representation will be at times explicit and at other times implicit.²⁴

To further illustrate how this way of carving the distinction fails to track ways of encoding content in a representation, notice that we can render a representation unconscious to the point of not even being possibly available to consumer systems by manipulating aspects of the cognitive architecture that lie beyond the confines of the individual representation. Cognitive science is full of cases like this. One particular case comes from the use of *transcranial magnetic stimulation* (TMS) (Lamme 2001; 2006). In these cases, stimuli that would normally be consciously perceived are rendered unconscious by the application of a precisely timed TMS pulse that disrupts the feedback mechanisms that would place the representation into a dynamic feedback loop with higher-level cortical areas. Again, the particular representation in question that may have become conscious is left untouched. What is changed is how

²³ Some theories, such as Tye's PANIC theory (Tye 1995; 2000), also require that conscious representations possess certain types of content. In Tye's case, conscious representations must possess *abstract non-conceptual intentional content*. However, in addition to the type of content conscious states possess they must also be *poised* to impact downstream systems, and in this way his theory relies on the functional difference I am describing here.

²⁴ Besides the global workspace, we find functional characterizations of the conscious / unconscious distinction in Lamme (2001; 2006), Prinz (2000; 2012), Dennett (1992), Dretske (1995), Rosenthal (2002), Koch (2004).

other parts of the brain are functioning. Since the representation is left the same, then the difference in consciousness cannot mark a difference in how the (untouched) representation encodes its content.

An alternative, but somewhat related, attempt to mark the distinction between implicit and explicit representations is given in terms of the amount of processing²⁵ needed to utilize information (Kirsh, 2006). According to this way of marking the distinction between implicit and explicit representations the difference is a graded one. Explicit representations are those whose contents are readily made use of with minimal effort and with minimal demands on limited processing resources. Implicit representations are, on the other hand, ones that require significant processing resources in order to access their content. In other words, explicitly represented content is capable of driving other behaviors with minimal additional processing while in order for implicitly represented content to drive further activity requires a *significant* amount of intervening processing. While there is something to this way of marking the distinction, there is nevertheless a significant shortcoming to using the amount of processing needed to extract the content as a means of dividing explicit from implicit representations.

Take the sentence ‘the police officer wore a hat’. Armed with a decent understanding of English the content of this sentence is readily extractable. It tells us something about what a particular person had on their head. In fact, the content of this sentence, on standard linguistic theories, is simply a result of a compositional process whereby the words contribute meaning to the sentence in a manner determined by the grammatical structure of the sentence. Now, take the sentence ‘police police police police police’ (from (Kirsh 2006)). On a first glance this sentence seems to be just a meaningless stringing together of the same word five times over. However, this sentence is in fact a well-formed sentence of English, whose content is determined through a straightforward compositional process, but as anyone coming across this sentence would notice, the content of the sentence requires a significant amount of mental

²⁵ This notion of *amount of processing* is left rather vague in these accounts. Typically, however, this amount of processing can be equated with the difficulty for the subject to extract the relevant information.

effort to extract. A grammatically correct interpretation of the sentence is *policemen who are policed by policemen also police policemen*.

The way that these two sentences carry their content is the same. They both possess their content as a result of the compositional semantics that dictates how the word level meanings are brought together to form the sentence level meaning. Yet they differ in how much effort it takes to utilize their content. If, as we're assuming, the explicit / implicit distinction is one about the way in which content is encoded by a representation, then the amount of information processing resources required for the extraction of that content *cannot* be what defines the distinction, since at most, the amount of information processing resources required to extract some content tells us about how some representation enters into the rest of the cognitive economy.

One final approach from the literature for marking the implicit / explicit representation distinction is given in terms of whether some content is *hardwired* into the operations of a system, i.e. implicitly represented, or whether it is explicitly represented through some variable aspect of the system.²⁶ As a matter of biological fact we have been endowed with a number of different neurological mechanisms that allow us to interact with the world in certain ways. Some of these mechanisms, it has been argued, have been shaped in just the right way that they seem to provide us with some rather sophisticated capacities. For instance, in the literature on object perception²⁷ a class of theories (e.g. Carey 2009; Leslie et al 1998) appeal to the operation of the *object tracking system* in order to explain how infants and adults come to form perceptual representations of enduring physical objects. At the core of these theories is a perceptual system that is designed by evolution to track stimuli that adhere to certain principles, such as continuity of form, spatio-temporal continuity of trajectory, etc. Only stimuli that adhere to these principles will be tracked by the object tracking system. The result is that even very young infants will appear to have an

²⁶ See (Bernal, 2005) for a discussion of this sort of approach to the explicit / implicit distinction with regards to the purported object-specific principles employed in infant object cognition.

²⁷ Similar conclusions could be drawn by discussing the literature on biological motion perception.

understanding of objects as being spatiotemporally coherent enduring entities. Yet, on many of the interpretations of this system, the object tracking system does not explicitly represent any of these principles, rather it only ever produces representations that pick out or refer to the objects they track (Bernal 2005; Pylyshyn, 1989, 2001, 2007). The system is hardwired to appear as though it is adhering to some rules, and in this way, some have described the object tracking system, and its operations, as providing the infant with *implicit knowledge* of objecthood. Nowhere is this knowledge explicitly represented, but it is embodied in the operations of the system.

In this way, we have something that seems to be a genuine difference in how a system encodes certain information. Implicitly represented contents are those that are possessed by a system due to the way that it is hardwired to operate. While explicitly represented contents are those that are carried by malleable or variable properties of the system. The purpose of this section isn't to dictate how someone should use the terms 'explicit representation' or 'implicit representation' as each of the various proposals covered so far seem to latch onto psychological distinctions that might have some explanatory use. Nevertheless, the worry with this particular means of drawing the implicit / explicit distinction is that it seems to trivialize the notion of representation. The only compulsion to claim that the principles of objecthood are represented at all is that the behavior of infants and adults seems to be best characterized as though it were following certain rules. However, this way of characterizing things applies to every lawful phenomenon in the world. A plant, as it grows towards the light, could be interpreted as having the implicit knowledge that light is good. Yet it seems as though this sort of interpretation isn't needed to understand how the system is actually operating. We can just describe the plant as having certain biological mechanisms that respond to light in very mechanical ways. Similarly, in the case of object perception, we needn't ascribe any representational component to the operation of the object tracking system, since we only need to say that it responds in a very mechanical way to retinal stimulation. Instead, as we'll see in the next section, the notion of representation is typically tied closely to the *states of a system*, how these states track the world around it, and how ascribing semantic content to these states

helps us understand how the representational system is put to use by consumer systems of the representation.

2.2. Minimal semantic interpretations and the implicit / explicit distinction

The approaches to marking the implicit / explicit distinction either trivialize the notion of mental representation (as in the case of taking hardwired operations as implicitly representing some principles) or simply fail to mark a distinction between forms of representation. In this section, I'll articulate a way of understanding the distinction that is not only found in the empirical literature but that also tracks a genuine distinction between ways that a representation can be said to carry some content.

Let's begin by drawing an analogy with linguistic communication where we can draw a distinction between *the literal linguistic meaning of an utterance* and *what is implied by the utterance*.²⁸ Take the following sentence:

1. Sally and Johnny are the only students in the course.

The literal meaning of the sentence is given by what we can call *the minimal semantic interpretation* (MSI) of the language. Given the expressive demands of a symbol system, the MSI of the system is the simplest assignment of semantic contents to the symbols of the system that given rules for the combination of these symbols provides for the full expressive power of the entire symbol system. The meaning of (1) is given by the assignment of semantic contents to the individual words that compose the sentence and the rules by which these symbols are combined. There needn't be any additional semantic interpretation of the sentence as a whole, apart from the interpretation of the individual symbols and a

²⁸ The distinction here is very similar to the Gricean distinction between what is said and what is communicated (H. Grice, 1975), however, we should resist the urge to identity the two distinction. According to Grice, *what is communicated* by a sentence are the implications of the sentence that rely on the acceptance of some auxiliary assumptions. If one were to reject those auxiliary assumptions, then the implications would not follow. However, the implications that I will be discussing here needn't rely on any auxiliary assumptions for their acceptance.

specification of how those symbols combine (that is, (1) is non-idiomatic). Take the literal meaning of (1), *that Sally and Johnny are students in the course and no other students are in the course*, as the *explicit content* of the sentence.

Notice, that the sentence as a whole only has a meaning, or evaluable semantic content, *in a context* given that the phrase 'the course' is only given an interpretation *relative to a context* (whether that is the context of utterance or interpretation doesn't matter). Due to the context dependency of how we interpret many symbols, the notion of a MSI needs to be modified to allow for this context dependency. The MSI of a system is, given a context, the simplest assignment of semantic content to the basic constituents of a symbol system that in conjunction with rules of combination provide for the full expressive power of that system.

Given the MSI of the language, in this case English, we can then describe two senses in which (1) has certain implications that a competent speaker of English would be warranted in inferring given (1). Some of these implications, such as (2) and (3), are guaranteed in a context free manner:

2. There are two students in the course.
3. If x is a student in the course and x is not Johnny, then x is Sally.

Both (2) and (3) can be directly derived from (1) regardless of the context in which (1) is used.²⁹ All that is needed to make these sorts of inferences is that the inference maker possess the appropriate concepts required to articulate the implications and that the inference maker grasp the principles that ground these inferences. If, in addition to the logical inferences that lead to (2) and (3) we allow for some analytic inferences, then the class of context free implications of (1) can be greatly expanded to include things, such as:

²⁹ In fact, all logical truths can be derived, and are therefore, implications of any sentence of English.

4. There are at least two individuals capable of learning in existence.³⁰

If we accept that 'x is a student' analytically implies that 'x is capable of learning and exists', then we can infer (4) from (1) independently of any *factual* information about the world that extends beyond semantic or conceptual relations between concepts. Therefore, if we accept that analyticities exist, then inferences like (4) from (1) will be guaranteed by the MSI of English and (1). Of course, however, the existence of analyticities is a highly contentious matter in analytic philosophy and their existence is doubted by many (Fodor & Lepore 1991, 1992; Quine 1951). So, perhaps inferences like (4) do in fact require factual knowledge of the world. Regardless, the discussion here can continue regardless of whether analyticities exist, as the important distinction is between the literal content of a representation and that which can be inferred from the representation.

This leads us to the next category of implications from (1). These implications are secured by bringing to bear other knowledge of the situation represented by the original sentence. In other words, these inferences are only derived in a context dependent manner that relies on information not specified by the MSI of the language and the logical structure of the language. Imagine that the context in which (1) appears is such that courses at that institution rarely ever have less than 10 students unless the course meets very early Saturday mornings. Given this additional information, we can infer from (1) the following:

5. Sally and Johnny are likely morning people.
6. The course most likely meets on Saturday mornings.
7. Other people are not interested in attending this course.

³⁰ If the inference from x is a student in a course to x is capable of learning and exists isn't one that strikes you as being an analytic statement, then substitute the sentence in (4) for whatever strikes you as more plausibly analytic. If nothing strikes you as a plausible analytic statement, then the discussion can continue on the assumption that there are no, or only very limited, analytic statements.

The MSI of the English does not by itself give us reasons for inferring (5) – (7). What unites (2) – (4) and (5) – (7) into one coherent group is that in every case the inference requires that we first interpret the literal meaning of (1) to derive any of (2) – (7). More specifically, we can say (1) *implicitly represents* X, just in case X is a context-free implication of (1). We can alternatively say (1) *implicitly represent X given C*, just in case X is a context dependent implication of (1). This distinction is important, since the context dependent implications of (1) are only warranted given some further facts that are not provided by the MSI.

Given the notion of implicit / explicit representation given here, we need to distinguish this from an outwardly similar appearing phenomena. Imagine a situation in which you and I need to coordinate our behavior, but we wish to do so in a way that goes unnoticed by onlookers. We might establish a *code word convention* where we use (1) to indicate that one should turn off their phone. When you use (1), I then know that you are telling me to turn off my phone. One might be tempted to say that in this case (1) implicitly represents that I should turn off my phone. However, this would be mistaken. The use of (1) as a code word fails to depend on the MSI of (1) necessary to understand the language in general. Instead, by adopting the convention that (1) be used to tell someone to turn off their phone, we no longer are using the original MSI. Instead, we have provided a new MSI.³¹ The actual phonetic string associated with (1) is now subject to two distinct interpretations. In this case, then, a usage of (1) will have two explicit contents, each associated with a distinct MSI.

With this usage of the implicit / explicit distinction in hand, we can understand how it can be applied to mental representations. While linguistic expressions consist in discrete units that can in some sense be interpreted however the language users like, as indicated by the use of (1) to mean that someone should turn off their phone, mental representations do not possess either characteristic. In providing an

³¹ In effect, the use of (1) in this *coded context* could be replaced with something like a wink of an eye or a random noise as long as we have previously agreed to use whatever “expression” in order to convey that one should turn off their phone.

analysis of the representational content of some mental mechanism we don't have clearly defined symbols that can then be assigned some semantic value. Instead, we must avail ourselves of *states of physical systems*. These states are the units that we can then assign semantic values to. However, in assigning semantic values to the system, we are constrained by our scientific practices to make the functioning of the system intelligible and by the causal efficacy of the states of the system in making the representation's contents available to downstream processes.

On the causal efficacy of the relevant states, it is crucial for a system to be a genuine representation that its contents be available to some consumer system. In the linguistic case, the actual physical marks on the paper, or the waves in the air, must be such that they are capable of causally impacting other systems that allow for the extraction of the encoded content. Some piece of paper may be my sister's favorite piece of paper, and we could plausibly try and ascribe some semantic content to the paper in virtue of having this property. But that would be a mistake. The property of *being my sister's favorite piece of paper* simply isn't something that changes the paper's causal powers thereby making some content extractable by consumer systems. In the case of mental representations realized by neural systems, whatever property of the system that we appeal to as being the semantically significant properties (i.e. the properties that encoded or bear semantic contents) must be capable of causally influencing downstream neural processes. So, while a neural system might change its mass through various metabolic processes, we could not take these states of the system as being semantically significant since downstream neural processes are not sensitive to these states of the system. That is, the mass of a neural system simply isn't the sort of state that is causally efficacious within the system (see chapter 3 for a more detailed discussion of this point).

Turning to the intelligibility constraint, let's take an extreme example. If we want to provide a semantic characterization of the functioning of the vestibular system in humans, we would be misguided to interpret the state of the vestibular system as having cities in France as their semantic values. While an

interpretation of this sort may be made formally consistent, in that there is a one-to-one mapping between cities in France and states of the vestibular system, it would not make human behavior intelligible since cities in France seemingly have no significance for the balancing behavior of humans. Instead, if the vestibular system is to be given a semantic characterization in order to make some feature of human behavior intelligible, then we need to assign semantic values to that system that have this intelligible making role.

This notion of intelligibility can be understood in terms of the difference that accurate or inaccurate representation has in the production of behavior. Cities in France simply aren't significant for a person's ability to balance themselves, and it goes without saying that accurately representing the state of cities in France is unlikely to be of any significance for human balancing. In other words, if we assign the states of the vestibular system a semantic interpretation that maps the states of the system onto cities in France, the fact that the system is accurately or inaccurately representing some state of affairs will have no explanatory value in our understanding of the cause of successful balancing. However, the direction of the gravitational pull is significant for human balancing in that if the system is not accurately keeping tabs of the current direction of acceleration, then the system will not be able to properly orient itself in its environment. So if we have two consistent interpretation of the representational role of the vestibular system, according to one where the system represents cities in France and the other that represents gravitational pull, then it seems clear that we should adopt the semantic interpretation that assigns direction of gravitational pull to the states of the vestibular system. None of this should strike anyone as controversial.

Taking the example of the human vestibular system, we have a system with various states – e.g. influence on cilia by fluid in the semi-circular canals. Let us also assume that this system can be provided

a MSI³² in that each state of the semi-circular canals represents the acceleration of the head in a particular direction. Furthermore, this assignment of semantic values to the states of the semi-circular canals makes the balancing behaviors of humans intelligible. We can then say that the states of the semi-circular canal explicitly represent accelerations of the head. Let us take a particular state of the semi-circular canals in the following way:

8. Head accelerates at $X \text{ m/s}^2$ in direction Y.

From (8) it can be inferred that:

9. Head accelerates at speed $< X+1$ in direction Y.

(9) can then be said to be *implicitly represented by N*, where N is a state of the semi-circular canal. Given some further contextual factors, say, that if one's head accelerates at a speed greater than $X-1$, then some damage is occurring to the organism, then it can be inferred that:

10. The organism is incurring damage.

Both (9) and (10) are implicitly represented by (8), however, (10) is only implicitly represented given some further contextual facts. In what follows, the distinction between contextual and non-contextual implicit representation will be of some importance, but for now just treat them as both forms of implicit representation.

This way of marking the implicit / explicit representation distinction explicitly appeals to a distinction in the way in which information can be encoded by a system. Take one further example, the implicit representation of numerosity by elements in working memory (e.g. Carey (2009)). It has been argued that the rudimentary numerical abilities of infants can be accounted for by the *implicit*

³² While the original statement of an MSI appeals to a combinatorial syntax for the system being interpreted, a non-combinatorial system like the vestibular system can still be given a MSI since the class of rules for the combinatorial syntax would simply be empty.

representation of numerosity by working memory representations. Infants, it is argued, are able to open *mental files*, or explicitly working memory representations, that demonstratively refer to a limited number of elements (between 2 – 5 elements). For instance, provided with a display of four puppies the infant might form a working memory representation of the following form:

11. *That₁, That₂, That₃, That₄.*

Where each 'That_x' explicitly represents a particular element. However, the entire representation, consisting of the four demonstrative representations, can be said to *implicitly represent that there are four elements* because an implication of (11) is:

12. There are four elements.

Number needn't be explicitly represented, in that number never needs to appear as the semantic content that is attributed to a mental representation, rather the implicit content is something that could be extracted or derived given the MSI of the initial representation, given that (11) along with some auxiliary assumptions would entail (12). Only later may the numerical content be explicitly represented. However, for the explanation of behavior, the MSI of the working memory representations of mental files needn't attribute to any single representation numerical semantic content.

Notice that for both explicit and implicit representations the information is only made available for further processing provided that the further processes are causally sensitive to the states of the system that encode the information (see Kulvicki (2005) for a discussion of the notion of extractability). Since the implicitly encoded information is carried in virtue of the explicitly encoded information the availability of both sorts of information depends on the availability of the explicitly represented information. Furthermore, in order for any downstream process to act on the encoded information it must be the case that the downstream processes not only are causally sensitive to the encoding properties but

that there is a mechanism that treats those encoding properties of the representation in a way that respects their semantic properties.

This point is one that was emphasized by Fodor in the defence of *methodological solipsism* (1980). A representational system in psychological explanation is first and foremost a mechanistic system, the functioning of which is governed by the causal relations amongst the parts of the system (Bechtel 2005, 2011; Craver 2007; Machamer et al 2000). While the interpretation of a representational system might be governed by non-local factors, as in the case of externalist theories, the operation of that system is local and mechanistic. As a result, whatever local properties are being mapped onto the semantic contents of the system must be such that the operations of the system can be explained by appealing to those local properties, as those local properties are the properties that enter into the mechanistic explanations that explain the transitions of the states of the particular system. In this way, the difference between implicit and explicit representation cannot be cashed out in terms of, nor does it even fall along the same lines as, needing further downstream processes for the extraction of information. In order for either explicitly represented information or implicitly represented information to be extracted the representation must be related to downstream systems that are appropriately causally sensitive to the encoding properties of the representation and whose causal operations respect the semantic properties of the encoded information.

2.3. Explicit and dedicated timing mechanisms – internal clocks

With the distinction between implicit and explicit representation in hand, we can now turn our attention towards providing a taxonomy of the various general approaches to solving the temporal coordination problem.

The two sorts of timing mechanisms described here fall within the explicit timing category; these mechanisms involve states that constitute explicit representations of time (or temporal features). The important difference between these two sorts of models is in *how* they represent time and *what* temporal features they explicitly represent. It is a difference in the mechanisms involved, or in other terms, it is a difference in the structure of the representational vehicle, that distinguishes these timing models, and, as we'll see, this difference in the structure of the vehicles results in a semantic difference in what these models are capable of representing.

Period timers

Consider the way in which a standard 12-hour analog clock keeps track of time. The underlying mechanism behind the functioning of the clock is an oscillatory mechanism, where the oscillator is a mechanism that exhibits periodic behavior with a fixed period. When one cycle ends, the underlying mechanics of the oscillator guarantee that the system will enter into a subsequent cycle with the same period as the previous cycle.³³ Given our time keeping practices, we take the hands of the clock, as they move around the clock, to be explicit representations of the time of day. Our time keeping practices assign times as the semantic values of the states of the clock – this assignment constitutes the MSI of clocks.

Keep in mind that 12-hour analog clocks are human artifacts. They are technological achievements brought about by the conventional practices of modern cultures. The appeal of period timing mechanisms when employed in cognitive science is that the very same sort of mechanism, an

³³ The idea of the “same period” has to be taken somewhat loosely. No clock is perfect, so unless we idealize the functioning of the system, we should talk about *periods that average a certain period*. For the most part, nothing hangs on this qualification, but there will be places where a stochastic or probabilistic understanding of oscillator periods will become useful – i.e. when explaining variance in timing behaviors by appealing to stochastic processes as a source of noise. See (Wearden, 2001) for a discussion of how variation in the periods of oscillatory mechanisms can be of use in explaining variations in timing behaviors.

oscillatory mechanism, can be plausibly appealed to in order to explain psychological phenomena. However, in order to understand the analogy between the cultural artifacts that are clocks and psychological period timers, not only is it important to notice the underlying mechanics of these clocks but we also need to know what the semantic contents of clocks are. At a rough pass, we know that we use clocks to tell us what time of day it is at the moment of reading the clock, but this leaves two possible interpretations of what the semantic contents of clock states are. Let us take a standard 12-hour clock that reads 5:30 as an example. The first way to characterize the content of this system is as follows.

13. It is currently 5:30.

14. It is 5.5 hours since the cycle began.

The first reading of the clock has the clock picking out a specific moment in time and saying of that moment that it has the property of being 5:30. The second reading of the clock, however, does not only mention a single moment in time, the present moment, but it states a relation – a temporal interval – that holds between two moments in time. For the time being we can remain neutral as to which of these readings is how we ought to interpret period timers, or clocks more generally.³⁴ We can even remain neutral as to whether there is a single semantic interpretation available for period timers or whether there may be cases in which period timers are used to express a semantic content in line with (1) while other cases may involve a semantic content in line with (2). Finally, we can even remain neutral for now as to what these temporal properties are that are picked out by clocks. For our purposes in this chapter we are merely laying out the mechanisms that are at the heart of these different models.

Despite their non-conventional nature, oscillator mechanisms realized in neural systems have many of the same features as clocks. Let's consider a concrete case to help our discussion. Circadian

³⁴ In fact, for the time being we can be neutral as to whether or not these two options are genuinely distinct options. One might cash out the predicate “is currently 5:30” in terms of *its being 5.5 hours after the beginning of the cycle*. If that is how we understand the property, then (1) and (2) are merely notational variants of one another. However, if we think we can give an analysis of the predicate in (1) that respects the surface appearance of a non-relational property, then (1) and (2) will turn out to have distinct metaphysical commitments.

behaviors in mammals are thought to be realized by a neural system that centers on the *suprachiasmatic nucleus* (SCN).³⁵ Individual neurons in the SCN exhibit oscillatory behavior in their spontaneous firing rates that follow a roughly sinusoidal pattern with an approximately 24-hour period. If we take the firing rate of SCN neurons to be the states of the system that carry information concerning the time of day, then we run into a problem in that considering the firing rates of individual SCN cells introduces a representational ambiguity in the system. For any given firing rate of an SCN neuron, there will be two times of day associated with that firing rate. As a result, the system could not by itself indicate one specific time of day. The situation is similar to seeing a 12-hour analog clock in isolation. The hands of the clock will point to particular places on the clock face, but without any further disambiguating information, you will not be in a position to know whether the clock indicated an AM or PM time and will be unable to provide a precise time as the semantic content (as opposed to a disjunctive content that specifies either AM or PM). There is further information that is needed to properly assign semantic values to the state of clock beyond what is given by merely looking at the clock face.

While this situation isn't problematic for clocks, since we often have disambiguating information that allows us to appropriately interpret the clock, the same can't be necessarily said for neural mechanisms.³⁶ Perhaps we needn't appeal to any time keeping devices that have less ambiguity than clocks, however, we have the theoretical machinery to grasp how networks of oscillator mechanism can override this ambiguity.

The most influential of these disambiguation techniques is seen in the Church-Broadbent Model (Church & Broadbent, 1990) that instead of employing a single oscillator employs a set of oscillators. Their original model (1990) involved 11 distinct oscillators with different oscillatory periods – the least

³⁵ This system is discussed at length in chapter 3.

³⁶ Perhaps some disambiguating information can be acquired through other sensory means, but for reasons discussed in chapter 3 this way of disambiguating oscillator information will not work for circadian oscillators. The circadian system is realized by an oscillator system that unambiguously indicated times of day.

being 0.2 sec and progressively doubling up to 204.8 sec. Temporal information is not encoded by any single oscillator, instead the unit of semantic significance, the unit that is assigned a semantic value, is the state of the entire oscillator array represented as a vector of oscillator phases where each oscillator has only two possible states, + or -. ³⁷

Using this set up, we can represent the state of the oscillator network that indicates that 15 seconds has passed as “+1 +1 -1 +1 -1 -1 +1 -1 -1 -1” and 40 seconds as “-1 -1 -1 +1 -1 -1 +1 +1 -1 -1 -1”. Given individual oscillators with a limited number of states, a system of these oscillators can successfully, and unambiguously, represent a number of temporal durations. ³⁸

Someone could object here, however, in that the state of the oscillator will still ambiguously pick out (at least) two distinct times. The objection goes that no matter how many distinct oscillators we employ if we consider the operation of the oscillators over a sufficiently long period of time, then the states of the oscillator will ambiguously pick out at least two times. Suppose, that the oscillators cycle through a complete series of states every 24 hours. If we interpret the states of the oscillator over a 24-hour period, then the oscillator states will unambiguously pick out a specific time within that 24-hour window. However, if we were to instead interpret the states of the oscillator over a 48-hour period, then the states of the oscillator will ambiguously pick out two distinct times (say, 10pm today or 10p tomorrow). In order to avoid this ambiguity, we would need some means of restricting the time interval over which the cycles of the oscillator are being interpreted.

Perhaps we could overcome this sort of ambiguity by appealing to external means of demarcating the intervals according to which the oscillator states are being interpreted (e.g. day / night cycles).

However, an entirely different approach would be to simply embrace the ambiguity. All that would be

³⁷ One could provide a semantic interpretation of each individual oscillator, however, if each oscillator were interpreted separately, then there would be a much greater level of ambiguity in the content of the system.

³⁸ See (Wearden, 2001) for an overview of this literature. The values and numbers of oscillators in the Church-Broadbent model were arbitrary, but also see (Wearden & Doherty, 1995) for how the choice of number of oscillators and their possible phases are important for the representational abilities of the system.

required is that the ambiguity be one that does not cause problems in the *use* of the system. If for instance the oscillator were used to tell the time of day, and the sort of ambiguity present in the system was one that disjunctively picked out a specific time of day on distinct days, then no behavioral consequences would follow from the particular sort of ambiguity. A single oscillator that ambiguously (or disjunctively) picks out a time every 12 hours would likely have problematic consequences for behavior if the oscillator was supposed to coordinate the activity of the organism with the specific time of day. However, if the ambiguity resulting in a disjunctive content that picked out times separated by 24 hours, then this would be unlikely to lead to any behavioral consequences, as the same time of day would be picked out.

Given these techniques for getting around the ambiguity introduced by appealing to a single oscillatory mechanism, we can assume that if period timers are involved, and there is no system independent means to disambiguate the information they carry (i.e. other sensory information such as ambient light levels), then we can assume that they are employing some technique, like the one found in the Church-Broadbent model, to avoid ambiguity within a particular temporal interval.

One standout limitation of period timers is that on either of the two possible semantic interpretations, period timers pick out specific times of day by either picking out that time directly (as in option 1) or by picking out that time by noting the temporal interval between some initial moment and the current time. These are the two possible ways of understanding what is explicitly represented by these timers. However, what is often important for animal behaviour is the ability to track the temporal interval between two arbitrarily chosen events.³⁹ If the animal is capable of remembering when two events occur, via noting the state of the oscillator system when those events occur, i.e. noting that *a occurred at T_1* and *b occurred at T_2* , where 'a' and 'b' are events, and ' T_1 ' and ' T_2 ' are times of day picked out by states of the oscillator, then the combination of these representations will implicitly represent the temporal interval

³⁹ Montemayor (2013) describes these limitations to period timers, although, his notion of explicit representation is the functional notion of needing further processing for the extraction of information and not the semantic notion employed here.

between a and b . It is only in conjunction with some memory store that we even get implicit representations of intervals because a memory store is required for the two time of day representations that bound the specific interval, and then this information is only extractable given the appropriate cognitive machinery to extract the interval information.⁴⁰

Nevertheless, if the period timer is combined with a memory system that records states of the oscillator system and the events that occurred simultaneously with those oscillator states, and the organism is equipped with the rudimentary mathematical ability to calculate interval lengths, then we have a powerful system that allows us to register both specific times of day and interval lengths.

Interval timers

The guiding analogy for understanding period timers was the 12-hour analog clock. The guiding analogy for interval timers is the hourglass. Hourglasses measure intervals between two events via the accumulation of some medium that accrues at a fixed rate and that indicates how much time has passed since the hourglass was first “activated”. If you want to measure how long it takes you to complete a puzzle, you can turn over the hourglass when you begin the puzzle, and then by noting how much sand, let's say, has accumulated in the lower bulb of the hourglass you can determine how much time has passed.

We can notice two striking differences between hourglasses and 12 hour clocks. First, the mechanisms that are involved, and grant the devices their time keeping abilities, are significantly different. The 12-hour clock keeps track of time by some underlying oscillatory mechanism, whereas the hourglass functions not through any oscillatory mechanism but instead by the accumulation of some

⁴⁰ For this reason, when period timers are employed to measure intervals, they are embedded within larger networks that invariably include a memory store (see Wearden (2001)).

medium.⁴¹ Second, hourglasses do not explicitly represent the time of day, rather, hourglasses explicitly represent the temporal interval between two events, a start event (*marker event*) and end event (*target event*).⁴² As a result, to accurately use an hourglass you must be able to recall the earlier event that demarcated the beginning of the timed interval and notice the later event that serves as the end of the time interval.

These lessons, comparing hourglasses with 12 hour clocks, can be extended to period timers and interval timers directly. Interval timers are modeled as accumulation processes that mark the temporal interval between a start event (*marker event*) and an end event (*target event*). As a result, interval timers require that there be a memory store that can recall the pair of events that demarcate the timed interval. Furthermore, what bears the informational content of the interval timer is the quantity of the accumulated medium – most likely in the case of neurally realized interval timers, the accumulated medium is an aggregation of neuron spikes in some memory store.

The function, then, of interval timers is to measure the temporal interval between two events. We can take this description to be giving what Marr (1982) calls *the computational level description*.⁴³ But what we need to really understand how interval timers differ from period timers is the *algorithmic level*. The standard model of an interval timer is given by the *scalar expectancy theory* (Gibbon 1977; Gibbon et al 1984) or more broadly *pacemaker-accumulator timing models* (Treisman 1963).

⁴¹ Wearden (2001) puts the difference in terms of period timers utilizing *qualitative encodings* of time in which different times are represented by different states of a system with a constant number of elements, versus, interval timers that employ a *quantitative encoding* of time in which “longer times are represented by *more* of something” (Wearden 2001).

⁴² One could of course use an hourglass to tell the time of day by starting the hourglass at the beginning of the day, however, the readout from the hourglass doesn't by itself tell you this. In this case time of day is represented implicitly via the explicit representation of a temporal interval.

⁴³ Marr (1982) distinguished three levels of description that were needed to properly characterize and study the visual system – *the computational*, *the algorithmic*, and the *implementational*. While Marr was specifically concerned with vision, the same explanatory schema can usefully be appealed to throughout the study of the mind and brain.

On these approaches the interval timer consists of a *pacemaker mechanism* and an *accumulator mechanism*. The pacemaker is a mechanism that produces some medium (some *ticks of the clock*) at some fixed rate that can be accumulated. In the case of an hourglass, the pacemaker is the upper bulb and the accumulator is the lower bulb with the sand being the accumulated medium. The pacemaker can be modeled as producing *ticks* at a regular rate (Treisman 1963) or may produce ticks at random but with some average rate provided a long enough interval (Gibbon, 1977).⁴⁴ Unlike traditional hourglasses, however, interval timers are thought to involve an additional component – a *gate* or *switch* that controls the flow of the accumulated medium between the pacemaker and the accumulator. This whole process is then rounded out by a memory store that can record the marker event that *opens* the gate between the pacemaker and accumulator and the target event that *closes* the gate. Given that the pacemaker produces *ticks* at a fixed rate (at least when averaged over a sufficiently long interval), the quantity of ticks recorded by the accumulator carries information about how much time has passed since the gate was opened. As a result, we can interpret the system as providing an explicit representation of the temporal interval between two events.⁴⁵

Since the point of this chapter is to present these timing models independently of any particular timing task, we would do well to note some of the flaws in these models. These models seem well suited for *prospective timing judgments* in which the animal knows ahead of time that it should judge how long a specific interval is between two events. To see the reason why the model is useful for these sorts of tasks consider the scenario where I ask you to see how long I can hold my breath. I give you strict instructions that when I start to hold my breath, the timing should commence. I then fidget around for a bit exhale completely and then inhale deeply holding my breath. You can then take the event of my inhaling deeply as the marker event that starts the timer, and then the timer ends when I gasp for breath.

⁴⁴ One can even construct an interval timer in which the pacemaker is an oscillatory mechanism and the accumulator records how many cycles the oscillator goes through.

⁴⁵ Notice that the interval timer proper doesn't represent the events that open and close the gate. Instead, the interval timer merely represents that some interval of some specific length has passed.

However, consider the other scenario in which I never give you instructions. The holding of my breath would have been just one among many events that you may have noticed occurring, but without any specific reason for doing so you wouldn't have started a timer based on noticing my having started to hold my breath. As a result, if given a *retrospective timing task*, in which the animal doesn't know what they should be timing, or even that they are timing anything, the interval timing mechanisms face a problem. We need a reason to open the gate between the ticker and the accumulator. We obviously can't open an interval timer for every event we perceive, but also, it's obvious that many of the timing tasks that are relevant for our behaviour do not involve explicitly prospective timing judgements.⁴⁶

2.4. No dedicating timing mechanism approaches

In order to understand the motivation behind the three approaches to explaining temporal coordination that we're about to discuss we need to take a step back and recall some very general points about scientific theory choice that most people seem to accept.

When confronted with some phenomena that we want to explain, if we have two theories that equally account for the observed phenomena, then, all else being equal, we should accept the *simpler* of the two theories. So far so good. The tricky business comes in when we try and understand what's meant by a 'simpler theory'. We can distinguish between (at least) two forms of simplicity. The first, *ontological simplicity*. The second, *law simplicity*.

A theory is said to enjoy ontological simplicity over another theory if the first theory posits fewer *kinds of things* in order to explain the phenomena. Famously, David Lewis (1986) argued that his theory of possible worlds realism, in which there exists an enormous number of possible worlds (perhaps even uncountably many) that are just as real as our actual world, was ontologically simpler than many of the

⁴⁶ Gallistel (1996) gives this line of reasoning that interval timers by themselves are unable to account for a number of timing tasks.

actualist accounts of possibility that tried to construct ersatz possible worlds out of abstract entities like properties or propositions.⁴⁷ Why was Lewis' possible world realism ontologically simpler? Well, even though the theory required that there exist many worlds, which in a sense seems extravagant, the theory didn't multiply the *kinds of things* out in the world. The non-actual possible worlds are just like our world – they are of the same kind. Along these lines, a psychological theory would enjoy ontological simplicity over another theory if that theory posits fewer kinds of psychological entities. For instance, a theory that posits fewer kinds of representational mechanisms would be ontologically simpler than a theory that posited many different representational mechanisms.

We can say a theory enjoys law simplicity if the laws that explain how the phenomena come about are simpler. This version of simplicity is to a certain degree more difficult to characterize. Simplicity of laws might be understood in the number of laws required to explain the phenomena. If, for instance, a distinct law is required for every state of a system in order to explain how that system evolves over time, then that theory would be *lawfully complex* in comparison to a theory that had a very small number of laws that could cover a wide range of possible states of the system.⁴⁸

Now, when it comes to providing an explanation for the temporal coordination problems, one motivation behind these explanatory approaches is that the explicit timing models are costly in terms of their ontological simplicity. On the explicit and dedicated timer models in addition to the standard perceptual and cognitive machinery that is often posited in order to explain other aspects of animal (and human) psychology, we must also posit an additional type of mental mechanism – a mental timer. And

⁴⁷ Lewis calls this form of ontological simplicity *qualitative simplicity* in contrast to *quantitative simplicity*. Qualitative simplicity is characterized with regards to the number of kinds of entities posited by a theory, whereas quantitative simplicity is characterized with regards to the number of entities posited.

⁴⁸ Now, we run into a difficulty here, given that we have a simple procedure for creating a theory with fewer laws given a theory that consists of many laws. Take the set of laws from the complex theory, now simply conjoin them in a long conjunction. The conjunction is logically equivalent with the set of individual laws, yet the revised theory consists of a single law. Obviously this won't do as a simpler theory. What is needed is some characterization of simple laws – laws that are not merely the result of conjoining individual laws. See Sober (2006) for a discussion of these issues.

not necessarily just one timer, but possibly many. If we could explain the temporal coordination tasks in some way that didn't appeal to dedicated timing mechanisms, then it would seem as though we would have on our hands an ontologically simpler model. We can characterize the two approaches to explaining temporal coordination in relation to the explicit timing models from above by looking at how sparse these *No Timer Models* are in comparison to the explicit models above.

Intrinsic timing models

Of the no-dedicated timing models we'll discuss here, the intrinsic timing models are the only ones that still explicitly represent time. According to these approaches, many of the systems involved in the representation of non-temporal aspects of the world (e.g. the color of a flash of light, the pitch of a tone, etc) have dual-contents. They represent the non-temporal features of the world as well as some of the temporal features of that very same aspect of the world. For instance, the very mechanism involved in the representation of a particular tone will also have states that allow it to carry information about the duration of that tone.

According to these models, the ability to keep track of time is ubiquitous in the brain as the ability to encode the temporal features of the world are intrinsic features of neural systems. In this way, there is no single timing mechanism, but rather, a number of timing mechanisms that are hyper-specific in that their timing capacities are tied to very specific and localized neural systems. The particular features of neural systems that intrinsic timing models appeal to differ from theory to theory. *State-dependent timing models* (discussed in detail in chapter 5) appeal to particular time-dependent spatial patterns of activation in neural systems to encode time (Buonomano 2000; Buonomano & Karmarkar 2002; Finnerty et al 2015; Ivry & Schlerf 2008; Karmarkar & Buonomano 2007). *Efficiency coding models* (Eagleman & Pariyadath 2009) appeal to the efficiency of neural signaling as a means of encoding duration. Finally,

ramping activation models (Lebedev et al 2008) appeal to the time-dependent ramping of neural activity to encode temporal content.

In all cases, whatever intrinsic property the models exploit, the appeal of these models is the ontological austerity of how they describe timing behavior. Dedicating timing models have to posit systems that are in addition to what is required to explain how animals come to represent other non-temporal aspects of their environment, yet intrinsic timing models attempt to explain these timing capacities by appealing to properties of neural systems that we have independent reasons for accepting.

Still these intrinsic models all suffer from similar problems. First, since their operations are tied to the particular neural systems involved in representing specific non-temporal aspects of the world, they have difficulty in explaining certain phenomena like cross-modal transfer (see discussion in Ivry & Schlerf (2008)). Specifically, it is well documented that subjects who are trained to discriminate durations with auditory stimuli show a benefit in their ability to discriminate durations in vision. However, the reverse is not true in that duration discrimination training in vision does not improve duration discrimination in audition (Alais & Cass, 2010; Bratzke et al, 2012).⁴⁹ If, as according to intrinsic timing models, the operations of the individual timing mechanisms are independent of each other, then there simply is no reason to suppose that training in one modality should enhance discrimination in another modality. Furthermore, there is no reason to suppose that there should be any asymmetries between how training benefits transfer across modalities.⁵⁰ As a result, there seems to be no a priori, or theoretical,

⁴⁹ Although see Lapid et al (2009) for cases in which cross-modal transfer from audition to vision do not occur. The data in this study, however, is up for interpretation. Other studies have found that there is a transfer of perceptual learning benefit from audition to vision (Alais & Cass, 2010; Bratzke et al., 2012), however, not only did the Lapid et al study fail to find this transfer, but they were also unable to replicate the within modality perceptual learning benefit that is also commonly found. This indicates that perhaps the failure to find a cross-modal transfer benefit was due to a more general problem with the methodology in their study.

⁵⁰ If we supposed that the training benefits were the result of a post-perceptual decision process, perhaps something akin to what occurs in signal detection theory, then we might expect there to be cross-modal training benefits as a single (type of) decision process might underpin the perceptual decision in the various modalities, and the training benefit afforded by the enhancement of this decision process might cover all of the various intrinsic time keeping mechanisms. However, we then run into a problem of explaining why there are asymmetries in how readily training benefits transfer across specific modalities. Of course one could likely come up with a story that makes sense of this

reason to suppose that training that modifies how one mechanism operates or is interpreted should also influence how other distinct mechanisms operate. Also, intrinsic timing models are plausible for representing only very short durations, as the intrinsic properties they appeal to often are not stable enough to underpin the perception of temporal properties beyond several seconds (at an absolute maximum) (Ivry & Schlerf 2008).

Cued Synchronization Accounts

To get a sense of how *cued synchronization models* explain how animals overcome the temporal coordination problem, let's begin with the picture of the mind as given by models that appeal to explicit representations of time. The mind contains a number of sensory / perceptual mechanisms that produce representation of the world. Once these perceptual mechanisms produce some initial representations of the world, then the content encoded in these perceptual representations gets taken up by cognitive and motor processes in order to accomplish a variety of different tasks. Along the way, the content or information provided by perception and manipulated by cognition is integrated with dedicated mechanisms that explicitly represent time and space, thereby providing the information needed to properly orient the organism within its environment. This story should seem old hat, because in a sense, it is very old hat. This just is the traditional representationalist picture of how the mind works and synchronizes behavior with the environment.

Cued synchronization accounts agree with almost everything in the old hat picture of the mind except for the need for the explicit representations of time. According to these approaches, all that is needed to account for the mind is all of the non-time-specific machinery of the old hat picture. With the

all, however, the main reason being cited in this argument is that the intrinsic timing models by themselves do not help us understand this aspect of our temporal perception.

ability to represent the *what* and *where* of things and events in their environment, minds can solve the temporal coordination problem.

But how could that be? How could one appropriately coordinate their behavior with their environment if they had no way of explicitly representing time? According to these approaches, we needn't form internal representations of time because the events in the world we interact with themselves exhibit temporal regularities⁵¹ – there's no need to represent time because the temporal structure of the environment scaffolds behavior in the appropriate ways needed for the coordination of behavior. For example, one means of explaining the migratory behavior of birds, like the Pied Flycatcher (*Ficedula hypoleuca*), a small European migratory bird, is to appeal to some internal clock mechanism that lets the bird initiate its migration from Sub-Saharan Africa where it winters to Northern Europe where it breeds. Perhaps the clock is capable of telling the length of the day and that allows the bird to know when it should migrate. Or perhaps the bird possesses a longer-period timing mechanism that allows it to know approximately what time of year it is. In either case, the bird's migratory behaviors seem to have a regular temporal pattern to them, and some internal clocklike mechanism could explain that pattern. However, as it turns out, there really isn't any need to posit an internal clock of this sort, since much of the temporal regularity in the bird's migratory behaviors are due to non-temporal cues that the bird is exposed to. Particularly, it appears that the bird's migration is due to noticing the temperature changes starting in its wintering grounds in Africa and its stopover location in the Iberian Peninsula (Both et al 2005).⁵² Since the changes in temperature at these locations have a temporal pattern to them, the bird by cueing into these temperature changes exhibits behaviors that themselves have a temporal pattern to them. The bird is

⁵¹ In some cases, the temporal regularities appealed to might be internal endogenous regularities of the organism, but the question would still remain as to what extent this sort of internal cuing could orient the organism properly without appealing to either a representation of time or the external cuing.

⁵² In fact, the explanation in terms of sensitivity to temperature better explains odd irregularities in the timing of the bird's migration, as the temperature in the Iberian Peninsula has changed less drastically than the temperature in the ultimate destination of the migration. See discussion in Both et al (2005).

successful in coordinating its actions with the temporal structure of its world without *any* temporal representations.

Now, once we have some animal behaviors that can be explained as being scaffolded from the temporal structure of events in the world, we can explain a number of other animal behaviors that use the initial animal behavior (the one scaffolded from some temporal pattern in the world) as a further scaffold. For instance, if birds begin to migrate due to certain changes in temperatures, then predatory behavior is likely to have a temporal pattern that builds off of the temporal patterns of their prey's behavior. This building of temporal patterns on top of temporal patterns can continue. Also, given the sheer number of temporally regular patterns in the environment, primarily the light dark cycle, it's likely that quite a bit of the temporal regularity in the animal world can be explained in this way.⁵³

On these approaches, the animal can still be said to be receiving information about the temporal features of the environment. If, some event, like sunrise, is being represented, and sunrise *always* occurs at a particular time of day, then by having an internal state that is sensitive to the appearance of the sunrise, then that internal state will implicitly represent the time of day given the contextual facts about the environment. Given that certain events correspond with particular times of the day, the representation of some event's occurrence can be used to infer that it is a particular time of day, given an understanding of the temporal relations. However, and this is an important aspect of the no-timer models, in order to act in a temporally coordinated manner, the organism needn't go through the inferential process to extract the temporal information from the representation of events. Rather, merely by representing the occurrence of an event, and knowing what to do when that event occurs (whenever that may be) the animal will exhibit temporally coordinated behavior. All the animal needs to internalize is a series of hypothetical imperatives, or conditional rules, that describe sequences of behavior and cuing events, such as, *<if sunrise, then descend from tree>*. Nevertheless, there is a legitimate sense in which the representation of

⁵³ A variation of this cued synchronization account can be found in Killeen & Gregor (1988).

events can implicitly represent the temporal features of the environment, in that if the animal possesses the appropriate cognitive machinery, then the animal could make the inference and explicitly represent the temporal information.

The cued synchronization accounts have the virtue that if they successfully explain how animal behavior becomes temporally coordinated with its environment, then it will have done so for a small cost. *Nothing* has to be added to the ontology of the mind. The ability to represent the *what* and *where*, an ability that is well studied in all areas of cognitive science, is enough to explain how from the theoretician's perspective animal behavior's seem to track the temporal features of their environments. And as a matter of fact, we can find many temporal regularities in the events in nature that an animal might be sensitive to.

Cued synchronization accounts seem to have some plausibility as it's really easy to find examples where we report the timing of our own behaviors as being the product of external cues. To use a case from my own childhood, as a child, living in the tristate area during the mid-nineties I can remember knowing when dinner was going to be ready without having to look in the kitchen or at a clock. Instead, I had internalized the following rule *<when Jeopardy ends, then help set the table>*. Looking upon my 10-year old self someone might have been taken aback by the temporal regularities that my behaviors showed, "He helps set the table most every night at 7:30PM". One way of explaining this is to posit something like an explicit representation of the time of day that is guiding my behavior – perhaps I look at a clock. However, the cued synchronization accounts have an economical, and most importantly, an accurate, explanation of my behavior. My behavior was cued by a certain TV-show and exhibited a temporal regularity not because I was explicitly representing the temporal features of my environment, but because the environment itself possessed a temporal regularity.

While in the best cases we can easily point to some cueing stimulus that regulates temporal behavior, the vast majority of cases seem to lack obvious candidates for cueing stimuli. In extreme

experimental cases, some timing behaviors persist in conditions in which all possible cueing stimuli have been removed (this line of reasoning is discussed in detail in chapter 3). Furthermore, some behaviors require the ability to keep track of the very contingent temporal structure of events in the immediate environment. With contingent events, perhaps trying to coordinate your behavior with the specific duration of a specific event, there will simply be no *pattern* in the environment to notice and as a result there can be no learning of any conditional rules that could explain how one coordinates their behaviors on the basis of some event in the world. Finally, in many cases, the cues for behaviors are the very temporal properties the representation of which we are trying to understand. For instance, Russell Church and Warren Meck (Church & Meck 1984; Meck & Church 1983) showed that rats could be trained to expect food at a particular location given either an auditory or visual cue of a particular duration. It is the duration itself that is the cue, and as a result, the animal requires some means of gathering information about the temporal structure of the stimuli. Cued synchronization simply can't help in this setting since the non-temporal cues fail to provide the required information to the rat. While cued synchronization accounts seem capable of explaining some behaviors, there is a significant difficulty in understanding how this approach can generalize to explain a wider range of temporal properties.

Non-representational Accounts

While cued synchronization accounts depart from the orthodoxy in the cognitive science of temporal coordination, they constitute only a minor departure from the traditional representational theory of temporal coordination in their denial of any explicit temporal representations. Whatever picture of the mind that emerges from the cued synchronization accounts, we can be fairly safe in claiming that it doesn't constitute a radical departure from the mainstream. However, more extreme approaches are possible. We can attempt to explain how temporal coordination of behavior occurs in the complete

absence of any temporal representations (*local non-representationalism*) – either explicit or implicit – or even more radically, in the absence of any representations whatsoever (*global non-representationalism*).

Both the local and global forms of non-representationalism are motivated (in large part) by parsimony reasons. Local non-representationalists claim that intentional explanations of *certain* psychological phenomena are explanatorily superfluous (Hutto & Myin 2013). Instead, particular temporal coordination tasks are explained in terms of particular organism–environment interactions and the dynamic features of the interaction. If any representational interpretation is possible for the particular system in question, then it happens to be the case that the representational interpretation fails to add to any explanatory goal.

Global non-representationalists motivate their position with many of the same concerns that motivate local forms of non-representationalism. One common approach by globalist non-representationalists is to lay out some conditions that must be satisfied in order for some system to count as representational, and then claim that the mind / brain fails to satisfy those conditions. What criteria are chosen to characterize representationalism varies from anti-representationalist to anti-representationalist and their individual preferences as to what is the most plausible candidate for representationalism.⁵⁴ This project, of showing that nothing in the mind / brain satisfies the conditions needed to count as a genuinely representational system, is given in conjunction with a positive sketch of how explanations should proceed in the cognitive sciences in the absence of intentional states. This leads us to the second motivation, global non-representationalists argue that there are non-representational ways of describing the functioning of the mind which render any representationalist account of how the mind works as

⁵⁴ Chemero (2000) directs his argument against a representationalism styled after Millikan (1989). Hutto and Myin (2013) direct their arguments against an information-theoretic account of representationalism styled after Dretske (1981, 1995). In both cases, the presumption is that the arguments weighed against these versions of representationalism scale to argue against all forms of representationalism as the authors take their target opponents as providing the strongest representationalist case.

explanatorily unnecessary. Furthermore, it's claimed that the non-representationalist explanations provide better more specific explanations / prediction of behavior.

The general feature of all of the non-representationalist approaches is that we shouldn't think of the mind as consisting of systems that manipulates representational states in the way that a computer manipulates computational states. Computation, it is said, occurs in an “abstract” time (van Gelder & Port 1995). The unfolding of the computations has no intrinsic connection with the temporal dynamics of the computation. However, the alternative approach endorsed by the non-representationalists, and especially those coming from the *dynamical systems approach to cognition* (Chemero 2000; van Gelder & Port 1995) understand the mind to be a dynamic system that is continuously evolving over time, and it is this being in time that is crucial to how minds arise from dynamical systems.⁵⁵ Temporal coordination then arises not because there is anything like an explicit representation of time, instead, temporal coordination arises due to the dynamical unfolding of the organism and its environment.⁵⁶

Consider the case in which a batter is receiving the pitch in a baseball game. The pitcher releases the ball, and according to the standard computational story, information picked up by the visual system is used to create a representation of the trajectory of the ball and the rate at which the ball is moving. A motor plan is then constructed instructing the motor systems to swing the bat at a particular moment and at a particular location (or not swing at all). The appropriate command is then sent to the motor systems and the batter swings. The dynamical approach instead would understand the batter as being in a dynamic causal relationship with the oncoming pitch. The light reflected off of the baseball and received by the batter starts a cascade of dynamical interactions between different aspects of the cognitive system – a

⁵⁵ It's important to notice, however, that not all dynamicists do away with the notion of representation. Some, for instance (Bechtel, 1998), view the dynamical account as providing a means of implementing a representational system in much the way that earlier connectionists viewed connectionist architectures as implementing a classical computational system.

⁵⁶ Again, this is only the case for those dynamicists that take the dynamical systems approach to cognition as *an alternative to* representational or computational models. For those that see dynamical accounts as explaining the implementation of a representational system, the adherence to a dynamical model of cognition is not in conflict with the representational explanation of temporal coordination.

feedback loop is produced between motor and perceptual systems. This interaction unfolds over time, and without the need for an explicit representation of when to swing. The successful timing of the swing is a result of the evolution of the dynamical system over time – if the process takes too long, the swing will be late, if the process is too quick, then the swing will be too early.

This is an admittedly sketchy explanation of how temporal coordination occurs in a dynamical system approach, one which surely doesn't do justice to the complexities and sophistication of their explanations, but the goal of this chapter is only to give an overview of the sorts of explanations that are produced by different approaches to temporal coordination problems. To give more details to the picture would require detailing the dynamics of each component cognitive process and detailing the way in which these different components interact. For any given timing task, we can assume that the details of the non-representationalist approach will differ depending on what systems are in causal interaction at a given moment.

It's important to note that the coordination between the contingencies of the environment and animal behavior is brought about by the causal connections that hold between the stimuli that impinge upon the sensory systems and the dynamical systems internal to the organism. In this sense, if the non-representationalist approach to timing tasks is successful, then we will have an extremely economic theory of the mind, but the problems arise in cases where it seems that the environment does not provide sufficient stimulation to properly constrain and scaffold behavior. This same problem arose for the cued-synchronization accounts. So, if there aren't even enough stimuli to allow cued-synchronization accounts to get off the ground, then the non-representationalist versions will also suffer.

2.5. Summing up

While this hasn't been intended as an exhaustive survey of the different explanatory approaches available for explaining how animal behavior comes to be coordinated with the temporal structure of the environment, it nevertheless gives us the background to go forward and investigate the various ways in which animals overcome the temporal coordination problem. In providing this taxonomy of positions, the focus on the distinction between *ways of representing time* (i.e. the explicit / implicit distinction) provides us not only with the basis for forming the taxonomy but for having a better understanding of how mental representations should be understood. In particular, the aspects of mental representation discussed here – how it makes animal behavior intelligible and the requirement that semantically significant states of the system be ones that can exert causal influence on downstream systems – will be put to work in later chapters.

Something to emphasize throughout the project of this thesis is to not take these models to be in all out competition. No single approach is correct *tout court*. Instead, some approaches may be successful in explaining certain psychological tasks whereas other approaches may be successful in others. The only competition that arises between the various approaches one might take is when multiple approaches attempt to explain a single phenomenon. Yet, as we'll see in chapter 5, even cases in which a single phenomenon seems to be under investigation, and thereby engendering a conflict between different models, it may turn out that what initially appears to be a single unified explanandum turns out to fragment into distinct phenomena to be explained independently.

Chapter 3: The Sense of Time

Most of cognitive science and philosophy of mind understands the architecture of the mind as having three main divisions. Sensory systems acquire information from the environment, cognitive systems extract and manipulate information from the sensory systems to allow for flexible thought and stimulus independent thinking, and motor systems translate the information carried by sensory and cognitive systems into motor signals that move creatures. With this general architectural picture in place we can ask about any mental process or system where in the cognitive architecture that process sits – is it a sensory, cognitive, motor, or some hybrid process? This paper will begin by asking where in the cognitive architecture we can find the representation of time.

Our initial starting point is the following: across the animal kingdom, and within any given individual creature, there are likely to be a number of different means by which the temporal structure of the world is represented. It's known that animal behaviors are sensitive to the temporal structure of their world over an incredibly wide time scale – over 10 orders of magnitude (Buhusi & Meck 2005). Humans, aided by cultural, mathematical, and technological achievements, can synchronize their behaviors with even greater time scales of incredible length and with incredibly high precision. In addition to the very different time scales that can be mentally represented, the temporal properties themselves that are mentally represented can be distinguished both semantically and in terms of the mechanisms that are involved in their representation. Semantically, representations of durations (e.g. *being 5 seconds long*) and locations in temporal sequences (e.g. *being earlier than* or *later than*) have a predicative role. They are applied to objects or events in much the same way that we apply predicates like *being green* or *being smaller than*. On the other hand, representations of particular moments in time (e.g. *now* or *March 15, 1992*) have a referential role. These representations refer to moments in time in a way that allows us to predicate properties of those times. Logically, or semantically, the representation of moments in time operates in the same way as other referential representations such as *here* or *Newark, NJ*. Mechanistically,

we also know that various aspects of our ability to mentally represent time come apart from one another since they are underpinned by distinct neural mechanisms.⁵⁷

We know that there are a variety of different means by which we mentally represent time. What we do not know, at the moment, is just *how varied* this underlying machinery in fact is. Given that there is such a variation in how we mentally represent time, there simply is no guarantee, *prima facie* or not, for thinking that all of the various components involved in temporal representation will fall within one or another category of the overall cognitive architecture. So, the question that we should ask is not *where in the cognitive architecture does the mental representation of time belong* but rather *where in the cognitive architecture does a specific capacity for the mental representation of time belong*.

Besides the question of in which division of the cognitive architecture some mental capacity belongs another broad question about cognitive architecture concerns the domain specificity or generality of any given mental process or system. Are the psychological capacities we know from introspection or through the observation of behavior the product of mechanisms or processes that are specific and dedicated to the capacities in question or do they follow from a more general capacity? Are there dedicated motor abilities for walking? Or do the movements involved in walking simply come from a more general motor ability to move one's legs? Similarly, is the sensory capacity for visual form detection due to a task specific mechanism or does the capacity arise from a more general visual capacity, say for the detection of luminance contrast?

In this chapter we are going to ask two questions. First, in which broad category (or categories) of the cognitive architecture will we find the particular capacities for the mental representation of time? Second, within any one of those broad architectural categories, is the capacity due to a task specific mechanism or not? In this paper, I will argue that at least some of our capacity for the mental

⁵⁷ See Buhusi and Meck (2005) and Hinton and Meck (1997) for an overview of how the mechanisms involved in various aspects of the mental representation of time come apart. However, for more detailed discussions see the arguments for the *fragmentary model of temporal perception* given in chapter 4 and 5.

representation of time is due to a genuine *sense of time*. That is, there is a task specific, and dedicated, sensory mechanism for the mental representation of time.

The chapter will go as follows: In section 1, I will provide an account of how to distinguish the senses from the rest of the cognitive architecture in terms of their role in an information processing architecture. In section 2, I will turn to arguments in the philosophical literature that attempt to show that there is not (or simply cannot be) a genuine sense of time. In section 3, I show how none of the arguments against the sense of time presented in section 2 succeed in establishing their intended conclusion. Finally, in section 4, I argue that when we turn to the literature on the circadian systems of mammals with the characterization of sensory systems given in section 1 that we find good reasons for supposing that mammals, as well as many other animals, have a genuine sense of time.

3.1. The senses in an information processing architecture

Folk wisdom, and philosophical tradition stretching back at least as far as Aristotle's *De Anima* (2004), has it that there are five and only five senses – touch, taste, vision, smell, and hearing. However, like many aspects of our folk understanding of psychology, the notion of a sensory system has come to play an important role in our *scientific* theorizing about the mind. Within those scientific practices that study, or appeal to, the sensory systems, it is common to find attributions of senses that go beyond the classic five *Aristotelean* senses. Given the wide range of uses of 'sense', in both folk and scientific contexts, it's not surprising that some philosophers have attempted to argue that the uses of 'sense' in scientific contexts to describe non-Aristotelean senses in a way manages to *change the subject* (Nudds 2004). The senses, they argue, are in some way analytically restricted to the five classic Aristotelean senses.

Now, it may strike some as though the disagreement as to whether there are only 5 or more senses is something of a verbal mistake. Defenders of only applying the term ‘sense’ to the Aristotelean senses simply have a more restrictive notion of ‘sense’ than those scientists that apply the term more liberally. All that is required, it might seem, to avoid this conflict is to simply specify *how* to apply our terms, and there are coherent accounts of the senses that apply to only the Aristotelean senses and coherent accounts that apply to the more liberal usage of ‘sense’. In this section, I will argue that the debate between the uses of ‘sense’ is not a merely verbal dispute and that liberal uses of ‘sense’ are not merely changing the subject. There is importantly a common core that lies at the heart of both the more restricted and liberal uses of ‘sense’ and it is by acknowledging this common core that we gain an understanding of how the senses function in our understanding, both scientific and folk, of the mind. In both folk and scientific cases, the notion of a sense is used to describe those aspects of the cognitive architecture that provide an immediate and ongoing tracking of the world and function as the initial stage in the information gathering process of the mind. The disagreement, however, simply arises in *how deep* into the cognitive architecture the independent sensory systems are supposed to extend and how these sensory systems are individuated from one another.

In order to see this common core, we need to distinguish between two questions that are often confused in the literature.⁵⁸ First, what distinguishes the senses from the rest of the cognitive architecture? In order to answer this question, we need to understand what is distinctive of the sensory systems as a psychological kind. An answer to this question will tell us in virtue of what audition and vision are of the same psychological kind while both being different in kind from motor systems. Second, what distinguishes one sensory system from another? In order to answer this question, we need to know something other than what is needed to answer the first. An answer to this question will tell us in virtue of what audition and vision, while being of the same broad psychological kind, are nevertheless distinct

⁵⁸ Matthen (2015) makes a distinction between these two questions in his overview of the philosophical literature on the senses.

types of sensory systems. As we shall see, part of the reluctance by some philosophers to admit of non-Aristotelean sensory systems is the result of confusing these two questions. Both folk and scientific understandings of sensory systems give roughly the same answer to the first question – the senses, as opposed to other aspects of the cognitive architecture, are involved in explaining how organisms gather information about the world. Where folk and scientific uses of ‘sense’ differ is in how they understand the individuation question. They appeal to distinct means of individuating the system. Yet, I will not argue that any one way of individuating the sensory systems *is the correct way*. Rather, there are a number of distinct ways of individuating the sensory systems that are responsive to our particular explanatory aims, and these means of individuating the sensory systems are closely tied to how to fully cash out what distinguishes the sensory systems from the rest of the cognitive architecture. As a first step towards understanding these questions, let’s turn to some examples of how the senses are attributed in scientific and folk explanations.

One of the overarching goals of the cognitive sciences is to explain behavior. Crucial to this explanatory task is the explanation of how individual creatures are able to coordinate their behaviors with the contingencies of their environments. It is in the service of explaining this coordination that sensory systems are attributed to organisms as a means of explaining the initial and ongoing contact that the mind makes with the world.

Let’s first look at research on the feeding habits of elasmobranch fish (these include sharks, skate, and rays).⁵⁹ These fish have what once seemed a puzzling ability to feed on creatures that are hidden in the sand on the ocean floor. What puzzled researchers about this feeding ability is *how* these fish were able to detect their prey. None of the classic Aristotelean senses seemed to be capable of providing the fish with the information required to locate their prey. However, through a series of nicely controlled behavioral and anatomical studies, it was determined that the information needed to find their prey was

⁵⁹ See Kalmijn (1982) for a detailed discussion of these capacities.

not provided by any of the Aristotelean senses, but was instead the product of a distinct system that could detect the electrical fields produced by the prey. As a result, researchers began to attribute to these fish an *electroreception sense*. That is, they began to attribute a distinct sense that would explain how a certain sort of environmental information could be gathered by these organisms in order to drive subsequent behavior.

A somewhat different example comes from research on *human proprioception* (Proske & Gandevia 2012). Even in the absence of other possible sources of sensory input humans possess the ability to know the relative position of their own body parts. Even in the absence of any light, you can know whether your hand is right in front of your face simply by having a sensation of where your body parts are. As it turns out, this fact, that we can know where our body parts are, cannot be explained on the basis of information from the classic Aristotelean senses, but instead, requires an appeal to a novel source of *sensory information* based in a system that monitors the location of the body. Researchers working on this topic have called this novel sensory system *the proprioceptive system*.

Another, more controversial example, is found in the literature on the human *vomeronasal system*.⁶⁰ It is known that a number of animals possess a distinct sensory system for the detection of the hormones of conspecifics called the *vomeronasal system* (Halpern 1987). Originally, while it was thought that human fetuses had a vomeronasal system (or at least the remnants of one), it was widely accepted that adults lacked a functioning vomeronasal system. It was thought that whatever system the fetus might have had was lost during development. However, more recent research has revealed that adult humans may in fact possess an intact and functioning vomeronasal system. Careful anatomical and behavioral studies (Taylor 1994) have shown that not only is there the anatomical machinery that would be expected for a vomeronasal system, but that activation of this system has behavioral consequences. Importantly,

⁶⁰ A longer discussion of the vomeronasal system and how its existence impacts philosophical discussions concerning the role of phenomenology in the individuation of sensory systems can be found in (Keeley, 2002).

this research shows that adult humans may possess a novel, non-Aristotelean, system for gathering information that is not made apparent through introspection. The discovery of this system was a thoroughly empirical one as there seems to be no notable phenomenal aspect associated with the operation of this system. Researchers in discussing this system have begun to call it the *vomeronasal sense*.

And finally, let us consider a case from science fiction. It's a common device in fiction to describe certain people as having a *sixth sense*. A sensory system that goes beyond the classic five Aristotelean senses and often is used to describe how someone is capable of having some mysterious and supernatural epistemic ability. In the movie *The Sixth Sense*, for instance, the main child actor has the ability to *see* dead people. In other movies, the term *sixth sense* gets used to describe people who can know when a disaster is going to happen or who have an uncanny ability to tell when someone is lying or bluffing at a hand of poker. The details about what is detected don't really matter, but what is important is the context in which a *sixth sense* is posited. In all of these cases, the additional sense is posited when people try to explain how some individual has the ability to know or detect something about the world that does not seem to be explainable by simply appealing to the Aristotelean senses. Nothing about neuroanatomy is involved in these cases. Also, nothing introspective seems to be involved either, as the attributions are often made to a third person. Rather, the basis for the attribution of a novel sense is due to explanatory pressures of explaining human behavior.

It is here where we find the common-core to both the folk and scientific uses of the notion of a sensory system. In all cases, the senses are attributed in situations where we need to explain how the individual organism comes to carry and act upon information about some aspect of the world. If we find that the organism is acting on some contingent aspect of its environment, and we cannot explain how the organism has information about these contingencies by appealing to the senses we already attribute to the organism, then we posit a novel sensory system. However, the senses cannot be distinguished from the

rest of the cognitive architecture by merely claiming that they acquire information about the environment, since other, non-sensory systems do this. The situation is a little more delicate.

The sensory systems do not merely carry information about the environment, they directly track the state of the environment as it changes. They provide an ongoing epistemic contact with the environment. The accurate deployment of a concept in a thought will also carry information about the environment. When I have the structured conceptual thought BEAR THERE NOW, the deployment of these concepts carries information about the state of the world. However, there is a way in which these conceptual resources do not track the world, but are deployed when other systems that do track the world provide the right sort of triggering information. For instance, I might deploy a thought of this sort when I see a large brown fuzzy creature standing on the trail in front of me. Furthermore, none of the information given to me by vision might be in some way identical to the information given to me by my thought BEAR THERE NOW. That is, the information given to me by my thoughts, that there is a bear there, needn't be identical to or reducible to the information given to me by vision, that there is a large brown fuzzy thing there. So, while the deployment of BEAR THERE NOW might carry information that is not found in any other mental representation, it is nevertheless the case that the deployment of these concepts are under the control of information provided by systems that *directly track the state of the world*. Concepts, and other mental representations, might carry information that is not found elsewhere in the cognitive system, but they are nevertheless triggered or come to be deployed so as to accurately track the world, through information that is given through purely sensory means.⁶¹ Sensory systems, are in a manner of speaking, on the frontlines of the information processing system. The initial representations

⁶¹ Some application of concepts might come to track the world through non-sensory triggering conditions. For instance, we may succeed in inferring the existence of something in the world based on purely rational grounds. In this case, we have a situation in which the deployment of a concept accurately tracks the world, but not on the basis of sensory information.

deployed in the sensory systems are triggered through means that do not rely on other representational states for their deployment.

A useful point of contrast here is with the so-called *number sense* (Dehaene 1997). I say ‘so-called’ because according to the understanding of sensory systems as systems that directly *track* the environment, the number sense is not in fact a sense. Dehaene takes the number sense to consist in a variety of systems that allow for rudimentary mathematical thought, such as the ability to quickly tell the approximate number of items in a collection. However, what distinguishes the “number sense” from the genuine senses is that the numerical representations that it produces are subservient to, or triggered by, the information carried by representations in more general sensory systems. When I quickly judge that one plate of cookies has more cookies than another plate (or that a plate has more cookies than the number of cookies I can feel inside of a black box), I deploy some of the systems included in the number sense, but it is on the basis of information gathered by vision that I make my mathematical judgment. The operation of the “number sense” requires the operation of the other sensory systems. A genuine sense would not require information gathered from other systems in this way.

A complication, arises, however. We know that there are a number of cross-modal interactions between sensory systems at even very early levels of processing (Eckert et al 2008; Falchier et al 2002). For instance, auditory information can influence how the visual system segments the visual scene into discrete objects over time (Shams et al 2000). This takes us into the very next problem in giving a general account of the senses. Sensory systems serve on the frontline of the information processing architecture of the mind, but there is a question about *how deep* into the cognitive architecture the sensory systems extend.

Consider a simplified account of the relation between sensation, cognition, and action (figure #1), in which we have the initial transducer systems that feed into a complicated system of perceptual processes that extract information from the activity of the transducers. This perceptual information is then

transferred to cognition for further processing and cross-modal integration, and then it is delivered to the motor systems to drive behavior.




Figure 1: The simplified feedforward architecture. Transducer activity feeds into several layers of sensory / perceptual processing, which feed into central cognition, which ultimately feeds into the motor systems.

On the account of the sensory systems that we have been developing so far, with the senses as the initial stage in the information processing architecture, the distinction between sensory systems and the rest of the cognitive architecture is simple to draw. We simply draw a vertical line prior to the stage at which sensory information becomes integrated in central cognition. The individual senses remain relatively isolated and stimulus driven as we move from transducers to central cognition.

Unfortunately, though, the architecture pictured in figure 1 simply doesn't exist. We have two reasons for thinking this. First, the operation of even “early” visual areas seems to be sensitive to the re-

entrant processing found in “later” visual areas.⁶² Evidence for re-entrant processing has been found in both behavioral studies (e.g. where higher-level task goals influence early sensory processing (Watanabe et al 1998)) and anatomical studies (e.g. neural projections from frontal areas to earlier sensory areas are common (Clavagnier et al 2004)). Second, we know that not only do “later” areas influence “earlier” areas, but that there are a number of cross-modal influences in which we have lateral influences from “early” areas of one system on the “early” areas of other systems. For instance, hearing two clicks will influence how the early visual system processes a flash of light. In the absence of the two auditory clicks, the visual system will parse the flash of light as a single flash. However, in the presence of the two auditory clicks, the visual system will parse the flash of light as two distinct flashes (Shams et al 2000). Furthermore, there are neuroanatomical studies that show cross-modal projections between early sensory areas (Eckert et al 2008; Falchier et al 2002). With all of this informational cross talk and back chatter the resulting picture of the architecture is less like that in figure 1 and more like that in figure 2.

⁶² “Early” and “late” need scare quotes here since once you admit of re-entrant processing, then the sense in which sensory processing can be divided into early and late starts to fall apart as you trace the flow of information in the system.




Figure 2: The complex interactive architecture. The information streams leading from the transducers are not segregated throughout the subsequent sensory / perceptual processing. There is extensive horizontal and feedback influences.

Whatever picture of the cognitive architecture we might adopt has to be one that allows for systems to influence each other at very early processing stages. So, if we are to draw the sensory / cognition boundary at a stage of processing that occurs prior to integration, then we seem to be forced to draw the boundary between sensory systems and the rest of the cognitive architecture at an incredibly early stage of processing – perhaps as far as the sensory transducers. This *might* be an acceptable place to draw the line between sensory systems and the rest of the cognitive architecture, but it seems difficult to maintain in light of how the notion of a sensory system actually gets used. Just look at any textbook on perceptual psychology and they will have chapters devoted to individual sensory systems and in their discussion of these systems they will include stages of sub-cortical and cortical processing.⁶³ The distinction between the sensory systems and cognition seems to be drawn at some intermediary point

⁶³ Really almost any psychology textbook will do.

between the transduction organs, which are clearly sensory, and some cognitive or associative areas that are (by most lights⁶⁴) non-sensory. However, as I will now argue, where we draw the line between sensory systems and cognition is not a fixed one, but is importantly tied up with drawing a line for the individuation of the specific sensory systems. The two questions, while conceptually distinct, come together in a way that each plays a role in determining the other.

How then are sensory systems individuated? What distinguishes one sense from another and what makes my visual system the same type of sensory system as yours or a dog's? Being able to answer this last question is crucial to many scientific projects. It's assumed by developmental cognitive scientists that the visual system (or any sensory system) they study in the developing neonate be the same sort of sensory system, namely a visual system, as the mature visual system whose development they wish to explain (despite the fact that there might be radical differences between the visual systems of the neonate and the mature individual). It is in fact the development of *the* visual system that they study. It's vital to psychology as a whole that the visual system in one person be the same in type as the visual system in other people, even with some abnormalities, in order to make any sort of generalization about sensory systems across the human population⁶⁵. Finally, it's vital to many ethologists that they be able to make generalizations or claims about sensory system types across the species boundary. For instance, many claims about the human visual system come from studies of the visual system of macaques and cats. Even more so are researchers that study the development of very general sensory capacities, for instance those who study things as general as the development of color vision (Braddick & Atkinson 2011; Brown

⁶⁴ Some people think that *everything* the brain does is sensory. These concept empiricists come in many guises, but for all intents and purposes they all agree that even amodal association areas should actually be interpreted as multisensory association areas which importantly do not count depart from processing purely sensory information. Examples of this sort of view are found in (Barsalou, 1999; Langacker, 1986; Prinz, 2002). For a critical discussion of the attempt to view these association areas as being multisensory as opposed to amodal (and therefore non-sensory) see Weiskopf (2007).

⁶⁵ Being of the same type does not require that the visual systems of the distinct individuals be *exactly the same* or *identical*. Variation is allowed. However, there must be something in common that makes the visual systems of distinct individual's *visual systems* and not something else. It's this something in common that we're after.

1990), whose research again requires that sensory systems be typed in a way that individual token sensory systems (the actual sensory systems realized in individual brains) can all be type-identical. If we are to give an account of how sensory systems are individuated (at least in our scientific explanations of the mind), then we must respect these taxonomic boundaries, since these boundaries are crucial to the explanatory tasks for which the senses are posited.⁶⁶

Perhaps one of the most influential recent philosophical pieces on how to individuate the senses is Brian Keeley's *Making Sense of the Senses: Individuating Modalities in Humans and Other Animals* (2002). In his paper he does not distinguish between the question of what distinguishes the senses from the rest of the cognitive architecture and the question of how to individuate the senses from one another, but his paper still provides us with a helpful starting point for discussing the question of individuation. Keeley proposes four criteria that can be used for individuating the senses. They are:

Physics: each sense responds to a distinct aspect of the environment as determined by physics.

Neurobiology: each senses is constituted by a distinct neurobiological mechanism.

Dedication: the sense must have the function of detecting the magnitudes specified by physics.

Behavior: the information acquired must be used for guiding behavior.

If we look at these criteria we find that only the first three, physics, neurobiology, and dedication, can be used to individuate the senses. The fourth, the behavior condition, actually plays no role in individuating the senses. Keeley, in stating this condition, is claiming that whatever system we are considering as a potential sensory system must be capable of driving behavior. Since all sensory systems are posited in

⁶⁶ Of course a philosopher could jump in here and insist that scientists are individuating the senses in a way that is simply incoherent. Scientists, they might say, are simply getting things wrong and they should stop. However, as we'll see there is nothing more incoherent in how scientists individuate the senses as do some philosophers. That is, there is nothing incoherent here, and unless there are outright contradictions or insurmountable theoretical puzzles that emerge from some use of language, then we should refrain from taking a normative approach and policing language use.

order to explain some sort of behavior, this condition does not distinguish between sensory types.⁶⁷ It is really the job of the other three criteria to individuate the senses. Let's take the simple (or at least highly studied!) examples of vision and audition and see how these criteria are supposed to function.

Vision and audition respond to distinct types of energy in the environment.⁶⁸ Vision responds to electromagnetic radiation within a certain range. Audition responds to mechanical energy. However, it is unclear whether the physics criterion is enough to individuate the senses since both vision and audition seem to be responsive to the same energy types as other sensory systems. The visual system can also be activated by mechanical energy (use your finger to push on the side of your eye). Also, the visual system, while it is responsive to electromagnetic energy, it is not the only sensory system that is responsive to this feature of the world. Thermoception capacities in animals like pit vipers (Buning 1983; Gracheva et al 2010) also are responsive to electromagnetic energy, although of a distinct range. If we are simply individuating sensory types in virtue of the energy types, as defined by physics, that they respond to, then our visual systems would be of the same type as the thermoreceptive system of pit vipers. Perhaps one could insist that they are of the same type, although any insistence of this sort obviously seems strained. Turning to audition, similar worries arise, and in this case, they arise within humans. The mechanical waves detected by audition can also be felt (if they are strong enough), yet it is clear that we should not admit that audition and touch are in fact one sense. That would flout any of the distinctions that appeals to sensory systems require.

⁶⁷ The behavior condition plays a specific role for Keeley. He's concerned with the role of phenomenology in understanding the mind and focuses on theories of sensory individuation that give a central role to phenomenology. However, as we'll see later in this chapter, Keeley gives reasons why phenomenology cannot be used to individuate the senses, when they are used in scientific contexts, because while we attribute sensory systems to a range of animals we simply have no idea what their phenomenology is like. So, instead of using phenomenology to understand the senses we must understand them in terms of the behavioral effects they have on the organism. It is here that the behavior condition becomes crucial. Sensory systems are not just systems that are responsive to the environment, but they are ones that can drive animal behavior. In humans, this sometimes comes along with some distinctive phenomenology, as in the case of the Aristotelean senses, but we needn't tie the sensory systems to this phenomenology. Rather, it is their ability to drive behavior, along with the other criteria, that makes something a genuine sense.

⁶⁸ The arguments in this paragraph and the next are closely related to those given by Macpherson (2011).

So, the physics criterion by itself won't work. But perhaps we can augment it by adding in the neurobiology criterion. Audition and touch might both respond to the same types of energy. And so do vision and thermoception. However, in both cases we have very distinct neurobiological systems at play. Take vision and thermoception. Both involve radically different sorts of transduction organs. The visual system of all animals have transducer mechanisms that employ photosensitive pigments whereas the heat sensors of the pit viper actually involve temperature changes in the receptor membrane that directly influences ion channels leading to changes in neural activity (Bullock & Cowles 1952; Gracheva et al 2010). Similar considerations apply to the audition / touch case. Both involve significantly different sensory organs. And yet, we again find reasons for thinking that neurobiology isn't enough. Take vision. The retina, the end organ of the visual system, is not a homogenous mass of photoreceptors. In fact, the retina of most humans contains 4 standard photoreceptor types – rods, S-cones, M-cones, and L-cones (Solomon & Lennie 2007), as well as photosensitive retinal ganglion cells that possess their own photosensitive pigment, melanopsin. But, not all human visual systems contain this array of photoreceptor types. Dichromats lack one of the cone types. Monochromats lack two of the cone types. Anomalous trichromats, while possessing three distinct and functional cone types, have photoreceptors that are tuned to different wavelengths than the cones of the typical human. Even more odd are cases exist in which the female offspring of a male anomalous trichromat will sometimes have 4 functional cone types (Jordan et al 2010; Jordan & Mollon 1993). All of these cases point to sometimes radical differences *in the visual systems of humans*. The neurobiology is different in all of these cases. If we consider non-human visual systems, then things get even more complex. Animals across the animal kingdom possess a wide variety of different photoreceptor types, mantis shrimp have up to 16! (Chiao et al 2000; Cronin et al 2001; Marshall & Oberwinkler 1999), and these animals all possess different systems that process the initial sensory information. And yet, in all of these cases, we still often describe these various sensory systems as all being of the same type. That is, despite their neurobiological differences we still consider them visual systems.

But, what binds the visual system in any single animal as being *a single system*? Why not consider the distinct photoreceptor types as giving rise to distinct sensory systems? Perhaps we could conceptualize the human visual system as containing 4 distinct sensory systems (one for each photoreceptor type), or 2 sensory systems (one for the rods and another for the cones), or perhaps we could distinguish sensory systems according to their projections to the LGN (so, there could be three systems, the magnocellular, parvocellular, and koniocellular sensory system). Perhaps one approach to unifying the visual system into a single sensory system would appeal to the similarities between the photoreceptors. One salient similarity is simply that they all employ opsin molecules to convert light into neural activity. In this regard there is a neurobiological fact that would group all of the distinct types of photoreceptors as being of a single kind. But, opsin molecules are also found in the pineal gland, so the presence of opsin molecules is not sufficient either for making the visual system distinct from other parts of the cognitive architecture (Peirson et al 2009; Velarde et al 2005). Furthermore, we could equally emphasize the differences in these biological system. They all employ opsin molecules that have peak sensitivities to distinct wavelengths of light. They also enter into different post-receptor processes (e.g. the M and L cones enter into the magnocellular and parvocellular systems, while the S cones enter into the koniocellular system).⁶⁹ Maybe another neurobiological similarity can be found that appropriately individuates sensory systems such that we have *a sense of vision* that is of the same type as the *sense of vision* possessed by a dog. We can actually stay neutral with regards to this point.⁷⁰

If we add the dedication criterion things start to change. We can nicely describe the photoreceptors in the retina as having a common function. They are all dedicated to detecting luminance and color, or dedicated to detecting the visual features of the world, or dedicated to the detection of distal

⁶⁹ There might be some reason to doubt this initial parsing of the human visual system into distinct pathways of this sort as there may be reason to think that there is a mixing of cone type information in the distinct retinal-LGN pathways in infancy. See Dobkins & Anderson (2002).

⁷⁰ My own two cents on this are that I believe that quite a bit can be made from this neurobiological criterion for sensory individuation.

objects and their properties. In this way, all of the photoreceptors can be grouped within a single sensory system. Now, audition also has the function of detecting the distal world, and yet, we needn't group it together with vision since they differ with regards to the other two criteria.⁷¹

We seem to have something that works here. We can individuate the senses by employing a host of different criteria, each of which is insufficient for the task when taken individually, but jointly succeed in distinguishing the sense. However, the success of these criteria for individuation depends on where we draw the boundary between sensory systems and the rest of the cognitive architecture. If we draw the sensory system / cognition border immediately after the transducers, then we run into problems with all three criteria. The physics criterion becomes difficult to maintain since, just taking vision as an example, all of the photoreceptor types while responsive to electromagnetic energy are responsive to distinct ranges of the EM spectrum. Similarly, the neurobiology criterion becomes difficult to maintain as the distinct photoreceptors are all distinct. And finally, the dedication criterion becomes particularly difficult to maintain as it's unclear what single dedicated function each of the different photoreceptors possess. Saying that they are all for seeing the world won't work, since that presupposes the very thing that we are trying to understand. We would need a non-question begging understanding of what it is to see the world and such an understanding could not employ the fact that seeing is done by vision. An evolutionary account of function also won't work as the distinct photoreceptors emerged various points in our evolutionary history in response to evolving ecological needs (Jacobs 2009). Similarly, if we adopt a historical approach to determining function (or for determining some neurobiological criterion), then we run into problems in that different evolutionary paths that stem from a common ancestor were driven by radically different ecological situations (e.g. the octopus and humans) (Gehring 2014). If we draw the

⁷¹ How we cash out the function of the sensory systems is actually a very tricky issue to figure out. Akins (K. Akins, 2014) nicely argues that in many cases we cannot understand the sense as picking out distal features of the world. Instead, she argues that the sensory systems track *narcissistic properties* that importantly are not found out in the mind independent world described by physics but are instead only found in relation to the organism doing the sensing. This is a thorny issue, especially for an account of the senses like Keeley's, but as we'll see we needn't decide this matter for what is going to come in this chapter.

sensory system / cognition boundary at the periphery, then we seem to have to individuate the senses at an incredibly fine-grain which would do damage to divisions that we wanted to respect.

But, of course, nobody would draw the sensory system / cognition boundary so close to the periphery. We need to look a little deeper into the information processing architecture (with regards to figure 2, we need to move the vertical division separating sensory systems from cognition farther to the right). Yet when we do that we run into further problems.

Once we begin to draw the boundary between sensory systems and cognition further into the processing stream we get further difficulties in individuating the senses. We want to avoid the problem of carving the senses too finely, but at this point we now have to deal with the ubiquitous cross-modal interactions that happen between cortical processes. A purely neurobiological account may be able to work if we choose to provide a coarse-grained account of neural anatomy. Audition might be just that sensory system that extends from the ears through the auditory cortex. Similarly, vision could just be that sensory system that extends from the eyes through the visual cortex. Now, in order to do this there better be a non-question begging way of individuating the auditory and visual cortices that do not appeal to a prior definition of the sensory systems. Perhaps they are the regions that receive input only from particular transducers, but it's unlikely that that criterion could hold, since we know that cross-modal interactions influence processing in early sensory areas. Perhaps there is some other anatomical marker that could cordon off the appropriate cortical areas, but the worry here is whatever marker we adopt would likely also separate sub-regions of the traditionally defined auditory and visual cortices.

Considering just the visual system, we find a number of distinct anatomical structures in various areas of the visual system (e.g. the striation in V1 and V2, the blobs and interblobs of early visual areas, various retinotopic maps vs tonotopic maps, etc), however, none of these anatomical structures are capable of carving off the *entirety* of the visual system. If we took these anatomical structures as means of distinguishing sensory types, then the visual system would fail to be a single sense. If we take a more

general type of anatomical structure, for instance the presence of projections from the thalamus, then we end up grouping together disparate sensory systems as the presence of these projections is common throughout the various sensory cortical areas. The result is either that we lump together supposedly disparate sensory systems, or we again carve them too finely.

The problem of jointly using many criteria to individuate the senses is that at any level that we carve off the sensory systems from cognition, at least one of these criteria seems like it will carve the senses too finely. It will distinguish between types of sensory system, when we want to think of those sensory systems as being of a single type. Adding further criteria will simply be of no help in those situations since they can only help to carve things even more finely.

Keeley's aren't the only criteria in town, though. There is another approach to individuating the senses that has a long and venerable tradition within philosophy according to which the senses are individuated according to some *introspective criteria*. A number of views fall under this heading, but they all share in common the idea that what is distinctive of each individual sensory system is some property of experience that is available to us introspectively. According to some, the senses are individuated by their distinctive contents (Aristotle 2004; Dretske 1995; O'Dea 2011). According to this *content criterion*, audition and vision are distinct sensory systems since they represent different features of the environment and this representational difference is something that we can be aware of through introspection. According to others, the introspectible property that individuates the senses is the *phenomenal character* that is associated with the individual senses. According to this *phenomenology criterion*, it may be the case that vision and audition might both represent the same sort of property in the world, they may both, for instance, represent the spatial location of an object, but they differ in that *what it is like to represent the spatial location in audition* is different than *what it is like to represent the spatial location in vision*.⁷²

⁷² Both sorts of introspective approaches to individuating the senses have met resistance. Against the content approach see (H. P. Grice, 1962; Lopes, 2000). Against the phenomenal approach see (Keeley 2002; Nudds 2004, 2011).

Note how different these introspective approaches are to Keeley's. Keeley was attempting to give a third-person account of the sensory systems, while these approaches are clearly first-person accounts⁷³. This difference will be important in what immediately follows.

One main motivation for the introspective accounts is that they seem to do well with the Aristotelean senses. It's common across cultures that people acknowledge the existence of the Aristotelean senses. This widespread acknowledgement happens even though there are varying states of scientific knowledge about the mind. In fact, even wildly misinformed accounts of the mind / brain, such as Aristotle's view that the brain cooled the blood, do not seem to cause problems with acknowledging the existence of the Aristotelean senses. How else, then, might all these various people in various epistemic situations be coming to make the same distinctions? Well, the front runner of an answer seems to simply be that it's clear in our experience of the world.

While perhaps this approach is helpful for discussions of how folk psychology might individuate the senses, it undermines a large portion of the scientific usage of the senses. As Keeley argues, if the notion of the senses is to play a role in our scientific understanding of the mind, then we need to divorce our account of the senses from phenomenology, since we readily attribute senses to animals while we have no access to their phenomenology. Second, we attribute senses to humans as well, like the vomeronasal system, that have no phenomenology. If the senses were individuated according to their introspective qualities, then we have an argument for why these non-phenomenal senses would fail to be genuine senses. The argument is as follows⁷⁴:

1. X is a sense with non-phenomenal sense. [assumption]

⁷³ There's another way that the difference between these approaches could be characterized. The introspective accounts are giving metaphysical accounts of sensory systems that also builds in an epistemic criterion according to how we determine whether some experience is given to us by one or another sensory system. The approach by Keeley doesn't necessarily build in any epistemic claims about how we should go about deciding whether two systems are of the same sensory type or not.

⁷⁴ The formalization here makes the argument much clearer than could be done with prose.

2. If something is a sense, then it must be a sense of a particular type. [senses as types]
3. If a sense is of a particular type, then it must have a distinct introspective qualities particular to that sense-type. (introspective individuation)
4. X has no distinctive introspective qualities. [from 1]
5. X is not a sense of a particular type. [from 4, 3]
6. X is not a sense. [from 5, 2]

If the individuation of the senses is based on some introspective property, then no non-phenomenal system could count as a genuine sense, and as a result, large swaths of the empirical literature concerning the senses would be misguided. A way of understanding this might be to say that if we individuate the senses according to their introspectible qualities, then we must draw the boundary between the senses and the rest of the cognitive architecture late enough in the processing so that the sensory systems will include whatever processes are involved in conscious experience. But we simply don't know where to draw this boundary in humans, and we have even less of a handle on how to draw this line when it comes to non-human animals. On top of that we simply have no idea how to determine whether there is some special introspectible quality that would individuate sensory system in non-human animals, since we do not have access to their conscious states.

We seem to be running into an impasse. The folk notion of sensory systems, that philosophers have typically analyzed, seems to rely on some introspectible quality and therefore draw the boundary between sensory systems and cognition rather late. However, this approach to individuating the senses simply doesn't apply well to non-human cases, and as a result undermines a swath of cognitive science that discusses the sensory systems of non-human animals. Perhaps this is where the *changing of the subject* objection arises against the third-person approaches to individuating the senses. However, if there is a core notion of sensory system that applies to both folk and scientific account of the senses, then we

can see this disagreement not as one of changing the subject, but there simply being multiple acceptable means of individuating the senses.

The argument for there being a common core notion of sensory system builds on the above characterization of senses as the initial stage in an information processing architecture – the senses for both the folk and scientists are those systems that gather information and directly track the world. Furthermore, in both folk and scientific cases, we as a matter of fact can individuate sensory systems without appeal to any introspectible quality, and therefore, there must be some other sort of feature that we can latch onto. Therefore, the above argument for the non-existence of non-phenomenal sensory systems fail (that is, premise three of the argument turns out to be false).

As mentioned above, it appears as though in folk contexts the senses are individuated according to their introspectible qualities. Simply reflect upon your experience and the divisions between the senses should be apparent. However, the folk notion of the senses is permissive enough to allow for the proper typing of sensory states as being of a certain sense despite there being no phenomenology associated with those states. A brief foray into recent history makes this clear enough. In the 1950's, in the midst of Cold War paranoia, the idea of subliminal perception was capturing the public attention. The fear was that the broadcasting of “foreign ideologies” through television, radio, and film could influence behavior without the public becoming aware of what was happening (Acland 2012). The fear was so strong that the National Association of Radio and Television Broadcasters asked its members to refrain from using subliminal messages in their broadcasts. Despite the fact that these subliminal messages were by definition something beyond our conscious experience and therefore had no associated phenomenology, people were still capable of conceptually distinguishing and being fearful of auditory vs visual subliminal messaging. The folk could do this. To put it differently, the folk could distinguish between subliminal perception of different sensory types. But for this to be possible, the typing of the subliminal experiences could not rely on the existence of any phenomenal or otherwise introspectible quality. The folk must have

a means for typing or individuating senses that appeals to a third-person available feature. This isn't to say that in some cases the senses are individuated by the folk in terms of some introspectible quality, but rather, it's simply to say that this can't be the only means by which the individuation occurs.

Premise three of the argument for the non-existence of non-phenomenal senses fails. There must be some other means that even folk senses can be individuated. What that is, we can remain neutral.⁷⁵ However, importantly the core notion of the sensory system remains constant. The senses play a role in directly tracking and gathering information about the world.

It is here where we can finally put the point. The two questions that began this section, what makes something a sense as opposed to some other psychological kind, and what makes one sense a different type than another sense, can be answered in a variety of ways. Furthermore, the two questions are not entirely independent. Given where we draw the boundary between sensory systems and cognition and motor systems will influence what sort of accounts of individuating the senses are plausible. That is, given a particular boundary between sensory systems and the rest of the cognitive architecture, certain accounts of individuating the senses will provide us with a particular taxonomy of the senses. However, if we adopt a different sense / cognition boundary, then that same account of individuating the senses will result in a different taxonomy of the senses. In some cases, we may wish to push the sensory boundary towards the periphery, and employ some third person criterion for individuating the senses, resulting in a likely very fine-grained taxonomy of sensory systems. In other cases, we may push the sensory boundary deeper into the architecture and adopt a first-person criterion, thereby resulting in a taxonomy that approaches (or is identical to) our folk notions. Furthermore, a range of intermediary positions are available. If you want to study how the visual system has developed in mammals, then you may draw the

⁷⁵ A plausible approach would be to employ some rudimentary (neuro)biological criterion. The folk can distinguish outward sensory organs and perhaps they use this to individuate sensory systems. Another approach is that of Nudds (Nudds, 2004, 2011) where the individual sensory systems are individuated by the epistemic activities involved in using those senses. Vision, according to Nudds, is for seeing. Audition is for hearing. Touch is for feeling or touching. Olfaction is for smelling. And taste is for tasting. Or perhaps there is an appeal to something like dedication or the content criterion.

sensory boundary at some intermediary position and draw upon third person criterion. The options are many. And yet, we have no reason to suppose that any of these combinations *gets it right*.

The brain is real and the operations in the brain are also real. However, how we produce a taxonomy of the brain and its operations are up to the goals of the researchers. Importantly, this isn't to adopt some form of radical anti-realism about the mind. Rather, it's admitting that there are real patterns in the world and we can group these together in a variety of different ways that track real clusters of behaviors and phenomena in the world. What dictates our general category of *sensory system* is a common-core commitment to the idea that the sensory systems are directly tracking and gathering information about the world. It is in this way that scientific discussions of non-Aristotelean senses do not fall prey of merely *changing the subject*. They maintain the important core, but are concerned with slightly different explanatory aims, in much the same way that the folk in different contexts will appeal to non-Aristotelean senses and non-introspective individuation criteria in the explanation of behavior.

3.2. Against the sense of time

Now that we have an understanding of the information gathering role of the senses in our theorizing about the mind we are in a position to see the arguments that some philosophers have put forward arguing that there simply is no (or cannot be) a genuine sense of time. It's somewhat interesting that it is actually fairly common to find people, both philosophers and scientists, claiming that there is no sense of time, however, it is much less common to find those claims backed up with any significant argument. So, the two arguments that I'll present here are ones that are not only explicitly found in the literature, but they are also arguments that seem to be behind some of the un-argued for claims that there is no sense of time.⁷⁶

⁷⁶ These two arguments aren't the only arguments that could be found in the literature. One could imagine someone running an argument against the existence of a sense of time on the basis of some general characterization of the

The Non-Causality Argument

The following quotes by the psychologist Lera Boroditsky and the philosopher Mohan Matthen give the general idea behind the non-causality argument.⁷⁷

All of our experience of the world is physical, accomplished through sensory perception and motor action. And yet our internal mental lives go far beyond those things observable through physical experience: we invent sophisticated notions of number and time [...] So how is it possible that physical organisms who collect photons through their eyes, respond to physical pressure in their ears, and bend their knees and flex their toes in just the right amount to defy gravity are able to invent and reason about the unperceivable and abstract?

(Boroditsky 2011, p. 333)

Should [the systems responsible for the representation of time] be regarded as transducers for a sense of time? That is, do periods of time cause them to emit a pulse that carries information about these periods of time? Both sides of the question can be argued. A negative answer might be a reason to exclude the sense of time.

(Matthen 2015, p. 573)

As I described earlier, an essential characteristic of the senses, which distinguishes them from other representational systems, is that the senses track the ongoing changes of the particular aspects of the world that they detect. In other words, the senses provide us with a *real-time information link with the*

sensory systems. Perhaps, by appealing to a particular account of sensory individuation, like Keeley's physics criterion that can be spelled out in terms of there being a causal connection to the world that individuates senses. However, as we saw there isn't any single means of individuating the senses, so any sort of argument of this sort will likely fail. The two sorts of arguments to be presented here both accept the general characterization of the senses as information gathering mechanisms that was argued for in the previous section.

⁷⁷ Similar arguments can also be found in Gallistel (1996) and Coull (2011).

environment. It is in virtue of the epistemic connection with the world given to us through the senses that we can then go and correctly deploy many of our other representational and conceptual systems. If we look at the classic Aristotelean senses, we find that this information link with the environment is established via a causal connection with the aspect of the environment that each sense detects and the sensory systems themselves. For example, it is through the causal interaction of photons impacting the retina that the visual system is able to acquire information about the changing environment. Similarly, for many non-Aristotelean senses, such as proprioception or electroreception, it is through the causal interaction between the to-be-detected aspects of the world and the sensory systems themselves that the appropriate information link is established.

However, it is commonly claimed in the literature on the metaphysics of time that time itself is not causally efficacious, as a result, there simply could not be a sense of time since no sensory system could ever stand in the right causal relation with time.⁷⁸ Putting the argument explicitly, we get the following:

1. The senses must keep track of the immediate environment in real time. [*assumption*]
2. The only way of keeping track of the contingencies of the immediate environment is to be in causal contact with those contingencies in the environment. [*the causality constraint*]

⁷⁸ See Benovsky (2012) for a discussion of the debate in the metaphysics literature on whether time can be causally efficacious. Benovsky argues that time in fact must be causally efficacious to explain the half-life decay of a number of elements. However, see (Lewis, 1973; Maudlin, 2002; Newton-Smith, 1980) for arguments that in cases where it appears as though time is in fact causally efficacious there is in fact a more parsimonious reading of the causal relations that does not require that time is actually causally effecting anything. To give just one example, someone might leave the potato salad out on a hot day and the salad spoils. We might say, the cause of the salad spoiling is the duration of time that it was in the hot sun. But, Newton-Smith (1980) for instance, argues that we can interpret the spoiling of the salad as not being an effect of the time in the sun, but as a chain of molecular events that bring about the spoiling (e.g. the bacteria in the salad feed on the ingredients of the salad, and thereby multiply, etc). This sort of story, while occurring in time, does not appeal to time, or any temporal properties themselves, as causes of any of the subsequent states. Furthermore, many have taken Shoemaker's (1969) thought experiment concerning frozen worlds as showing that the mere passage of time isn't causally efficacious. The details of all of these arguments trade not only on how to provide understand time but also importantly how to understand causation. Making sense of these issues would be a project all to itself.

3. If there is a sense of time [i.e. a sense that keeps track of the temporal structure of the events in the world], then we must be in causal contact with time. [*from 1 & 2*]
4. Time is not itself causal. [*the non-causality assumption*]
5. Therefore, there is no sense of time. [*from 3 & 4*]

Now the purpose of this chapter is to understand something about the senses and not about the metaphysics of time. So, we can simply take the claim that time is not itself causally efficacious as a fair assumption. Furthermore, if we were to abandon this claim about the causal inefficacy of time, then the arguments considered here would simply fail out of the gate. But, by giving them their space, and to evaluate them for what they say about sensory systems and not about the metaphysical nature of time, we actually can gain insight into a very peculiar aspect of how we mentally represent time. So, in the subsequent section when I consider how this argument fails, I will focus on *the causality constraint* and not *the non-causality assumption*.

The Integration Argument

Kant famously argued that intuitions of space and time serve as pre-conditions for the possibility of any experience of an objective world whatsoever. The central idea being that it is by embedding experiences, or the deliverances of the individual senses, into a spatial and temporal framework that one is able to have experience of and make sense of the outer world. In this picture, the representations of space and time are in an important sense prior to the individual senses, since it is only by embedding experiences within a spatio-temporal framework that we can have experience of an objective world, and as a result, the representation of space and time are not proper parts of any of the individual senses. While it's difficult to evaluate Kant's own arguments for the priority of temporal representations, Matthen (2014) has recently argued that while Kant's arguments likely fail, a Kantian-inspired argument can be

run where the representation of time is not sensory since it provides a framework for the interpretation of the individual senses.

The starting point for Matthen's argument is to notice that, as a matter of fact, representations of time serve as a "common measure" for the deliverances of the individual senses. The information acquired by the individual senses is ultimately organized within a unified temporal framework to create a unified representation of the external world.⁷⁹ For instance, when I see and hear a thunder storm, the flashes of lightning that I see and the crashes of thunder that I hear are placed within a single temporal order, so that at times they are simultaneous, but at other times the lightning precedes the thunder.

Importantly for Matthen, not only does the representation of time serve as a common measure for the various sensory systems, but the way in which temporal properties are attributed to events in the world gives us further reasons for thinking that the representation of time is at one remove from the sensory systems. Adopting the account from Phillips (2010), Matthen argues that temporal properties are attributed to events in the world through a *metaexperience*. The individual sensory systems, in tracking the world, bring it about that the creature has certain experiences that are specific to those individual senses. Furthermore, these experiences themselves are temporally structured in that the experiences themselves occur at particular moments in time, have durations, and stand in temporal relations to other experiences and events in the world. The attribution of temporal properties to the events detected by the individual senses occurs by some system that notes the temporal relations that hold between the experiences of worldly events (e.g. *earlier than*, *later than*, and *simultaneous with*), and then these temporal relations are *exported* and attributed to the events out in the world originally detected by the individual senses. It is through these metaexperiences that operate over the individual senses that we are able to coherently integrate the various deliverances of the individual sensory systems into a single unified temporal order. Since the temporal exportation process operates over the individual senses it is not

⁷⁹ This sort of integration is the focus of chapter 5.

itself a sensory system nor is it a part of any of the senses. In this way, Matthen argues that due to the role of temporal representation in structuring and ordering the deliverances of the individual sensory systems we have reasons for denying that the representation of time is due to a genuine sense of time.

We can put the argument more explicitly as follows:

1. The deliverances of the individual senses are organized / integrated within a temporal framework of *earlier than*, *later than*, and *simultaneous with* relations. [*time as organizing framework*]
2. The represented order of events is due to a mechanism that tracks the temporal relations between the experiences of those events, and then the temporal relations that hold between those experiences are exported and attributed to the experienced events. [*temporal exportation*]
3. Therefore, the attribution of temporal properties is through a mechanism that takes experiences as its inputs. [*from 1 & 2*]
4. The senses do not take as their inputs other sensory systems or representational states of the organism. [*established in the previous section*]
5. Therefore, there is no sense of time. [*from 4*]

In what follows, we'll see that the argument as given by Matthen fails to establish its intended conclusion since as it stands it simply isn't valid. Even if all of the premises are true (although, see chapter 4 for an argument that premise 2 is false), the conclusion still would not follow as there may be other mechanisms for temporal representation that do not fit the story given here. But, with these two arguments on the table, let's now turn to see why they both fail.

3.3. Towards a sense of time

In this section I will argue that both *the non-causality argument* and *the integration argument* fail to establish that there is no sense of time. Once we notice how these arguments fail, we will be in a

position to determine whether the existing psychological / neurological mechanisms possessed by animals provide them with a genuine sense of time.

Against the non-causality argument

As we've been discussing, the senses are distinguished from the rest of the cognitive architecture by their information gathering role. They directly track and gather information about the ongoing contingencies of the environment. As some have described them, the sensory systems are servile to the environment (Akins, 1996) in that the states of the sensory systems are importantly tied to the ongoing and changing state of the environment. When we open our eyes, the sensory system cannot help but begin to respond to the environment. The same goes for touch and the various other sensory systems. None of this is to say that the organism itself doesn't contribute to the process of perception. But, it is nevertheless the case that at the very periphery of the sensory system the senses are servile to the environment.

It's for these reasons that a causal connection between the environment and the senses seems to be necessary. How else might a sensory system keep track of the changing contingencies of the environment if not through a causal connection? There's a sense that if there were no causal connection, then any correlation between the states of the sensory system and states of the world would be mysterious. Again, if we look at the classic Aristotelean senses we find that the capacity that the senses have to keep track of the environment as it changes and as we move within the environment is due to the causal sensitivities of the sensory systems. If time simply can't be a cause, then we can't be in the appropriate causal connection. Notice how strong this argument is. It's not simply that as a matter of fact some creature lacks a sense of time. Rather, the conclusion drawn from the non-causality argument is that there simply *could not be by necessity* a sense of time.

However, as Dretske emphasized in his *Knowledge and the Flow of Information* (1981), for there to be an informational link between a representational system and some state of the world (or in Dretske's terms, between a signal and its source) a causal link need not be present. For a system to carry

information about some state of the world simply requires that there be some nomic relation between the state of a representational system and the state of the world that meets the following condition:

A system, *s*, in state *R* carries the information that *t* is *F* (where *t* is a particular instantiating the property *F*), just in case, *t*'s being *F* is guaranteed by *S*'s being in state *R* (given appropriate background channel conditions⁸⁰ and the laws of nature).⁸¹

In other (more epistemically loaded) words, the fact that *s* is in state *R* allows one to know that *t* is *F*. Often what guarantees this nomic relation is a causal connection. However, if the appropriate information link can be established in the absence of such a causal connection, then we would still have the appropriate information link.

Dretske gave an example that tried to show that an informational link between two states of the world, say *s* and *t*, needn't depend on there being a causal influence from *s* to *t* (or vice versa). His example is the following: Consider a pair of TVs, *TV1* and *TV2*, located in isolated rooms and connected to a single closed-circuit television network. Both TVs are capable of only showing the signal that is being sent from a single central broadcast (see figure 3). Now, despite the fact that *TV1* and *TV2* exert absolutely no causal influence on one another, there isn't even any physical means by which they might do so, there is still an information link between them. We know that *TV1*'s being in state *I* carries information about the state of *TV2* since the two TVs are being controlled by a single source.

⁸⁰ These channel conditions are the physical states of the world that when in place mediate the nomic relation between *s* and *t*.

⁸¹ This definition could also be given in terms of conditional probabilities. A system, *s*, carries information about *t* if the states of *s* eliminate some of the possible states of *t*. Dretske in giving his analysis of semantic or representational content in terms of information required that the conditional probability of *t*'s being *F* given that *s* is *R* be 1 since he needed to rule out cases in which *s*'s being *R* is an ambiguous signal that could indicate either *s*'s being *R* or some other state of affairs entirely such as *q*'s being *F*.




Figure 3: Common-cause information channel. Notice that there is not causal connection between TV_1 and TV_2 yet they nevertheless carry information about one another.

Now, Dretske's example shows that there needn't be a causal connection between the information carrying system and the state of the world that it carries information about but it does so by appealing to a common cause. However, in what follows, we'll see that this common cause condition doesn't even seem necessary.

Since a causal connection isn't necessary for an information channel to be established, one can understand the non-causality argument to be a sort of abduction argument. Since the senses put us into real time informational contact with the environment, and causal connections seem to be the only means of establishing these connections, then it seems as though there can be no sense of time, since a causal connection is not available. However, if one can show that we possess psychological mechanisms that establish the appropriate informational link with time, then we will have done part of the work required to

avoid the conclusion of this argument. What remains to be shown is that the information link established is an automatic one in the sense that the relevant feature of the environment is in control of the information gathered by the senses. Importantly, in what follows I will show that this additional constraint is also satisfied, even though time does not exert causal control over the sense of time, it is nevertheless true that the sense of time is nomically related to time, and is in that sense not under the control of the individual organism that possesses the sense.

Against the integration argument

As Matthen argued, it appears as though at least some of our access to the temporal features of mind-external events is indirect. We export the temporal relations in which our experiences stand and attribute them to the events those experiences are of. As a result, the attribution of temporal properties is through a process that operates over the individual senses, and not a genuine sense in and of itself. Whether or not Matthen is correct in his account of how the temporal relations of earlier than, later than, and simultaneous with are attributed to the deliverances of the individual senses can actually be sidestepped here. We can avoid the conclusion to his argument in either of two ways while still granting him all of his premises. (In fact, in the next chapter we'll see that the exportation mechanism doesn't explain how we perceive temporal relations.)

The first way of denying his conclusion is to argue that the mechanism he proposes may not be the only mechanism available for the attribution of temporal ordering relations. The temporal exportation mechanism he describes could simply be one of several mechanisms that attributes temporal relations to the world. The resulting situation would be similar to the situation in which we are asked to detect the texture of an object. Merely pointing to the fact that someone can come to infer the expected feel of an object based on how it looks, doesn't suffice to deny that we have a genuine sensory access to texture

given by touch. We often have multiple routes according to which we might detect properties in the environment, yet this fact doesn't prohibit the possession of a sense for any of these features.

The second approach is to argue that perhaps the mechanism Matthen discusses is the only mechanism for the attribution of temporal ordering relations (although for this argument it needn't be, but we can grant this and still run this line of argument), but there may be mechanisms that attribute other temporal features to the external world that do not employ any sort of temporal exportation. Furthermore, of these other mechanisms there may be a mechanism that can be properly described as a genuine sense of time. Both of these approaches would amount to a denial of the exhaustiveness claim Matthen needs, but it is the second approach that I will follow in the coming section. There exists a mechanism, distinct from the mechanism discussed by Matthen, that has the features of a genuine sense and is used to detect the approximate time of day. It is by noting that we have multiple means of representing time that we open up the possibility for understanding which of these representational abilities are genuinely sensory or not.

3.4. Circadian rhythms and internal clocks

Let's begin with a very general point about animal behavior that we've already mentioned. In order for animals to be able to successfully navigate their environments they need to be able to coordinate their behaviors, at a variety of time scales, with the temporal structure of the events in their environment. Some behaviors need to be coordinated with the temporal properties of events, regardless of when they occur. For instance, a baseball batter has to be able to coordinate their swing with the spatio-temporal trajectory of an incoming pitch regardless of when that pitch is occurring (i.e. regardless of what specific point in time or time of day the pitch is occurring). Some behaviors, however, need to be coordinated with the temporal properties of events that occur at particular times (e.g. foraging behavior often must be timed to specific times of day when food sources are available). In this section we'll look at coordination of the

second sort. Specifically, we'll be looking at relevant literature on the role of the circadian system in the coordination of circadian behaviors.

Throughout the animal kingdom we find that many animals exhibit patterns of behavior and patterns of internal activity that have roughly 24-hour periods.⁸² These behaviors include things like sleep-wake cycles, eating patterns, hormone regulation, body temperature regulation (Moore 1997), and even patterns pertaining to how effectively new memories can be formed (Ruby et al 2008). While the existence of a number of daily patterns with nearly the same 24-hour period may seem to imply that there is some sort of mechanism (perhaps even a single mechanism) that represents the approximate time of day and allows for the regulation of these pattern that inference could fail.

One could deny that animals need to represent the time of day in order to coordinate their behaviors with the temporal structure of the environment. Take for instance the feeding behavior of certain fish. As many fishermen might know, in certain places fish can be found in certain locations at certain times of day. Now, one might explain this predictability of fish behavior in terms of the fish having a robust internal representation of the time of day that they use to navigate to a particular location at a particular time. However, another explanation of the regular pattern in fish behavior offloads the explanation from an internal representation of time to a causal sensitivity to rhythmic patterns in the environment (in chapter 2 I called this the *cued-synchronization model*). One possible example can appeal to the rhythmic pattern of tides in particular locations. In some places, tides have a nearly 24-hour period, and fish might simply be cuing into the fact that at certain tides, food sources will be located in particular locations, and as a result the fish will congregate in particular locations at particular times of day. No representation of time is needed to explain the regular pattern of fish behavior. Instead, the fish are simply

⁸² In fact, circadian rhythms seem to be ubiquitous in not just the animal kingdom but also the plant and bacteria kingdoms. Ultimately, I will argue that many *animals* have a sense of time that is centered on the circadian system, but I would likely stop short of attributing a sense of time to non-animal organisms. In those cases, the mechanisms that underpin the circadian rhythms of activity simply do not have the functional characteristics needed for being a sensory system.

causally sensitive to some aspect of the environment that itself exhibits a 24-hour period, and the circadian behavior of the fish is simply scaffolded off of this external rhythm. So, what seems to be a sophisticated situation in which a rich system of internal temporal representations were required actually can turn out to be a rather simple situation in which no temporal representation (and possibly, no representation whatsoever) is needed.

Furthermore, once we begin to notice that there are patterns in the environment that can be cueing and driving temporally regular behaviors in animals we open the door for a complex iterative process in which further complex behaviors are further scaffolded off of the initial environmental regularities. Suppose, again, that the temporal patterns of fish behavior can be explained by appealing to a causal sensitivity to environmental patterns. If that were the case, then we can likely expect to find further animal behaviors that build off of the fish behavior. Predators that prey upon fish will also exhibit behaviors that are coordinated with the temporal structure of the environment by merely tracking the location of the fish whose movements are themselves temporally structured. Once we notice how patterns can build upon patterns we can actually explain how a complex of temporally structured behaviors might simply be built on top of some environmental rhythm.

In fact, the prevalence of *contextual cues* in the environment led to many researchers denying that animals synchronized their behaviors with the environment by representing time itself. Certainly, this line of reasoning has something to it. In general, we choose scientific theories that posit the least amount of novel machinery. Since we already, for many other reasons, have to posit that animals have means of keeping track of the non-temporal aspects of their environment, if we can also explain their circadian behaviors by merely appealing to these other systems, then we would have a simpler theory than one that posits mechanisms specific to temporal representation. In other words, researchers tried to explain the temporal patterns of animal behaviors by positing mechanisms that keep track of the *what* and the *where*. In some cases, surely something like this happens. As a child, I would know that dinner was going to be

soon when Jeopardy began playing on TV, and more often than not, this would result in me, or someone in the family, setting the table. So, there would be a temporally regular behavior in the household, but we can explain the temporal regularity without appealing to any temporal representation. Rather, the regularity was due to a response to the regularity in the TV broadcast.

However, there are good reasons for thinking that many, if not most, circadian behaviors cannot be explained as a result of some temporal patterns of external cues. The main reason is simply that many of these circadian behaviors, even those that are learned, persist in the absence of any external cues. In this way, the experiments that revealed the temporal mechanisms that underpin mirror those that revealed the mechanisms that underpinned the electroreceptive capacities of elasmobranch fish. First, eliminate the influence of any of the existing sensory system, then, if the behaviors persist there must be some other means by which the animal is keeping track of its environment.

Consider the following sorts of studies reported in Gallistel's *The Organization of Learning* (Gallistel, 1996). In one study (Holloway & Wansley 1973), rats were administered shocks (e.g. negative reinforcement) when they attempted to enter a chamber containing sugared milk at a particular time of day. The very same rats were also given positive reinforcement when they tried to enter the very same chamber but at a different time of day. So, the rats were receiving feedback that could only be interpreted as not conflicting if the rats were able to keep track of *when* they were receiving feedback. That is, if they were not keeping track of when they were receiving feedback, then the feedback they were receiving would be contradictory and the positive and negative reinforcement should cancel each other out (if not fully, then at least there would be a weakening of any reinforcement effect). However, the results were that the positive and negative reinforcement would only compete if they were receiving the feedback at the same time of day (on separate days, of course). The conditioning of the rat behavior was clearly sensitive to time of day.

What wasn't clear from these conditioning experiments was whether the resulting circadian behaviors (i.e. the feeding pattern with approximately 24-hour rhythms) were still the product of some external cue. Further studies by Rosenwasser, Pelchat, and Adler (Rosenwasser et al 1984) showed that rats that were trained to anticipate food in certain locations at particular times of day would continue to show these anticipatory behaviors even when they were put into *free running conditions* in which all contextual cues were removed. The circadian behaviors of the rats simply could not be explained by appealing to any external cues since there simply were no available cues in the rats' environment.

However, Gallistel points out that some theorists attempted to maintain that *some still undiscovered cues* could explain the free running behavior of rats. However, even this sort of *know not what external cue* theory is ultimately untenable once you realize that the circadian rhythms of individual animals kept in free running conditions will begin to drift away from one another. Within a short amount of time, often simply within a day or two, animals that were originally showing synchronized behaviors began to show behaviors that slowly drifted away from synchrony. But, importantly, their behavior was still rhythmic (or had a nearly 24-hour period). Some animals would show circadian behavior that was operating on a contracted schedule (e.g. they would perform some action every 22 hours, as opposed to every 24 hours). Some animals, on the other hand, would show circadian behavior that was operating on an extended schedule (e.g. they would perform some action every 26 hours). Furthermore, once the period of behavior for an animal was determined, then the timing of other of that animal's circadian behaviors could be predicted as falling on that animal specific period (Bolles & Moot 1973). This shifting away from each other could not explained if each animal was responding to the same *unmentionable cue* unless in addition to some unknown cue we posited a specific unknown cue for each animal. However, this is clearly starting to strain any desire for a parsimonious theory.⁸³

⁸³ The astute reader will notice that in free-running conditions the circadian behaviors of animals began to drift, thus indicating some role for external cues in synchronizing behavior. We'll discuss this influence of external cues on circadian behaviors slightly later.

Instead, the explanation that was given, and that is widely accepted today, is that the circadian behaviors of a wide-range of animals are in many cases best explained by the operation of an internal time keeping mechanism with a fixed period. That is, explanations of circadian behavior often appeal to the operation of an internal time keeping device while downplaying the influence of contextual cues. This still, however, left two options open for how to understand the operation of this internal time keeping mechanism. The first option describes the time keeping mechanism as an *interval timing mechanism*, while the second option describes the mechanism as a *period timing mechanism*.

According to the explanation based on an interval timing mechanism, circadian behaviors would be explained in the following way: A given event is used as a marker and a timer is started with the occurrence of that event. Subsequent behavior is then governed by noting the temporal interval between the marker event and some reading of the interval timing mechanism. For example, in time-place-learning experiments, a rat may learn when to find food at a particular location by measuring the temporal interval between some marker event and some feeding event.

According to the alternative period timing mechanism, circadian behaviors would be explained in the following way: No marker event is required. Rather, there is some oscillator system with a regular period and particular times of day are recorded as corresponding to particular states of the oscillator. Taking the time-place-learning example in hand again, the feeding behavior of the rat could be explained by directly appealing to the rat encoding the particular time of day at which food is found in particular locations in terms of states of the oscillator. In a sense, this means of keeping track of the timing of events involves a time-stamping process.

Importantly, the differences between an interval timer and a period timer are noticeable both in terms of the implementation of the timers and in the content of the timer. In terms of implementation interval times typically involve some sort of accumulation process. Considering the classic pacemaker-accumulator mechanism of Treisman (1963), the pacemaker will produce neural spikes at a regular rate

and the accumulator system keeps track of these spikes in format that is neurally significant⁸⁴ for downstream processes. On the other hand, period timers needn't involve any accumulation process. Rather, they simply appeal to the rhythmic change of some neurally significant properties of a system that can be exploited to keep track of time. For instance, we cannot describe the accumulation or the rhythmic states of a time keeping mechanism as being *purely temporal properties*. The accumulation cannot merely be that the system has been turned on for 5 *seconds*. Similarly, the rhythmic properties of an oscillator cannot be *that the clock is now at 5pm*. Pure temporal properties simply are not the sorts of things that neural systems can exploit to guide their behavior, and any temporal information encoded by a timing mechanism must be such that it can be exploited by downstream systems.⁸⁵ Otherwise, the positing of such mechanisms would do nothing for explaining animal behavior. Crucially, whatever properties of the timing mechanisms we appeal to as being the semantically significant properties must also be causally efficacious within the overall cognitive architecture.

Also, the two sorts of timers differ in their semantic contents. Interval timers invariably appeal to a marker event – their content is always of the form $\langle e \text{ is } I \text{ since } m \rangle$, where e is the target event, m is the marker event, and I is the temporal interval between e and m . Period timers, on the other hand, do not reference any marker state. Instead, period timers predicate of particular moments some temporal property – their content is always of the form $\langle t_p \text{ is } p \rangle$, where t_p is a particular moment in time and p predicates of t_p a certain temporal property (we'll say more about this later). If an interval timer model is to account for circadian behavior, then we need to give an account of what the relevant target and marker

⁸⁴ Recall, from chapter 2 it was argued that representational systems are constrained by the downstream accessibility of information. In the case of neural representations, the information encoded by a system must be stored in a manner that downstream neural systems can access that information. That is, the information must be encoded in a *neurally significant* format.

⁸⁵ Surely, there are properties in the vicinity of these pure temporal properties *that are* neurally significant. Suppose that a neural system in a particular state, say *state A*, results in the system producing 50 spikes per second. As a result, being in neural state *A* for 2 second, will be causally efficacious since the system will have produced a total of 100 spikes over that interval.

events might be for any circadian task. If no such account can be given, then we ought to abandon the interval timing model for circadian behavior.

With an understanding of the typical implementation of these timing mechanisms and the differences in their semantic contents we notice that we have ways of adjudicating between these different models of circadian behavior.

Several considerations count against adopting an interval timing model at both the neuroscientific level and the behavioral / functional levels. First, the interval approach still requires some initial marker event according to which subsequent behaviours can be synchronized. In the studies discussed, all attempts were made to remove any contextual cues that could be used in order to regulate circadian behaviour, as a result, it seems as though there would be no marker events according to which one might initiate the interval timing mechanism. However, its possible that the animal's own behaviours could serve as marker events in order to synchronize circadian behaviours. Perhaps, the animal knows that 4 hours after it wakes it should look for food in location A and that 6 hours after looking for food in location A it should look for food in location B. And so on for the rest of the animal's behaviour. However, there is a feature of the interval timing mechanisms that is not found in period timing mechanisms. As the length of the interval the interval timer is measuring increases, the timer will increase in noise and become less accurate (Dehaene 2003; Gallistel & Gelman 1992; Malapani & Fairhurst 2002; Wearden 2001). In some cases, this sort of noise in the model is appreciated as we find a corresponding noisiness in timing behaviors. In particular, many interval timing studies have shown that animals throughout the animal kingdom discriminate durations in accordance with Weber's Law, according to which the just noticeable difference between two stimuli is a constant ratio of the values of the two intensities (e.g. as the overall intensity or duration of a stimulus increases individuals will require a larger difference for another stimulus to be judged as different). As a result, we should find that a degree of inaccuracy in the circadian behaviors that simply isn't found in the above studies (Ko et al

2003). The circadian behaviors are timed with equal accuracy regardless of *when in the day* they are occurring. So, on the behavioral / functional level, we seem to not find the signatures associated with interval timing mechanisms.

While the behavioral / functional data were telling against the interval timing model, the final nail in the coffin came from the direction of implementation. What really changed the way in which researchers thought about the circadian system was the neuroscientific findings concerning the neural systems that regulate the circadian behaviours. The first task in understanding the circadian system was the task of localization. It was discovered through a series of studies that many circadian behaviours in mammals are regulated by a specific neural region called the suprachiasmatic nucleus (SCN), a region of the hypothalamus, alongside a variety of more peripheral oscillatory systems whose rhythms are synchronized through the SCN. It was shown in a number of studies that damage to the SCN results in disruptions to the circadian rhythms of animals.

For instance, one telling study involved the long-term monitoring of two chipmunk populations (DeCoursey et al 2000). The first population underwent surgery in which their SCN's were removed. The second, control, population underwent control operations in which the animal underwent much the same operation as the chipmunks in the first population underwent, however, they were left with intact SCNs. The chipmunks were then released back into the wild and monitored over the course of 18 months. After the end of the 18-month monitoring period, the control group had a significantly higher survival rate than the experimental group that had lost their functioning SCNs. It was determined that one contributing factor in the lowered survival rate of the experimental group was that the chipmunks with excised SCNs were much more active at night when predatory weasels were out on the hunt. As a result, the SCN excised chipmunks were noticed at a much higher rate by the predatory weasels and for that reason they were victims of weasel predation at a much higher rate than the control group of chipmunks. Due to the

loss of the SCN circadian behaviours were not properly synchronized with the environment, and as a result, the chipmunks in the experimental group suffered from a greatly reduced survival rate.

Another telling experiment involved the behaviour of hamsters kept in free-running conditions (Ralph et al 1990). Through selective breeding practices two populations of hamster were bred with distinct circadian cycles. One group of hamsters possessed the normal circadian system with a nearly 24-hour period, while the second, mutant, hamster group possessed a circadian system with a greatly reduced period of approximately 16 hours. Upon the ablation of the SCN in hamsters of either group circadian rhythms halted – sleep/wake cycles, daily activities, body temperature regulation, and hormone regulation all became erratic. However, circadian behaviours were recovered upon surgical transplantation of healthy SCN tissue. Perhaps most interesting however is that when mutant SCN tissue was transplanted in a hamster from the normal population circadian rhythmicity was restored but the circadian patterns of the normal hamster displayed the 16 hour period of the mutant SCN. Similar findings were found when the mutant hamsters were given normal SCN transplants.⁸⁶ The most natural explanation of these behaviours is that it is the SCN that plays a crucial role in determining and regulating circadian behaviours.

One final example is useful to understand the role of the SCN. It is known that mammalian body temperatures fluctuate over a circadian cycle. In a study by (Megumi Maruyama 2007) it was shown that the SCN facilitates the formation of heat stress memories. Two groups of rats were exposed to elevated temperatures for a fixed five-hour period of the day. One group of rats had bilateral damage to their SCN whereas the control group had an intact SCN. Both groups exhibited lowered core temperatures for an extended time of the day, as a means of combatting the heat stress. However, only the control group of rats with intact SCNs exhibited lowered core temperatures that coincided with the time of day during

⁸⁶ A fascinating aspect of these studies was that upon receiving a transplant SCN, animals would show circadian behaviors but upon anatomical examination it was found that the neural connections between the SCN and the rest of the brain were not properly formed. The influence that the SCN had on behavior was seemingly driven through *humoral factors* released into the bloodstream. This seems to be a manner in which a neural system may convey information to other regions of the brain in a way that bypasses direct neural connections.

which the training heat stress stimuli were administered. The rats with SCN damage exhibited lowered core temperature but this lowered core temperature was not specific to any time of day. The SCN damaged rats were registering the heat stress but were unable to synchronize their body temperature with when the heat stress occurred.

Now, merely localizing the SCN as the mechanism that regulates circadian behaviour doesn't adjudicate between the interval timing explanation and the period timing explanation. However, it is by looking at the manner in which the SCN functions that we differentiate between the two models. Through a series of experiments the mechanism by which circadian timing occurs was discovered. While the details are still being worked out (see Bechtel (2011b)), the general story is now well known, and a simplified version of the model can be given here. The general mechanics of the SCN is that of a molecular clock that is governed by the reciprocal interaction of the transcription and translation of multiple genes and proteins, but for our purposes we can simplify the issue and consider a system that consists of only a single gene / protein pair⁸⁷ Consider the gene PER (short for period). During gene expression, PER is transcribed to produce the protein *per*, however, as the level of *per* in the cell increases, the expression of PER is inhibited, thereby reducing the production rate of *per*. As the protein levels decrease, the gene PER is once again readily expressed. It was then discovered that through some at present unknown mechanism, the level of *per* in the cell directly related to the spontaneous firing rate of the SCN neurons⁸⁸. As a result, no accumulation process was found in the functioning of the SCN, instead, what was found was a cyclical behaviour in the SCN neurons. Furthermore, the SCN's behaviour is controlled by the endogenous behaviour of the SCN neurons, and not through some causal interaction

⁸⁷ Most current models of SCN functioning in mammals involve the reciprocal interaction of five or more gene / protein pairs (Bechtel, 2011b), however, the simplifications made in this chapter do not mischaracterize the general point.

⁸⁸ While the exact mechanism by which *per* levels influence spontaneous firing rates is unknown, see (Vasalou & Henson 2010) for an account.

or registering of an external state, and as a result, there is no marker state according to which the functioning of the SCN depends.

The spontaneous firing rate of the SCN is then used as a master regulator that regulates a variety of other circadian systems in the body. For instance, as the firing rate of the SCN increases, it will either activate or inhibit downstream neural populations that themselves will directly influence various conditions in the body that are behaviourally relevant. In some cases these systems will involve the production of hormones during certain times of day, such as the production of melatonin at night (Moore 1997). In these cases, the behaviour of the SCN is used to provide coarse synchronizations with the environment – sleep related hormones are released at night and this becomes behaviourally important because if the animal were not asleep at night, as in the chipmunk case described above, then the animal will incur a reduction in its ability to survive. In other cases, these downstream processes will involve the coordination of behaviour to much more specific times. For instance, it is known that certain birds navigate by calculating target paths by integrating the location of the sun and the current time of day (Cassone et al 2009; Cassone & Westneat 2012). This sort of navigation requires a much more specific representation of the time of day than merely day / night representations, furthermore, merely appealing to the location of the sun (i.e. contextual cues) is insufficient for the determination of a flight path.

It should be noted that since the circadian system, like any clock, acts as an oscillator, going through a complete cycle approximately every 24 hours, that the system occasionally needs to be recalibrated or entrained. This is particularly important when as the seasons change and the day / night ratio changes. This recalibration mainly occurs through the input from a dedicated visual pathway that indicates the overall light level of the environment. Unlike most visual processing, which begins with input from the cones and rods, circadian entrainment begins with photosensitive retinal ganglion cells. These cells input directly to the circadian system and do not enter into the pathways that lead to the visual cortex. Importantly, as is witnessed by animal behavior in free-running conditions, the light entrainment

from retinal ganglion cells only serves to recalibrate the system, as the circadian system continues to function even in the absence of light / dark information. Much like the way that a clock that is recalibrated on occasion still tracks time, the circadian system, even though it is recalibrated by light / dark information, still tracks time of day, and not the light / dark cycle.

The need for this occasional recalibration is what explains why in normal situations we might find a whole group of animals to have synchronized behaviors but when those animals are put into free-running conditions the timing of their behaviors begins to drift apart. This fact can often be used as a way of raising objections to my discussion so far.⁸⁹ The objection states that perhaps what this system is representing is not the time of day, but rather, the location of the sun.

Let's take this objection that the circadian system isn't representing *time* per say, but is instead measuring the location of the sun and see how it fails. As we'll see in chapter 5, there's actually something very plausible to this suggestion yet that doesn't undermine the claim that the circadian system represents the time of day. Many of our own cultural time keeping mechanisms, the analog clocks you see on the wall, for instance, were for many years calibrated to *local noon*. Clocks were synchronized to read 12:00 pm at the point when the sun was highest in the sky. So, in this way, we *could* have interpreted these clocks as representing the location of the sun, and not the time of day. However, this isn't how we *made use* of the clocks. Given that the time of day and the location of the sun would often go hand in hand (less so, of course as we move away from the equator), the movement of the hands of the clock could carry information about both the time of day and the location of the sun. However, the information that we would access from these devices would be used to coordinate the timing of our actions.⁹⁰

⁸⁹ These objections have been raised by audiences at the 2016 Eastern Meeting of the American Philosophical Association and the 2014 Society for Philosophy and Psychology Meeting. Also, thanks to my commentator at the Eastern APA, David Sackris, for raising this objection in detail.

⁹⁰ Yes, we could also coordinate our time sensitive actions by following rules like, "when the sun is at X location, do action Y". In this way, however, we would be using the location of the sun in much the same way that we use a clock. The trajectory of the sun itself would be the periodic system whose states we use to tell us the time. Regardless, the system would still be representing the time of day.

Furthermore, even in the absence of sunlight, i.e. when the sun was simply gone, we could still use the clocks to coordinate our behaviors. Similarly, for animals. The circadian system can still be used, for instance in time-place-learning, when the sun is gone. Finally, while input from the retina is perhaps the strongest source of calibration for the SCN, it is not the only source. One significant source, and one that many people know from merely battling jetlag, is food (Schibler et al 2003; Stephan 2002). If food can also calibrate the system, and have similar results as the calibration through light levels, then the system really cannot be interpreted as representing the location of the sun. The very same considerations would force us to attribute a contrary interpretation in which the SCN is tracking the presence of food, but this conflicts with the SCN tracking the location of the sun. Rather, the more general interpretation, that the SCN is tracking the approximate time of day, makes all the other circadian behaviors intelligible. And as a result should be our preferred interpretation.

Someone might object however that on purely empirical grounds, we shouldn't take the SCN to be the system that regulates all of these distinct circadian behaviours. This is due to some findings that certain circadian behaviours persist in animals that have had their SCN damaged through selective surgical procedures (P. A. Lewis et al., 2003). For instance, it has been shown that hamsters show conditioned place preferences to specific times of day even though they have had bilateral SCN damage (Marchant & Mistlberger 1997). However, in these cases the question never shifts away from understanding the internal oscillator unpinning circadian behaviors. Instead, the only shift in the question is a shift to what other oscillator systems might be in place, since the behavioral / functional markers of an oscillator system are still present – there are no available contextual cues and there are no signs of increased inaccuracy as length of time increases (ruling out a contextual account and an interval timing model respectively). Further data shows that the SCN is still involved to some degree in the timing of these other oscillators since a damaged SCN that is still capable of sending signals, although misleading signals, to downstream processes seems to impair the functioning of these other oscillator systems. The resulting view that seems to be coming out of the chronobiology literature is the SCN serves as a master

clock that is capable of synchronizing a variety of other mechanisms in the body and brain that display circadian behaviors. The important point for this paper is that the way in which circadian behaviors are implemented is through oscillator systems that directly and automatically keep track of time. Whether those oscillators are wholly located in the SCN or not is only of secondary concern.

What we have here is an argument against the argument from integration. Recall that that argument had an implicit premise that the mechanism described there that takes the temporal relation between experiences and attributes that temporal relation to the objects that those experiences are of was the only mechanism that we possessed to track the temporal properties of the world. However, as has been shown so far in this discussion, the mechanisms that subserve the circadian behaviours are not like that at all. Instead, they function by representing the time of day through the functioning of an internal oscillator(s). Nothing like temporal exportation happens here. Instead, we seem to have system that tracks the time of day without extracting the time of day information from any other sense. So while it may be true that the way in which we consciously attribute temporal properties to many objects and events at short time scales is through something like what Matthen (2014) and Phillips (2010) discuss, the point remains that there is an additional system that keeps track of longer periods of time involved in circadian behaviour (although, see chapter 4 for an argument against the exportation mechanism).

3.5. The semantics of (internal) clocks

In order to use the empirical literature on the SCN to show that the non-causality argument fails, we need to first say something about the semantics of clocks. Once we have a general story about the semantics of clocks on the table, then we'll be in a position to give the meta-semantic account of how that content or information is acquired in a direct and automatic manner that does not depend on a causal relationship between time itself and the states of the clock.

Let's begin with a general point. In order to get clear on the semantic content of any representation it's important to look at the correctness conditions for that representation. For instance, the sentence "Matthew is a student" is correctly asserted, or is true, just in case Matthew, the referent of 'Matthew', has the property of *being a student*, the semantic value of 'is a student'. The sentence is false, just in case, Matthew, or whatever object is referred to by the sentence, fails to have the property of being a student.⁹¹ Individual names or predicates do not have correctness conditions in and of themselves.⁹² To put the idea in Fregean terms, a predicate or name by itself does not express a complete thought, or proposition, that can be evaluated for truth. In order for a sentence, or a representation, to have correctness conditions it must be evaluable against the predication of a property to an object.

Given that we take clocks to be correct or incorrect it follows that clocks do more than merely refer to an object or pick out a property. Clocks must predicate of particular objects some sort of property, but what object clocks refer to and what property they attribute to those objects is less than obvious. In order to clarify this matter, we can look at *the way clocks are used* as a guide to their content, since whatever content we attribute to clocks must be in line with the way that we make semantic, or informational, use of clocks. Otherwise, if there weren't some connection between the uses of clocks and their contents, then our uses of clocks would remain mysterious.

Let's start with the following claim, called CLOCK, in order to understand the semantics of clocks:

CLOCK: It is 5pm.

⁹¹ The sense in which whatever referred to by the sentence is "an object" takes object in an incredibly broad sense. Anything that can be referred to or predicated of counts as an object. We shouldn't restrict this semantic structure from applying to things that do not live up to the standards of genuinely material objects.

⁹² We can of course talk about correct or incorrect applications of individual names or predicates. For instance, I might incorrectly call some philosopher 'Bruce' in which case I would have incorrectly used the name. But in these cases the correctness of the application of a name, or predicate, is really short hand for something like *that is Bruce*. The correctness is still given at the level of full subject predicate wholes.

For CLOCK to be true, the following must be the case: the referent of ‘it’ must stand in the relation expressed by ‘is’ to whatever is expressed (or referred to) by ‘5pm’. Given what we’ve said so far in this section, this much should be obvious.

Let’s begin with understanding ‘it’. What does ‘it’ refer to in CLOCK? Something important about clocks is their indexical nature. A properly functioning clock does not tell you something about some arbitrary time. Clocks tell you something about what time it is *now*. To clarify this point, consider the following scenario involving Bonnie. Bonnie is a part in a complicated bank heist and she has explicit instructions to cut the power to the bank’s security system at 5pm sharp. Since the timing of this task is of the utmost importance, she brings with her two clocks as a redundancy just in case one of the clocks stops working. As she is sitting there in front of the electrical wires that feed power to the bank’s security system she goes to consult her clocks but accidentally drops both of them. When she picks them up and looks at them, she realizes that one clock reads 4:45pm while the other clock reads 4:48pm. Both clocks can’t be accurate, but how are we supposed to characterize this difference? What makes one inaccurate?

The intuitive answer, and the answer that is crucial for how Bonnie is using the clocks, is that one of the clocks is inaccurate since it is misrepresenting what time it is now.⁹³ While it’s a matter of debate in the metaphysics of time as to what particular moments in time amount to and whether there is anything special about the current moment in time, however, for our purposes here we can remain neutral towards these metaphysical debates. Whatever metaphysics of time we adopt, there will be something that we pick out with the phrase ‘this moment in time’, and it is to that sort of entity that clocks attribute particular temporal properties.⁹⁴

⁹³ There are clear parallels between the semantics of clocks and the role of indexicals that John Perry isolates in his paper *The problem of the essential indexical* (1979). In both cases, the indexical representations are crucial for guiding behavior.

⁹⁴ Whether particular moments in time are in some sense primitive or whether they are derived from other aspects of the world is a particularly pressing issue in the philosophy of time. For discussion see Meyer (2013).

To further make the case for the indexical nature of clocks, we can see that clocks possess a property that has often been restricted to demonstratives in the philosophy of language (Recanati 2012). Clocks are immune to error through misidentification. A sentence like ‘John is running’ can fail to be true on a particular use in two ways. Either John fails to be running or the sentence, as it is applied to some situation in the world, does not actually pick out John. I may misapply the sentence to someone who looks like John, but isn’t in fact John. However, some sentences are unable to fail to be true due to this second sort of situation. These are sentences that are thought to be immune to error through misidentification. Consider a demonstrative statement “that is running” that I utter without any presuppositions about the thing that I am referring to. I intend to demonstratively refer to whatever it is that I’ve selected. This particular demonstrative utterance may be wrong because whatever I demonstratively pick out fails to be running. Or I may be under an illusion and fail to pick anything out at all. However, this statement cannot be wrong as a result of not picking out the right object (in the same way as when I mistook someone as John). In this way, clocks also are immune through error through misidentification, since they cannot be describing the wrong time, since they always pick out now. However, unlike demonstratives, clock representations cannot even fail to pick out an entity. Whenever a clock is representing a time it is invariably referring successfully, since the clock’s operations will necessarily be occurring in time.

Having pinned down what ‘it’ refers to in CLOCK, we must now turn to understanding the overall logical structure of CLOCK and how this logical structure depends on how we interpret ‘is’. ‘Is’ famously has two distinct interpretations (Russell 1905). There is the ‘is’ of predication, as in ‘2 is prime’ in which we attribute a property to some object. There is also the ‘is’ of identity, as in ‘2 is 2’, in which we say that one object stands in the relation of being identical to an object (namely, itself). In cases in which the is of identity is being used, we can typically add the phrase ‘identical to’ after the ‘is’ in the sentence without any problem (e.g. ‘2 is 2’ can be read as ‘2 is identical to 2’). Correspondingly, there are two possible interpretations of CLOCK that depend on differing uses of ‘is’.

I think there are clear reasons for rejecting the identity reading of CLOCK. If we read CLOCK as expressing an identity statement, then we are making the claim that now = 5pm. Yet to understand the truth of this statement, we need to look at the referent of the indexical, and we get a particular instant in time, call that instant i_1 , so the claim states that $i_1 = 5\text{pm}$. However, suppose that a day goes by, and again we utter CLOCK, in that context of utterance we would understand clock to be saying of a particular instance of time, i_2 , that $i_2 = 5\text{pm}$. Yet, transitivity of identity fails in these cases since $i_1 \neq i_2$, since these are times on different days, and therefore are different times, yet we take each utterance of CLOCK to be true. As a result, we can't coherently hold the identificational reading of CLOCK, instead the logical structure of CLOCK is one of predication. As a result, we need to find some property that is being picked out by '5pm'.

This should strike most people as odd. What property is picked out by 'is 5pm'? It seems like 'is 5pm' shouldn't pick out any objective property of the world since our timing practices seem to be the product of merely contingent convention and as a result the division of any length of time into something like hours, minutes, and seconds, is just some sort of societal construction. If there really isn't any property out in the world that is expressed by 'is 5pm', then it would be impossible to understand clocks as being accurate or not (given that accuracy is typically understood in terms of something like correspondence with reality), yet we do in fact take clocks to have satisfiable accuracy conditions, and therefore there must be some property out in the world that is expressed by 'is 5pm'. In order to make sense of all of this we need to notice the way in which our timing practices provide a framework according to which clocks can be understood as making accurate or inaccurate claims about the objective world. The divisions (e.g. the temporal properties) that are picked out by our timing practices are just those that are imposed on temporal reality by our timing practices themselves.

Let us consider a spatial analogy, and then we will make connections between the spatial case and the temporal case. Consider a long road along which we are moving at a fixed rate and along this road

there are rest stops at regular intervals. We might want to know what the distance is between any location on the road and the next rest stop. Imagine that in order to determine how far between two rest stops we are we decide to extend a semi-transparent overlay with a regular 24 color pattern over the length of the road in which the entire 24 color spectrum is repeated through every interval between rest stops. We can then describe our location between any two rest stops in the following way “here is blue” or “here is orange”. When we say 'here is orange' we are saying that our current location has the property of being in a particular position in a sequence of locations. The predicate 'is orange' has as its extension all of those objective spatial locations that fall under the orange overlay. Accordingly, these descriptions provide us with a characterization of our objective location relative to any two rest stops, even though the subdivision of the road was imposed by something of our choosing. Once that convention is in place, we can then speak accurately or inaccurately about our location along the road. Notice importantly, that in employing this colored overlay on the road as a means of dividing the length of the road, we do not end up adopting some anti-realism about space. The spatial locations are all independent of our way of carving up the intervals between rest stops. The only thing that we are imposing is a means of classifying these locations such that we have a means of describing the objective spatial locations in terms of the class of locations they fall under given our means of carving up the road.

We can do the same with the time of day. Consider each passing day, or sunrise / sunset cycle, as being analogous to the regularly spaced rest stops along the road. We then simply establish some representational system that allows for that interval to be subdivided into discrete units. The degree to which we subdivide this interval is simply a product of the specificity we need in order to characterize our time keeping practices. But there is no circularity here in trying to account for our time keeping abilities in terms of arbitrary subdivision that are then explained by our time keeping practices. The time keeping practices go hand in hand with how we subdivide the temporal interval. The two things come together since you cannot have one without the other.

When it comes to how we analyse the semantics of 24-hour clocks, we essentially lay out a “transparent” overlay along the temporal dimension dividing up the objective stretch of time into hours, minutes, seconds, etc. We then use this subdivision of the world to predicate of the present moment. When it comes to internal clocks, however, someone might object that as theorists we are anthropomorphizing circadian behaviour in that we describe the content of the circadian system in hours, minutes and seconds. In general there is a risk in the cognitive sciences that as conceptually sophisticated adult humans we may be systematically mischaracterizing the informational content of not only sub-personal states but of non-adult human animals in general. There is some reason to worry about this concern, in general, however, the concern is misplaced here.

We may reformulate the content of the circadian system using our language as theorists without mischaracterizing the content of the circadian system. Our theoretical language may involve different predicates as the circadian system, yet in attempting to describe the circadian system we won't cause any problems unless we start attributing to the system a greater amount of sophistication than is present. For instance, perhaps we do not want to attribute to the circadian system the same level of precision in time keeping as we would to a watch with a second hand. However, we can constrain the content we attribute by limiting our attribution of content to the minimum precision required to explain the circadian behaviour and the functioning of the circadian mechanism.

3.6. An information-theoretic account

As mentioned earlier, a system s carries information about some state of the world, t 's *being* F (i.e. some object being some way), provided that t 's *being* F is guaranteed by s 's *being in state* R and as a matter of fact s *is in state* R . Since we have an understanding of the semantics of clocks, we need a theory of content that is capable of explaining how it is that the neurological mechanisms of the circadian system come to bear their semantic content. While theories of content aren't often explicitly discussed by

cognitive scientists, there is a common approach in the empirical sciences that lends itself (at least initially) to a *resemblance theory of content* (see chapter 4 for a more detailed discussion of resemblance theories). In particular, many researchers in trying to explain the temporal content of some psychological state will point to the temporal properties of the neurological state.⁹⁵ For instance, in order to explain the temporal content of a perceptual state, say *hearing a loud tone as preceding a quiet tone*, theorists might simply point to the fact that the neural state representing the loud tone temporally precedes the neural state that represents the quiet tone.

However plausible some sort of resemblance theory might be for the representation of duration or temporal ordering the situation is rather different when it comes to representations of the time of day (although, see the discussion in Chapter 4 for why even in those cases resemblance is implausible). Resemblance theories are inheritance theories. A representation will represent some property X only if the representation itself possesses that property X (think for example how a paint swatch at the hardware store will represent the color of paint in a paint can by itself being that color). But what would the inheritance story be for the circadian system? As discussed earlier, the content of the circadian oscillator will be something along the lines of *<it is now t_p >* where ' t_p ' picks out a temporal property that applies to a particular time of day. So, if an inheritance theory is correct, then the representation will have this content only if the representational vehicle itself possesses the property t_p . But notice, that this will trivially be the case, as the representation itself is located in time, it will necessarily have the property of being at the particular moment in time that it is supposed to be misrepresenting. In this way, someone could argue that every aspect of the world carries this content as every current state of affairs in existence will exist currently.

⁹⁵ For a critical discussion of this approach in the cognitive sciences see Dennett & Kinsbourne (1992) and more recently Grush (2006). For a discussion of this approach as it appears in the philosophical literature see Phillips (Phillips, 2014a, 2014b) and Lee (2014a, 2014b) as well as the extended discussion of the topic in chapter 4.

Several problems arise when trying to use an inheritance theory for the circadian clock. First, consider cases of misrepresentation (i.e. where the clock fails to accurately keep track of time and produces maladaptive behaviors). Notice that a paint swatch might misrepresent the color of paint in a can by there being a misalignment between the actual paint in the can and the color represented by the paint swatch. The reason is that there is nothing that by necessity guarantees the color of the paint swatch (i.e. the property of the representational vehicle) will match the color of the paint in the can. However, in the case of the time of day there is a necessary connection. The time of day applies to everything that's occurring at that moment (and in a particular location). So, there can't be the misalignment between the content of the representation and what is supposed to be represented. Therefore, no possibility of misrepresentation. Second, it's unclear how the time of day at which a neural state occurs could itself be neurally significant. If neural systems could just access these time of day properties, then there would be no need for a dedicated clock. But as is evidenced by the behavioral effects of damaging the SCN, there is a need for a dedicated clock, so it's unlikely that the neural system can actually access time of day properties directly.

A distinct theory of content is needed, yet there is something importantly right about the resemblance theorist's attempt to explain the content of the circadian system. An important aspect of how clocks come to have their contents is that clocks not only represent time, but to use a term from Brian Smith (1988), clocks participate in time. Vital to a clock's ability to keep track of time is that the mechanics of clocks are themselves governed by processes that occur in time, and that the occurrence of these processes in time is regular. A grandfather clock is able to keep track of time because of the regular way in which the movement of a pendulum progresses through time. Similarly, it is the progression of the mechanics of the circadian system that allow the circadian clock to keep track of time. Contrast both of these cases with an analog clock whose hands are stuck at 5:30. This clock is never correct – not even twice per day – it simply fails to represent the time. The participatory aspect of clocks is vital to their keeping track of time, yet, the fact that clocks represent time is not due to any sort of resemblance

between the clock and the temporal features of the environment as shown by cases of misrepresentation.

The importance of the participatory nature of clocks is that it is their participation in time that allows them to carry information about time.

Perhaps one could employ a causal theory of content (Fodor 1987; Prinz 2004; Stampe 1977), however, given that it is a controversial idea in the metaphysics of time to think that particular moments in time are causally efficacious it seems that a causal theory of content can be ruled out a priori. I propose to understand the content of the circadian system in information-theoretic terms. Recall the earlier characterization:

A system s in state R carries the information that t is F (where t is an object instantiating F), just in case, t 's being F is guaranteed by S 's being in state r (given appropriate background channel conditions and the laws of nature).

All that is required for an information relation to hold is that there is a relation between pairs of property instantiations. In this case, the relation holds between states of the circadian system and states of the world, as described in the previous section. When the circadian system is in state S_1 , the organism is in a position to know that the current time of day is T_1 , because there is a nomic relation that holds that whenever the time of day is T_1 , the circadian system will be in state S_1 . As time progresses from T_1 to T_2 , the circadian system will itself progress from state S_1 to state S_2 , thereby putting the organism in the position to know that the time of day is T_2 . The endogenous activity of the circadian system guarantees the appropriate nomic relation between the states of the circadian system and the time of day, in the very same way in that the physical behavior of a pendulum clock guarantees that when the clock's hand moves, the movement of the hand indicates the current time.

What we have so far is an account of how the circadian system successfully represents time, but the true test is whether the circadian system is capable of misrepresentation. This brings up a general problem for information theoretic accounts of content. Pure information theories seem unable to explain misrepresentation since if the state of the world that typically corresponds to a state of the representational system fails to occur when the representational system is in the particular corresponding state, then the representational system will simply fail to carry any information. On a strict account of information, there is no such thing as misinformation (Dretske 1981). Either the system carries accurate information or it doesn't carry any information at all. In these cases we can augment a pure information theoretic account by distinguishing between indication and representation (Dretske 1991). A system indicates that some state of affairs is the case, just in case the system carries information about that state of affairs. In other words, indication is a factive relation that holds between an information bearing system and that which it carries information about. However, genuine representation occurs when a state that typically carries information acquires the function of carrying that information and can thereby be said to represent what it typically indicates. A system will acquire the function of representing something when the system typically carries information about some state of the world, and then the representation is in some sense selected for and maintained in order to signal the state of the world. A clock has the function of telling the time, because typically, when everything is functioning properly, the state of the clock carries information about the time *and we use the clock on the assumption that it is successfully carrying this information*. Therefore, when the background channel conditions that underpin the information channel are damaged, and the system misfires and fails to genuinely carry information, we can still treat the system as though it were successfully carrying that information.⁹⁶ Similarly, the circadian clock has the function of representing the time of day, because it has the ability to carry this information, and the adaptive value of the circadian clock resulting in the circadian clock being selected for as a means of

⁹⁶ Misrepresentation, then given this analysis, can be understood in terms of a failure of channel conditions. For instance, the speedometer on your car may misrepresent the speed at which you are travelling due to some damage in the channel that mediates the location of the needle on your dashboard and the movement of your wheels.

serving this function, *and downstream systems operate on the assumption that the circadian clock is successfully carrying this information.* As a result, even when the clock fails to carry information about the environment, the clock is still used by the organism as though it were carrying information. This is all because the clock's internal states typically signal the current time of day, and in cases of misrepresentation, the internal state of the clock is understood as signalling the time of day even though the clock fails to carry any information.

What we have now is an account of how the circadian system comes to have its particular contents. In addition, we have the answer to the original worry posed by the non-causality argument. Recall, the non-causality argument against the sense of time made the following claims: the senses perform their information gathering function in virtue of a causal relation between the senses and that which they are senses of. Time cannot be a cause. Therefore, there cannot be a sense of time. However, we are now in a position to see how this argument fails. The circadian system is a system that is capable of directly establishing a real-time information link with a particular aspect of the environment. But there is something interesting to notice here. The reason why there can be a sense of time is that the circadian system participates in the very thing that it represents. As time progresses in the environment, the internal operations of the circadian system themselves develop and unfold in time. It is this sort of parallelism that guarantees the real-time information link. This is something that is plausibly unique to time and temporal representation. If we look at any of the other senses, there is no similar parallelism. Time, in this sense, is a unique aspect of our environment. However, this strangeness does not create any roadblocks in our understanding of how animals come to represent time. Instead, it is this very feature of time that makes it possible for there to be a sense of time.

3.7. Conclusion

As we typically understand the mind, in both scientific and folk contexts, the senses play a crucial role in explaining our epistemic access to the world around us. They are the first step in how we come to acquire information about the contingent goings on of our environment. However, we find that there is quite a bit of disagreement over how to even distinguish the sensory systems from the rest of the cognitive architecture and also how to distinguish one sensory system from another. In many ways, the Aristotelean senses are often treated by some philosophers as being special in that only they deserve to be categorized as genuine sensory systems and that scientific appeals to non-Aristotelean senses are in some sense *changing the subject*. But, as we've seen, there is in fact a common core to the notion of sensory system that extends through both the folk and scientific uses of the term. The senses are the information uptake systems. Disagreement arises not over this common core idea but over the place to draw the line between cognition and sensory systems and how to individuate one sensory system from another. I have argued, however, that there simply is no single place where to draw these divisions. For various explanatory needs we may need to distinguish the senses from the rest of the cognitive architecture at a certain place as opposed to another place. Similarly, for various explanatory needs we may need to individuate the senses according to different criteria. These two questions are related, and how we determine one will constrain how we answer the other, but importantly both questions are amenable to fluid answers.

Once we have this understanding of the sensory systems in hand, we can then turn to the question of whether there is a sense of time that parallels our other sensory systems. We found that many researchers have tried to deny that there is in fact a sense of time, but argument for this position are far less commonly found. But we were able to isolate two arguments that are explicitly found in the literature. Both arguments focus on the point that there is no system for temporal representation that picks up temporal information in the direct manner that is typical of the sensory systems. In my argument, I

have shown that this simply isn't the case. The systems that coordinate animal activity with the approximate time of day, the circadian systems, pick up information about the temporal structure of the world in a way that cannot be explained by appealing to the informational capacities of the other sensory systems.

In arguing for this claim we actually came across a rather interesting point that has to do with temporal representation. While the other sensory systems require a causal connection between the aspects of the environment they track and the states of the sensory system itself, the circadian system doesn't require any sort of causal connection of this sort to maintain its information gathering role. It is because the circadian system's operations occur in time at a regular rate that they are able to keep track of and represent the approximate time of day in a format that is accessible to downstream neural processes. So, unlike Dretske's example of two TV's that can carry information about one another without having any causal influence on one another, the case of does not even require that there be a continuous causal path (i.e. TV1 to broadcaster to TV2) between the information carrying system and that aspect of the world that it carries information about. Clocks, as Brian Smith pointed out, not only represent time but they also *participate* in time. As a result, time may stand alone as the only sensory system that does not require a causal connection to its content.

Ultimately, however, the goal of this chapter is to point out one way in which animals come to coordinate their behavior with their environments. In this case, the coordination occurs at the relatively long time scale of daily patterns and is subserved by a dedicated system that gathers information about the approximate time of day that can then be used to coordinate a number of behavior. Importantly, though, this system only accounts for a very small aspect of how we can come to represent the temporal structure of our world, but by having looked at this system, we get a first glimpse at the ways that the representation of time is similar to and also differs from the representation of the non-temporal aspects of our world.

Chapter 4: The Temporal Structure of Experience: Against the Atomism

/ Extensionalism Debate

In the previous chapter we looked at the role of the circadian system in coordinating activity with the temporal structure of the environment. Specifically, it was argued that the circadian system counts as a genuine sensory system that provides a wide range of organisms with information about the approximate time of day. However, the extent to which the circadian system is implicated in coordinating animal behavior with the temporal world is rather restricted as it is not capable of producing the fine grained temporal discriminations that are apparent in our experience of the world and that are needed for the coordination of fine-grained motor activity.

In this chapter, we'll turn to the ways in which humans⁹⁷ gather information about time at the much shorter time scale of milliseconds to seconds. It is here where we find the temporal capacities that pervade our experience of the world. From perceiving the slight variations in phonemic structure, music, and movement to being able to navigate a busy sidewalk or catch a ball to feeling the painful sting of an awkward pause in conversation or a mis-delivered joke, all of these activities and experiences require that we be able to keep track of and notice the temporal structure of the events around us at very short time scales through multiple modalities.

As we'll see, there is no single system or mechanism that provides us with the ability to perceive or experience the temporal features of the world at this time scale. In fact, 'temporal experience' does not even pick out a single phenomenon or capacity. Rather, 'temporal experience' admits of a number of polysemous interpretations that are held together by the fact that the various uses or interpretations of the term apply to phenomena whereby we pick out, track, or represent the temporal structure of the world around us. Instead of the unitary approach, where temporal experience is a single phenomenon, I will

⁹⁷ The restriction is to humans at this point since we'll largely be talking about temporal *experience*.

argue for *the fragmentary model of temporal perception* according to which our capacity to perceive time is the product of a great many different mechanisms that represent distinct aspects of the temporal structure of our world.

While the argument for the fragmentary model of temporal perception rests on largely empirical considerations I will show how this model of temporal perception has direct implications for philosophical accounts of temporal perception. In particular, as was briefly mentioned in the previous chapter, there is tendency within philosophy and the cognitive sciences to explain the temporal content of perception by appealing to the temporal structure of the perceptual process itself. For instance, you may ask someone to explain in virtue of what their experience of a thunderstorm presents the flash of lightning as being prior to the crash of the thunder. A perfectly common reply would simply be to appeal to the fact that first they had an experience of the flash of lightning and then they had an experience of the crash of thunder. In this way, the temporal structure of experience has a role in explaining how experience has its particular temporal contents through a commitment to *the mirroring constraint*.

Mirroring: the temporal contents of an experience mirror the temporal structure of that experience itself.

Now, quite a bit of the philosophical literature on temporal experience is centered around debates about the truth of *the mirroring constraint*. *Extensionalists* claim that the mirroring constraint, in a sense to be further described below, must be true. *Atomists*, on the other hand, deny the mirroring constraint and claim that the temporal contents of experience rarely (or never!) mirror the temporal structure of experience itself. One thing, however, that both views have in common is that they agree that the

mirroring constraint is something that is either true or false across the board.⁹⁸ That is, there should be a univocal answer to whether the mirroring constraint is satisfied or not for temporal experience in general.

The main conclusion of the fragmentary model for the debate over the mirroring constraint is simply that we should not expect a univocal answer to whether the mirroring constraint is true or not. Temporal experience is not a unified phenomenon and therefore we shouldn't expect a unified response. Second, and perhaps more importantly, I argue that once we look at the distinct aspects that make up the fragmentary model of temporal experience we'll see that in fact there isn't a single answer to whether or not the mirroring constraint is satisfied. Rather, distinct aspects of our temporal experience are subserved by mechanisms that employ radically different strategies for representing the temporal structure of our world. Some of these employ representational strategies that respect the mirroring constraint while others do not.

Once we see that there isn't a single answer to the question of whether or not temporal experience satisfies the mirroring constraint the debate over mirroring begins to lose much of its philosophical appeal. Instead, I argue, that the genuine puzzlement about temporal experience that originally motivated discussions of the mirroring constraint is largely left untouched. Instead, in order to properly understand how temporal experience is possible and how it is that we come to perceive the fine-grained temporal structure of our world we need to more closely look at how theories of content can explain these various representational capacities.

The chapter will go as follows: In section 1, I say a little to clarify the target of this chapter and describe what we are trying to explain when we talk about temporal experience. In section 2, I lay out the

⁹⁸ Explicit statements of this assumption are in fact hard to find. Perhaps because the idea of temporal experience not being a single unitary phenomenon simply doesn't arise in the literature there is no need to assert this assumption. However, evidence *that* people in the debate hold this assumption comes from the argumentative strategy that is so common in the literature. Both atomists and extensionalists will attempt to find individual counterexamples to their opponent's view and then infer that their position therefore must be true. This sort of inference can only be justified on the assumption that temporal experience is a unitary phenomenon.

debate over the mirroring constraint. In sections 3 through 5 I lay out the standard arguments for extensionalism and atomism and argue that they fail to establish their intended conclusions. In section 6 we turn towards the fragmentary model of temporal experience and how it undermines the prospects of a univocal answer to whether or not mirroring is satisfied.

4.1. What is temporal experience?

Almost every introduction to a discussion on temporal experience begins with an appeal to introspection.⁹⁹ This discussion will be no different. Simply reflect on the way that the world is presented to you. Watching people walk down the street, you become aware of how their bodies move, how quickly or how slowly some people walk, you become aware of how people react to the honking of the horn at the intersection. As you see a person ride on their bike down the street, just by looking at their trajectory you can predict that they'll likely run into the garbage truck that's slowly pulling out of the alley. When you are standing there at the crosswalk waiting for the light to change the very duration of your wait seems to be capable of weighing heavily on you. Examples like this are everywhere. Consider a simple case of experiencing a thunderstorm pass by. When the storm is in the distance you will see the flash of lightning well before you hear the crash of thunder. As the storm approaches, you begin to notice that the delay between the flash of lightning and the crash of thunder becomes shorter and shorter. Eventually, the delay between the two types of events becomes so short that the two events seem to overlap in their time of occurrence and eventually they appear to occur simultaneously. All of these examples involve the temporal structure of the world being presented to us with much the same immediacy with which we are presented with the non-temporal properties that we take to populate our experience of the world.

⁹⁹ A reasonably short list of discussions that begin this way include (Broad, 1923; Dainton, 2008, 2011; Grush, 2007; Hoerl, 2009; Phillips, 2010).

Even beyond cases in which the temporal structure of the world clearly seems to impact our experience of the world, there are other cases in which we know that keeping track of the temporal structure of events is vital to how we perceive and experience the world. As Buonomano and Karmarkar (2002) mentioned, even the perception of speech involves keeping track of the fine grained temporal structure of your world. The phonetic difference between the Hendrix line “kiss the sky” and the command “kiss this guy” is merely one of timing. Stretching out one phoneme and not another will change the words that we hear. The difference between “kiss the sky” and “kiss this guy” might not strike us upon introspection to be a difference in the temporal structure of the utterance, but we know that this is all it consists in.

These are the sorts of perceptual phenomena that fall under the banner of ‘temporal perception’ or ‘temporal experience’. However, as is evidenced by many of these examples there is a difficulty in clearly laying out what counts as *temporal perception proper* and what begins to fall under the responsibility of *temporal cognition*. At some stage, it seems safer to say that our ability to track the temporal structure of our world is genuinely perceptual. Take motion. The perception of motion seems to require the perception of time, since the perception of motion seems to simply be the perception of the change in something’s location over time. We also think that motion is something that can be genuinely perceived with the same immediacy as color or shape. In fact, for many people a safe way of operationalizing the contents of perception is through the phenomena of *pop-out* (Wolfe & Horowitz 2004). A stimulus, or stimulus feature, pops out when the time that it takes us to positively detect a stimulus amongst a background of distractors remains constant despite changes in the number of distractor items. Since motion can cause pop-out in adults, we have some reason to think that motion is genuinely perceived. If motion perception involves some temporal representation, then temporal representation can be perceptual.

At short time scales, time does seem to present itself perceptually. However, at some stage it seems as though the representation of temporal structure shifts away from truly perceptual processes to

something closer to cognition. Take the experience of the thunderstorm again. As the storm approaches, the delay between the flash of lightning and the crash of thunder becomes shorter and shorter, and the slow narrowing of the delay is something that seems to be apparent in experience. As the storm passes, the once simultaneous lightning and thunder begin to drift apart. At some point, when the storm is in the distance, the delay between the two starts to become so great that several seconds separate the two experienced events. At some point, it can even become difficult to tell which crashes (or rumbles) of thunder go with which flashes of lightning. Eventually, it even seems as though you need to exert some effort in keeping the flashes of lightning in mind to make the connection between the lightning and the rumbling thunder. Even more extended than the slowly moving storm, and something that complicates our understanding of temporal experience, is that we'll use the term 'experience' to describe our access to the very short events of the individual flashes of lightning, to the moderate length events of the slow drifting apart of the onset and offset of the flashes of lightning and crashes of thunder, to the long scale events of experiencing the thunderstorm as a whole (i.e. think of cases in which someone might ask you whether you've ever *experienced a hurricane*).

We apply 'experience' to the representation of a number of different events at different time scales. Yet, it's not clear whether all of these "experiences" are truly perceptual. It seems clear that the experience of the thunderstorm as a whole involves quite a bit of cognitive machinery. It also seems clear (although some have rejected this¹⁰⁰) that our "experience" of the very short temporal structure of the individual flashes of lightning seems to be a genuinely perceptual phenomenon. However, it's unclear how to make sense of the intermediate cases. Appeals to the philosopher's tool of *what it's likeness* also doesn't seem to help, as in all these cases there seems to be a way in which the "experience" of the world at these different time scales has *something that it is like* to it (e.g. there is something that its like to go through a hurricane as opposed to a quick storm. The actual ordeal of the hours long storm is crucial to

¹⁰⁰ Examples of people who have rejected that even at this time scale temporal experience is a perceptual phenomena have included (Chuard, 2011; Gallistel, 1996; Prosser, 2007; Reid, 1855).

what it's like and this what its likeness cannot simply be reduced to the what it's likeness of the individual components of the storm).

Also, we can't appeal to the boundary of the sensory system to delimit what "experiences" count as genuinely perceptual. As we saw in the previous chapter, the boundary between the sensory systems and cognition is not a fixed one and can be shifted depending on what our explanatory aims are. Instead, for our purposes here, let's simply restrict ourselves to the shortest time scales from milliseconds to perhaps a couple seconds. At this scale, we are more secure in claiming that we are discussing perceptual phenomena. But further whether or not temporal experience is a truly perceptual phenomenon isn't the most important thing for us to figure out. We want to understand how we come to have temporal experience. If this phenomena lines up with perception, then so be it, but if it doesn't line up, then we still have the task of understanding our temporal experience. So, let us take 'temporal experience' in this broad sense that is not necessarily restricted to falling entirely within the boundaries of perception.

4.2. The classic dialectic and the mirroring constraint

After specifying what they mean by temporal experience, most philosophers proceed to motivate their particular discussion of temporal experience by stating a puzzle that calls into question the very possibility of temporal experience. While various authors have put this puzzle in slightly different terms, we can follow Sean Kelly's formulation. He says:

How is it possible for us to have experiences as of continuous, dynamic, temporally structured, unified events given that we start with (what at least seems to be) a sequence of independent static snapshots of the world at a time?

(Kelly, 2005, p. 2)

Kelly's worry, as it's been shared by (or at least mentioned by) many others in some form or other, can be summarized in the following way: Experience puts us into immediate and direct contact with how the world is *right now*. Imagine asking someone "what are you experiencing *RIGHT NOW!*?" as they watch a ball fly through the air. If they report accurately, they will report seeing the ball at a particular location. They might even point to a specific location along the ball's trajectory. In this way, experience seems to be composed of a series of snapshots of the world at a moment.¹⁰¹ If in a moment's time you ask the person the same question, they likely just point to a location a little bit further along the trajectory. However, if experience is to ever represent a temporally extended event (or temporally extended property of an event), like a duration or a temporal sequence or ordering, then experience must present us with more than a static snapshot of the world, as static snapshots by definition cannot present us with dynamic changing phenomena.¹⁰² As a result, we have a tension. We seem to have temporal experiences, that is, experiences that present to us the temporal structure of the events around us, but given the snapshot model of experience, we seem to never have experiences that actually present us with this temporal structure. This is the puzzle. How does a series of snapshots come to have temporal content that extends beyond a durationless instant?

¹⁰¹ Another motivation of the snapshot model can be found in the work of the early modern empiricists (e.g. Hume (2000) and Locke (1689)). According to these empiricists we can distinguish between *complex* and *simple* impressions (what we would now call perceptual experiences). *Complex impressions* are those that have aspects that *could be experienced by themselves*, while *simple impressions* cannot be divided into parts that could be individually experienced. Now, imagine the experience of a succession of events (say, lightning and thunder). This experience will count as a complex one because it's conceivable that we could experience the lightning without having an experience of the thunder. Consider now the experience of the lightning. This event, as perceived, will have a duration. This duration is also complex because it consists of a beginning, middle, and end, each of which could be experienced in isolation (e.g. the particular pattern of light at the beginning of the flash of lightning can be experienced independently of the particular pattern of light at the middle or end of the flash). We can continue this process until we arrive at the simple experiences that present the world to us as it is at a single moment in time. In this way, all experience is built up of static snapshots of the world.

¹⁰² We can take a snapshot of a dynamic event, but it's not clear at all whether the dynamic aspect of the event will be presented to us in the snapshot. Take a snapshot of a shark jumping through the air, and you will have taken a snapshot of a dynamic event, but the actual speed (the actual temporal structure) of the event will not be made apparent in the snapshot.

Interestingly, perhaps only Kelly actually takes this tension to indicate a genuine puzzle that simply lacks a satisfactory resolution. Almost everyone else avoids the tension by either denying that experience is exhaustively composed of individual snapshots of the world or by denying that we in fact have experiences with temporal content.

Those that deny that we have temporal experience (Chuard 2011; Gallistel 1996; Reid 1855) typically do so by simply denying that we have temporal experiences that are genuinely *perceptual*. However, as we're framing the debate, the problem we are interested in is how to account for our temporal experience when we take experience in a broad sense. We cannot prejudge the situation and insist that our temporal experience is purely perceptual. So, even these people that deny that there are perceptual experiences of time still maintain that in a sense we have experiences of time in the broad sense that I am using it.

For those that do not deny that we have temporal experience, we find that their discussions quickly fall into a discussion of the mirroring constraint.¹⁰³ They deny that the snapshot model of experience provides an exhaustive analysis of experience. There is something more to experience than just snapshots of a static world.¹⁰⁴ So, according to them our experiences are capable of possessing temporally extended contents. The explanatory task for these authors, then, is to explain how it is that experience comes to have this temporally extended content. But, this simply rewords the initial question

¹⁰³ The historical reasons for this don't actually change the nature of the discussion to follow. However, if we look back in the literature we find that much of the initial debate over temporal experience (from Locke (1689) onwards) consisted of a debate over whether a *succession of experience* amounts to an *experience of succession*. That is, whether the actual temporal structure of experience could explain the temporal content of experience. This was particularly salient as the early modern empiricists, like Locke (1689) and Hume (2000), had views that endorsed this identification of a succession of experiences with an experience of succession. For our purposes, however, we can simply pick up from once this debate is firmly established in the literature as the appropriate sort of historical analysis would take us far afield from our main concerns.

¹⁰⁴ Views that attempt to supplement the snapshot model by positing some relation between snapshots that binds them together into something further than mere snapshots fall into this category of view.

that was behind the initial puzzle about how it is possible that we can have temporal experiences, that is, how we can have experiences with temporally extended contents.

In trying to answer this question is where we confront the mirroring constraint. As mentioned earlier, there is something very intuitive in how we might explain how experience has its temporal contents. When asked why an experience of a succession of events, say an experience of a flash of lightning preceding a crash of thunder, has its particular temporal contents, it seems clear that it has to do with the fact that first there is the experience of the flash of lightning and then there is the experience of the crash of thunder. Extensionalists (Dainton, 2008, 2011; Hoerl, 2009; Martin, 2002; Phillips, 2010, 2014a, 2014b; Rashbrook, 2013) build on this intuitive explanation of the content of experience. Phillips (2010) goes as far as calling the view *the naïve view of temporal perception*. According to extensionalists, the temporal structure of the world that is represented in experience is mirrored by the temporal structure of experience itself.

Atomists, on the other hand, deny that mirroring has to hold. In fact, some (Grush 2005; Lee 2014b) argue that mirroring never or rarely ever holds. Instead, according to atomists, the temporal structure of events that is represented in experience can come apart, often in extreme ways, from the temporal structure of experience itself. While it may be that we *experience the flash of lightning as preceding the crash of thunder*, atomists will deny that *the experience of the flash of lightning must precede the experience of the crash of thunder*. Instead, they will typically claim that if experience presents the lightning and thunder as standing in some temporal relation, then we will have an experience of both of these events at once – that is the experience of the lightning and the experience of the crash will occur at the same time.

Now, the debate over the mirroring constraint requires an important clarification.¹⁰⁵ We cannot restate the debate, as Dainton (2014) does, as being over whether or not experiences themselves are temporally extended, since doing so would trivialize the debate. To show why this would trivialize the debate, let's assume at least the minimal relation between mental states and the physical whereby mental states are realized by physical events. If mental states are realized by physical events, then it's plausible to assume that the temporal properties of mental events will match those of their realizers – e.g. if experience *e* is realized by physical state *p*, then as long as *p* obtains *e* will be occurring. One of course could insist on there being a more complicated relationship between realizer timing and the timing of mental states (perhaps mental states only exist at the last moment of the realization), but such a view would require some strong motivation, which I do not see.¹⁰⁶

If we combine this claim about the relation between mental state and realizer timing with the claim that the physical realizers of all experiences are temporally extended physical processes, then we would have the conclusion that all experiences are themselves temporally extended. Given that we have reasons to think that experiences are realized by temporally extended physical or neural processes¹⁰⁷, then it would simply follow that all experiences are themselves temporally extended. If we were to understand the atomist as denying that experiences are temporally extended, as Dainton describes them, then their

¹⁰⁵ This clarification can be found in (Lee, 2014b).

¹⁰⁶ Michael Pelczar (2015) seems to have some view of this sort alternative sort. In his book he defends a form of idealism where experiences lie outside of space and time and it is only within and through experience that something like the physical world is constructed. How then is realizer timing and experience related? Well, on his view it's a rather complicated story. But, importantly, we don't have to go down the path of trying to tease apart his view. His argument relies on an inference (in the broadest of strokes) from there being a failure of the mirroring constraint and us not having any guide to the temporal structure of experience to the claim that experiences lie outside of space and time. However, as will become clear in this chapter, we do in fact have some alternative guide other than introspection for deciphering the temporal structure of experience, and therefore needn't accept his conclusion.

¹⁰⁷ Look at many of the sorts of accounts that arise in the literature on neural correlates of consciousness (e.g. Crick & Koch 1990, 2003; Dehaene et al 1998; Lamme 2006; Llinás et al 1998). This involve information being maintained in certain processes for some amount of time. Or more generally, neural representation is often understood in terms of the firing rates of neurons and are only accessible via the binning of incoming signals by downstream neurons over a short temporal interval (Purves et al 2001; Sincich et al 2009). All of these processes require time.

position would be false. But notice that this would be so for nothing having specifically to do with temporal perception. Rather it would follow from a claim about experience in general.

Instead, the debate is truly about whether the temporal structures represented in experience *mirror* the temporal structure of experience itself. An experience has a temporal structure that mirrors its temporal content just in case each temporal stage that is represented by the experience is realized by a temporally distinct physical process and these physical processes stand in the same relations to one another as the events that are represented by the experience are represented as standing to one another. If experience presents the flash of lightning as preceding the crash of thunder, then this represented temporal structure will be mirrored by the temporal structure of the experience just in case the realization of the experience of the flash of lightning precedes the realization of the crash of thunder.

The extensionalist position is then one that holds that there is a mirroring between the temporal structure of events that are represented in experience and the temporal structure of experience itself.¹⁰⁸ And the atomist denies this mirroring by denying that we can divide the realizer of an experience of a temporally structured event (or series of events) into temporally distinct parts that each realize distinct aspects of the overall experience (see figure #4).¹⁰⁹

¹⁰⁸ It needn't be the case that for all temporal stages of the process (however thin we might slice them) that there is a semantic interpretation of the system. Rather, all that's required is that there be some temporal stages that are interpretable.

¹⁰⁹ We can further refine our analysis of mirroring by making precise the sort of temporal properties in experience that are mirrored in the structure of experience (Lee, 2014a; Phillips, 2014b). For instance, *metric mirroring* involves the mirroring of not only temporal ordering relations but also durations. While *topological mirroring* involves the mirroring of temporal relations. However, for our purposes here we can ignore these variations as the arguments to be presented apply to all types of mirroring.




Figure 4: Mirroring and non-mirroring experiences.

With the debate in hand, it's now time to turn to the arguments that are found in the literature in favor of each view. I will argue that all of the existing arguments fail to establish that temporal experience is either atomistic or extensionalist. With the arguments for these views out of the way, we'll turn in the later sections to arguments that neither position is correct as a general characterization of temporal perception. Instead, once we realize the 'temporal experience' does not pick out a single unified phenomenon, we'll be in a position to see that some aspects of our temporal experience are atomistic while others are extensionalist.

4.3. The case for extensionalism: introspection

In defenses of extensionalism we find that introspection looms large (Dainton, 2008, 2011; Hoerl, 2009; Phillips, 2010, 2014a, 2014b; Rashbrook, 2013). While the specific ways in which introspective

evidence is used to establish the truth of extensionalism differ depending on which author we're considering, all of the introspective arguments have essentially the same structure. In this general form the argument runs as follows: upon introspection we find that experience presents us with temporal contents that simply *could not be accounted for* by atomism, therefore atomism is false based purely on introspective evidence alone. In order to evaluate these arguments, then, we're forced to ask two questions. First, just what is this introspective evidence? Second, does the presence of this evidence really require an extensionalist theory? Let's take each of these questions in turn.

Given our initial characterization of temporal experience in section 1, we can be fairly brief about the first question. Experience presents us with a temporally structured and dynamic world. In experience, the temporal structure of the events around us, their durations, their ordering, etc, is represented. Furthermore, there is a sense of the passage of time in experience such that at any given moment we are not merely aware of the present moment, but we have some *current* awareness of what has just passed and according to some we have some current awareness of what will likely occur.¹¹⁰ All of the examples so far discussed involve experiences that attribute temporal properties to objects and events that are distinct from the experience itself. However, there is a somewhat compelling feeling that when we introspect our experiences we are not only confronting the temporal structure of the world as it is presented to us in experience, but that we are also being made directly aware of the temporal structure of experience itself. It is this final intuition that seems to motivate Phillips' description of extensionalism as the *naïve view* since it's this intuition that seems to be fundamental.

Certainly, one aspect of the naïve intuition is correct. We often do make true introspective judgments about the temporal structure of experience. Focus on this page of text and shift your eyes from

¹¹⁰ This aspect of our temporal experience is often called *the specious present* following James (1890) who attributed the term to a mysterious "E.R. Clay". However, I'm choosing to not use the term as over the years the quite a lot of baggage has started to come along with the term. Specifically, many authors tend to build in some form of atomism (Grush 2006) or extensionalism (Dainton 2014) into their characterization of the view. For an interesting history of the development of the specious present doctrine see Andersen (2014).

the top line of text to the bottom line of text. Continue to move your eyes to fixate on one and then the other. When you do this it's tempting to make the judgment that first you are experiencing the top line of text, and then you are experiencing the bottom line of text, and back and forth. It seems, then, that here we're in a position to make a judgment about the temporal structure of experience itself. It's likely that our judgments in these cases are correct. However, it's not clear how much can be made of this sort of case.

Do we have *direct* access to the temporal structure of experience? In the case just described, there is a clear cut candidate worldly event that mediates our judgments about the temporal structure of experience. We experience our movements, our guided shifts in eye movement, and we note the temporal relations between these worldly events, and then *infer* the temporal structure of our experience.¹¹¹ In other cases, the same sort of inference seems to be what drives our judgments about the temporal structure of experience. We judge that the experience of the lightning precedes the experience of the crash of thunder since experience presents these events represented by experience as standing in this order. We readily make this sort of inference. Furthermore, as we'll see in the next section, it is because we make this sort of inference that a number of temporal illusions are so striking. Given that we do not directly experience the temporal structure of experience itself, and instead we infer from the presence of certain temporal contents to the structure of experience, we then simply need to ask whether the inference is warranted. Could the very same contents be presented to us in experience but not in a way that satisfied the mirroring constraint.

Ian Phillips (2010) has notably argued recently that this inference from the perceived temporal structure of the world to the temporal structure of experience is infallible. His argument goes as follows:

¹¹¹ This point is closely related to the topic of transparency although it's not identical to it. Transparency is typically taken to be a general claim about whether or not we are ever aware of the *intrinsic properties of experience*. Transparency advocates (Byrne 2011; Dretske 1995; Tye 1995, 2000) deny that we ever have any access to the non-representational or intrinsic features of experience. But notice, that these claims about temporal perception can be made without endorsing a universal account of transparency.

Consider again the experience of the flash of lightning (L) preceding the crash of thunder (T). Perception makes us aware of the temporal relation between L and T, and on the basis of this information, we form the judgment that the experience of L preceded the experience of T. Furthermore, as Phillips puts it, when we experience L before T it certainly seems to us that the experience of L precedes the experience of T. This just is the naïve intuition. But, building on the Cartesian idea that we are infallible when it comes to our own experiences, Phillips claims that there cannot be any *seems / is* distinction when it comes to experience itself. That is, if experience *seems* to be a certain way, then experience *is* that way. Phillips calls this general principle about experience *the seems → is principle*. Notice the strength of Phillips' claim. On the basis of the introspective evidence alone with the *seems → is principle*, it follows that extensionalism is *necessarily* true. That is quite the claim to be made on the basis of introspection alone.¹¹²

Unfortunately for Phillips' defense of extensionalism, we have reasons to doubt the *seems → is principle*. Lee (2014a) surveys a number of responses to Phillips' argument, however, we'll take a somewhat different tactic in this section. The *seems → is principle*, and in general any inference from the contents of experience to the structure of experience itself, requires some connection between the vehicles of experience and their contents. In order to show that the initial introspective data simply do not favor extensionalism or atomism over the other, I will argue that no viable theory of content that would explain how the contents of experience map onto the vehicles of those experiences depends on atomism or extensionalism. As a result, we have no theoretical reasons for supposing that atomism or extensionalism is true on the basis of the contents of experience itself.

As mentioned in the introduction, there is very little consensus in the philosophical literature on how to best understand the metasemantics of mental representations. Without any theory of content as *the*

¹¹² Others have also noted that Phillips' argument seems to be make an implausibly strong modal claim. See Frischhut (2014) and Lee (2014a).

theory of content that explains how the representational contents of perception map onto their representational vehicles, we can proceed by considering the main theories of content. However, even that process would turn out to be rather lengthy. Instead, we can get the same payoff by dividing existing theories of content into two general types. On the one hand, we have *vehicle neutral theories of content* and on the other hand we have *vehicle dependent theories of content*. Roughly, the divide between the two classes of theories has to do with whether the content determining properties of the representational vehicle are *local* or *situational*.

All representational theories of mind have to posit some complexity to the representational vehicles that the theory appeals to in explaining psychological phenomena. According to many (Carruthers, 2006; Fodor, 1975, 2008; Marcus, 2003) in order to explain productive representational capacities¹¹³, like conceptual thought, we must accept theories that posit mental representations that exhibit some level of sentence-like syntactic structure where basic representations are combined in various systematic ways to produce complex thoughts. Beyond conceptual thought, many have thought that perception itself requires an underlying representational system that has some complexity in the structure of the underlying representational vehicles.¹¹⁴ We can see that this way of positing structure to the underlying representational system is a direct response to the particular requirements of explaining some psychological phenomena.

In this way, all representational theories must posit some complexity in the underlying representational vehicles in order to explain particular psychological phenomena.¹¹⁵ The distinction

¹¹³ That is, representational capacities that allow for an in principle infinite number of representational states.

¹¹⁴ Salient examples of this are Anne Treisman's *Feature Integration Theory* (Treisman & Gelade 1980) and Biederman (1987) and Marr's (1982) computational models of object recognition.

¹¹⁵ The debates over connectionism vs classicism that occurred in the 80's and 90's was to a certain extent over whether or not this claim is true. Some extreme forms of connectionism claimed that connectionist systems, while representational, lacked any syntactic structure in that they could not be decomposed into semantically meaningful units. For these theories, the semantic content contained in a connectionist system was distributed across the entire network. Many of the debates, however, centered around whether such a structureless representational system could account for the productive and systematic properties of conceptual thought.

between *vehicle neutral* and *vehicle dependent* theories of content comes in when we try to understand whether the semantic content of the individual representations, not the sentence-like representations but the word-like representation, places any demands on the structure the underlying representational vehicles.

Vehicle neutral theories of content explain how individual representations acquire their content in terms of the *situational* properties of the individual representational vehicles. In other words, they place no requirement on how representational vehicles are structured. Instead, the explanation of how representations come to have their contents is given by appealing to how the individual representational vehicle relates to things beyond itself. Examples of these theories include *causal theories* (Fodor 1994; Prinz 2002; Stampe 1977), which explain the content of a representation in terms of its typical (or privileged) causes, *informational theories* (Dretske 1981; Skyrms 2010; Usher 2001), which explain the content of a representation in terms of its informational content (see Chapter 3 for a discussion of these theories), and *teleological theories* (Dretske 1995; Millikan 1989, 1995; Ryder 2004), which explain the content of representations in terms of their selection histories. In every case mentioned, the content of a representation isn't explained by something internal to that representation, but instead by how that representation relates to something beyond itself.

Vehicle dependent theories, on the other hand, explain how representations acquire their content primarily on the basis of how those representations are internally structured. Most notably amongst this category of theories are *resemblance theories of content* (Hume 2000; Locke 1689)¹¹⁶. According to this class of theories, in addition to any syntactic structure needed to explain the use of complex representations in the production of behavior, vehicle dependent theories require that the individual

¹¹⁶ Interestingly while still many people talk about resemblance theories, it's rare to find someone that explicitly endorses the view. Rather, as being discussed here, people often exhibit some sort of implicit commitment to the view.

constituent representations also have a particular structure to them in virtue of which they possess their particular semantic content.

Now, let's return to the issue at hand, the mirroring constraint and the debate between atomists and extensionalists. The particular argument that we've seen on behalf of the extensionalist claims that the introspective evidence concerning the temporal content of our experience demands that we accept the mirroring constraint. However, we've seen that this argument hinges on whether or not there is any theory of content that would explain why we should demand a particular structure to experience in order to explain how experience can represent the temporal structure of our world. In what follows, I'll argue that only vehicle neutral theories of content provide viable explanations of how experience comes to have its content, and that no vehicle neutral theory would provide the extensionalist with the connection they require between the content of experience and the structure of our experiential vehicles.

Let's start by seeing how a vehicle neutral theory of content can explain how experience could have the temporal content that it does without requiring one way or another that the mirroring constraint be satisfied. Consider the following scenario:

Remote Timing: A control operator needs to monitor the order and temporal delay between a pair of distant events *a* and *b*. Since the operator cannot directly witness *a* and *b* the operator must be able to make her decision on the basis of the readout of a local representation of the events. Now, in the scenario the operator has two redundant systems available to her.

Mirroring system: the first system represents the temporal relation between *a* and *b* by presenting a pair of tones, *tone₁* and *tone₂*, such that *tone₁* stands for *a* and *tone₂* stands for *b*. The temporal relation between *a* and *b* is represented by the very temporal relation between *tone₁* and *tone₂*.

That is, the temporal structure of the representation of the temporal relation between a and b mirrors the actual represented relation between a and b .

Non-mirroring system: the second system represents the temporal relation between a and b by a pair of lights. The first light can take be in either of two states, blue or green. The second light can be in a range of different states corresponding to slight variations in the brightness of a white light. The color of the first light, blue or green, indicates whether a or b occurred first. The brightness of the second light indicates the temporal delay between the two events (if the light fails to turn on, then the two events were simultaneous). Here we have a system that while representing the temporal relation between events a and b fails to satisfy the mirroring constraint as the system does not decompose into temporal parts that correspond to or mirror the temporal stages of the represented event.

Our operator can use either system with equal results. Further, as mentioned, these two systems differ in their respect of the mirroring constraint.¹¹⁷ That there could be these two types of systems shouldn't be a great surprise to anyone. But, the purpose of this toy example is to show that our theories of content are sufficiently flexible to explain how either system of representation acquires its content. Since it would be a tedious task to go through each vehicle neutral theory, I'll discuss Dretske's information theoretic account of content as a representative example of this type of theory.

¹¹⁷ Another interesting difference between the two systems is that the mirroring system could be interpreted as only implicitly representing the temporal relation between a and b whereas the non-mirroring system requires an explicit representation of time. But, we should resist the urge to infer that this pattern generalizes such that implicit representations of time respect the mirroring constraint while explicit representations do not.

Information theoretic accounts of content begin with the notion of *mutual information*. While different authors have taken this formal tool in different directions, Dretske gives a highly influential version of the theory in his *Knowledge and the Flow of Information* (1981). According to Dretske, the information content of a system is determined by the following biconditional:

A [system] r carries the information that s is F if and only if the conditional probability of s 's being f given r (and k), is 1 (but given k alone, less than 1).

(Dretske, 1981, p. 51)¹¹⁸

As long as certain background channel conditions, k , are in place, the state of the representational system, r , will increase the probability of some other system's being a certain way to 1. In other words, r 's being a certain way guarantees that some state of affairs out in the world obtains.¹¹⁹

When we look at both of the systems that the operator has at her disposal we find that in both cases given that certain channel conditions obtain, e.g. that the systems are properly connected to the events a and b , if the systems enter into particular states, then the events a and b will stand in a particular temporal relation to one another. Furthermore, neither system differs in the information they carry about the temporal relations between a and b . By simply hearing the reports of the operator about the temporal relation between a and b , we'd have no way of knowing which system she was using.

The analogy should be clear. If the extensionalist attempts to argue for their position on the basis of the introspectively available contents, then they need a reason to think that no vehicle neutral theory of content should be what accounts for the temporal contents of experience. Otherwise, the fact that experience has a certain content will always underdetermine how experiences themselves are temporally

¹¹⁸ The analysis given here is a direct quote from Dretske but is equivalent to the analysis given in chapter 3.

¹¹⁹ The demand that the conditional probability be raised to 1 is rather demanding and has been a source of problems for information theories of content. However, see Usher (2001) for how an information theoretic account of content can be developed that weakens this requirement.

structured. The most natural way for the extensionalist to establish their conclusion is to adopt some form of a resemblance or vehicle dependent theory of content.¹²⁰

As we've already discussed, there does seem to be a naïve pull to explain how experience has its temporal contents by appealing to something that is very reminiscent of a resemblance theory. Recall the intuitiveness of explaining why experience might present the flash of lightning as preceding the crash of thunder by appealing to the temporal order of our experiences of the flash and of the lightning. However, it's important to note that this sort of explanation is very foreign to how we try to understand most cases of mental representation. Consider what you might say if someone asked you to explain how their experience comes to represent something as being red? For the scientifically naïve, there seems to be little that they could report. Experience represents something as being red, simply because something looks to be red. Perhaps closer to the temporal case, we can also consider what you might say if someone were to ask you in virtue of what your experience represents something as *being to the left of something else*. Again, for the scientifically naïve there seems to be little to say. Surely, there seems to be much less of a pull to say that my experience of the one thing is to the left of my experience of the other.

For the scientifically informed the answer they give might be somewhat different. In the case of red, we might appeal to the differential stimulation of different photoreceptors and some degree of cortical processing that leads to color opponent systems, which ultimately define our perceptual color space. In the case of space, a more complicated story might be given that appeals to various spatial maps that are produced by the perceptual system. Here we find scientific explanations of spatial perception that *may* in fact be amenable to a resemblance theory. The various spatial maps of the perceptual system (for instance the retinotopic maps of the early visual system) preserve the spatial relations between retinal

¹²⁰ Of course an extensionalist could deny that any theory of content is sufficiently worked out to run this sort of argument. However, then it's unclear how they could make an inference from the content of experience to the temporal structure of experience. Some story about how contents map onto the structure of representational vehicles is needed for the argument to go through.

locations in the spatial configuration of the individual maps, such that two adjacent locations in a retinotopic map will encode information about spatially adjacent retinal locations. However, not all spatial properties are plausibly represented through a resemblance as the cortical maps are typically understood as being two-dimensional cortical networks, yet we often represent three-dimensional spatial relations. So, not all spatial relations could be accounted for through resemblance.¹²¹ Importantly, though, to whatever extent the spatial properties of vision (or perception more generally) are explainable through a resemblance theory, it remains that the viability of this theory is a purely theoretical one. Introspection tells us nothing about these systems.

So, resemblance isn't typically appealed to in the explanation of mental content. Even though introspectively it seems as though the timing of experience has something to do with its temporal content, to simply insist that it does have a role to play here is simply question begging.

Now, the extensionalist is actually in a rather odd position here. If they attempt to establish the truth of extensionalism on the basis of the introspective evidence by arguing for a resemblance theory of content, then they cannot make any direct appeals to the temporal structure of experience (whether introspective or scientific) in support of a resemblance theory as the temporal structure of experience is the very issue that we are trying to understand. But, then that means that they have to give general reasons for accepting a resemblance theory of content. But again, resemblance simply isn't a well accepted account of mental content in general, so there is a problem for the extensionalist here. They simply do not have the resources to establish the resemblance theory that their argument needs.

¹²¹ In fact, the existence of cortical maps doesn't entail that they represent their environment through resemblance. The same considerations as mentioned above for the case of the mirroring and non-mirroring representational systems and the vehicle neutral theories of content apply here as well. We can explain how a maplike spatial representation encodes spatial content in a way that does not appeal to resemblance. It might simply be that the maplike configurations of cortical regions could arise for other reasons. For instance, perhaps there is an advantage in having nearby neurons encoding nearby spatial locations as this facilitates activation of neurons that code for a particular region of space, thereby making processing of motion (for instance) more efficient.

To make concrete the sorts of worries that arise for resemblance theories in general let's briefly go over a few of the common problems. To begin with, as we've mentioned, resemblance isn't typically thought to explain the contents of mental representations in general, and more specifically there are a wealth of temporal representations that also seem to not fit well with a resemblance theory of content. Take temporal concepts, like those expressed by the sentence "the big bang occurred roughly 13.8 billion years ago", it is not expected that anything like resemblance explains the contents of these temporal representations. Clearly the thought doesn't take 13.8 billion years to think! Similarly, for non-temporal representations we don't take resemblance to explain the content of experience. Grush (2006) has a nice example to illustrate this point. Consider the case of *brain gray* (the term given to the experience of a roughly gray expanse in the absence of any retinal stimulation). Imagine a resemblance theorist attempting to explain the experience of brain gray by pointing out that the brain itself is gray. They might explain that their theory actually predicted the properties of the representational vehicle, and lo and behold, they were correct! Clearly, this isn't a good explanation of the phenomena of brain gray. Mere resemblance isn't enough to explain mental content. Whatever property of the representational vehicle explains the content of the representation must be capable of varying with what it represents and must be capable to causally influencing downstream processes. Neither of which seems to be captured by the color of the brain.

Similar worries arise in that the properties of the vehicle may fail to reflect the content of the representation (and vice versa). Consider that experiences, i.e. token representations, instantiate a wide range of properties, including both temporal and non-temporal properties, yet most of these properties fail to ever, even accidentally, appear as contents of the representation. Take any token experience. That experience will occur at a particular moment in time (e.g. *August 29th 1997*) and bear a host of temporal relations to other events, both mental and non-mental, (e.g. *being later than the experience of drinking coffee I had earlier today*) yet neither of these sorts of temporal properties possessed by this representation ever appear as contents of the representation. Some sort of explanation is needed to explain

why only some of the properties of the representational vehicle are taken on as contents of the representation.

A defender of a resemblance theory for temporal representation might say there is an important difference between the case of brain gray and the color of the brain and the temporal properties of neural systems. It's plausible that some temporal properties of neural systems, say *being in a certain state for a certain amount of time*, are in fact capable of influencing downstream processing in a way that the color of the brain isn't. That a neural system is gray (or blue, green, or any color) by itself has no impact on the behavior of downstream processes (neurons, in this sense are color blind). But, if a system is in a certain state, say it has a certain level of activity, for a certain amount of time, then it will have influenced downstream systems in a way that differs from how the system would have been influenced if the activity had had a shorter duration.¹²²

This is, however, where the extensionalist is backed into a dialectical corner. If in order to support their claim for a resemblance theory they need to engage in an empirical investigation of the temporal structure of experience, to see if these are viable candidate properties for explaining the temporal content of experience, then they undermine their introspection based argument for extensionalism. Their arguments, in particular Phillips' argument based on the *seems* → *is principle*, were meant to be based on introspective evidence. However, it seems as though they simply cannot support their conclusion on this basis alone. They need to appeal to an empirical investigation. But this leaves open that the resulting picture of how the temporal content of experience might relate to the temporal structure of experience might be in line with the atomist and violate the mirroring constraint.

¹²² In fact, on interventionist accounts of causation (Woodward 2005) this sort of manipulability is all that is required for there to be a causal relation.

4.4. The case for atomism

The extensionalist, we just saw, attempted to establish their position on the basis of introspection alone, but their argument fails. In order to use the introspective evidence in the manner that they want, they need to appeal to some theory that connects that introspective evidence, i.e. the temporal contents of experience, with the temporal structure of experience. But, they simply cannot provide reasons for thinking that such a theory is true without appealing to some other, non-introspective, means of determining the temporal structure of experience. In this section we'll now turn to looking at the arguments for atomism. These arguments come in two groups. The first, and by far the more common approach, is to appeal to the existence of a variety of temporal illusions. The second approach appeals to what Geoffrey Lee calls *the trace integration argument* (Lee 2014b).

The arguments from temporal illusion

The arguments from temporal illusion all have the same structure. Through certain manipulations, or in certain circumstances, we come to have temporal illusions that present the temporal structure of some events in the world as being a certain way. However, the temporal structure of the world, as it is presented to us in these illusions, cannot plausibly be mirrored by the temporal structure of the experiences themselves. As a result of the failure of mirroring in these cases extensionalism fails and atomism must be true. In what follows, I will sketch out some of these illusions and the arguments for why they show that mirroring fails and then I will look at how the extensionalist needn't accept the conclusion of the atomist's argument on the basis of the existence of temporal illusions alone.

Probably the most famous, and now classic, type of temporal illusions involves the influence of *postdictive effects*. In these cases, subjects perceive a sequence of events that are arranged such that the perception of later events influences the way that earlier events in the sequence are perceived. In essence,

the perception of the later events causes the perceptual system to re-write history by re-representing the earlier event as occurring in a way that it in fact did not occur. While there are quite a few different types of examples along these lines, we can discuss one in particular, *the cutaneous rabbit illusion*, since this example has been discussed at length in the literature on mirroring (Grush 2005, 2007).

The set-up for the illusion is the following¹²³: subjects are sitting with their forearm exposed. They then receive a total of 15 taps at three locations along their arm – 5 taps at the wrist, 5 taps at a location 5cm above the initial set of taps, and then another 5 taps another 5cm above the second location. Each tap is identical in its duration and each tap is separated from the previous tap by the same temporal interval (between 40-60ms depending on the specific trial conditions).

One would expect that subjects would report simply feeling 15 taps with 5 at each of three locations. However, that is not what subjects report. Instead, they report feeling 15 taps that are equally spread out between the initial location at the wrist (where the first 5 taps were located) and somewhere slightly beyond the location of the final 5 taps. The taps seem to *hop* up the arm (like a rabbit).

Initially, this might seem to be a purely spatial illusion. What the subject is misperceiving are the *locations* of the taps, and not their timing. However, the temporal aspect of this illusion arises when we try to account for the perception of the taps hopping up the arm. Let's label the moment of each tap t_1 through t_{15} . Now consider what the subject might experience at t_1 if no subsequent taps were to occur. They would simply experience a single tap at the wrist as occurring at t_1 . Now, ask the same question for t_2 through t_5 . At each of these moments in time subjects will accurately perceive the location of the taps. Now, ask yourself what the person experiences at t_6 and later. The situation begins to become more complicated. At each of these later moments, the subsequent taps cause subjects to perceive the earlier taps as occurring in different locations. So, at t_{12} , say, the subject will perceive the location of the tap at t_{12}

¹²³ For the original discussion of the cutaneous rabbit illusion see (Geldard & Sherrick, 1972).

as occurring at t_{12} but at some location between 5 and 10cm away from the wrist (between the second and third tapping location). But, also, at t_{12} , and these later tap times more generally, there is a re-writing of the earlier tap location. So, at t_{12} the taps that occurred at t_2 - t_{11} are re-represented as having occurred at locations between 5 and 10cm away from the wrist.

Grush (2005, 2007) argues that in order to account for this experience of the hopping taps it must be the case that at a moment, say t_{15} , the entire temporal interval prior to t_{15} , which included the previous 14 taps, must be represented since only with the representation of that interval can the perceptual system post-dictively change the perceived location of the taps at those times. In other words, for the later taps to influence how we perceive the earlier taps, it must be the case that the perceptual system at the later moment in time is representing the entire temporal interval beginning at t_1 . If it weren't representing these earlier times, then the perceptual system would be unable to represent the location of the taps at those earlier times. Since we have the representation of the entire interval with its particular temporal structure occurring at a later moment in time, we seem to have a violation of mirroring. At t_{15} we have a representation that encompasses the entire interval at once. Since mirroring fails we have a counterexample to extensionalism, so atomism must be correct.¹²⁴

In response to the illusions involving post-dictive effects, extensionalists typically give the same response. Experience, they say, is delayed (Dainton, 2014; Phillips, 2014b). Between the moment of initial stimulation of the sensory systems through the initial stages of perceptual processes and the subsequent conscious experience of the world there is a delay that allows the systems responsible for experience to consolidate the sensory information and produce a coherent experience of the world. Once the information from the sensory systems has been consolidated in this supposed buffer, then it is presented in experience in a way that respects the mirroring constraint.

¹²⁴ Other postdictive effects, like the color phi phenomena (Dennett 1992; Dennett & Kinsbourne 1992) raise similar problems for mirroring and extensionalism.

While the picture the extensionalists provide is a coherent one, some atomists (namely Grush (2006)), ask whether the picture they provide is ultimately plausible. Consider the sort of delay that would be required for the cutaneous rabbit illusion. If the perceptual system has to wait until the entire sequence of taps is produced before it relays the relevant information to experience, then the delay between sensory stimulation and experience would be on the order of several hundred milliseconds (the taps in the original 1972 study were separated by intervals between 40-80ms). If that's the case, then experience would be unable to drive our actions, as a delay of several hundred milliseconds would severely damage our ability to act in the world. Grush is surely right about something. If our experience drives action, and experience is delayed by several hundred milliseconds, then we shouldn't expect to live very long lives. However, the extensionalist simply denies that it is *experience* that drives actions at this short time scale. Yes, perception can drive action, they might say, but experience is something else.

The extensionalist's claim has some independent plausibility as there's quite a bit of evidence that actually makes many doubt that it is experience that drives action. In non-temporal cases, the literature on the two visual streams (Goodale & Milner 1992; Ungerleider & Mishkin 1982) is often taken to show that the systems involved in bringing about conscious experience are distinct from those that drive action. The evidence that is typically taken to show the operation of two distinct visual streams involve cases where some spatial illusion causes subjects to report that the world looks a certain way to them, but their actions do not reflect this illusory spatial content, but rather reflect an accurate representation of the world. For instance, in the Titchener Circle illusion, subjects will report that two circles of equal size look as though they differ in size depending on whether those circles are surrounded by smaller or larger circles. Despite the fact that subjects report the center circles as being of different sizes their motor behaviors do not reflect this difference in perceived size. In both cases, subject will make the same hand movements (e.g. opening their thumb and forefinger to the same extent) when they are asked to grab the circles. Kozuch (2015) and Clark (2001) both take this sort of phenomena to show that the content that drives experience must be distinct from the content that drives action. Now, of course, these examples concern spatial

content¹²⁵ so it would be misguided to generalize to all experiential content not being involved in driving action. Surely, an architectural difference in the visual system cannot be used to make any claims about whether information in the auditory processing stream is divided into two streams, one for action and one for experience. But, at the very least, the extensionalist can rest on prior precedent that experience may not be driving action in order to make their claim that the mirroring constraint is salvaged by delaying experience.

So, as it stands, the extensionalist seems to have a way out of the problems that postdictive effects raise for the mirroring constraint. Whether this approach is ultimately plausible, however, remains to be seen. If the temporal illusions of this sort are reflected in our actions, then the extensionalist position will fail, however, if, like the cases of spatial illusions, our actions do not reflect the illusory temporal content in our experience, then the extensionalist will have further support for their view. The empirical question remains to be answered.

A second form of illusion that atomists have attempted to use in support of their view are cases of *subjective time dilation*. These sorts of cases are familiar enough outside of any experimental setting. In some cases, perhaps under extreme levels of stress like in a car accident, time will seem to slow down. An event that may have objectively only had a duration of a second may be experienced as lasting several seconds. Take for instance that feeling when you are driving a car during a rainstorm after a long drought you slow down to stop at an intersection and your car begins to hydroplane. The actual duration of your car losing traction may have only been a split second, but as you notice that you've lost control, it may feel as though the whole event actually lasted much longer. Anecdotal evidence like this is fairly common

¹²⁵ Other category specific deficits also occur that point to functional distinctions between the two streams. For instance, damage to the more anterior portions of the ventral pathway typically cause problems with recognition (see the literature on visual agnosias (Farah 1990)). However, these other sorts of functional differences are less relevant to our current discussion.

and can stretch from very short to longer time scales. However, experimental studies of time dilation tend to focus on very short durations.

As was the case for the postdictive phenomena, there are quite a few different sorts of scenarios and experiments that bring about some degree of time dilation, but they are all used in the same way by the atomist. So, for our discussion here, let's focus on *the oddball illusion* (Tse et al 2004).¹²⁶ Subjects, when presented with a train of identical stimuli (same in all regards, including duration and ISI), will perceive an oddball stimulus that only differs in some non-temporal aspect (say, a stimulus that differs in color, size, or pitch) as being of a significantly longer duration than the standard stimuli. In some cases, the dilation of the oddball results in the oddball appearing to have a duration that is as nearly 50% longer than the standard stimuli (Tse et al 2004).

As Lee (2014a) points out, if the mirroring constraint were satisfied, then the experience of the oddball stimulus would have a duration that is significantly longer than the experiences of the standard stimuli. If this is so, then our perception of the world would begin to *lag behind* the actual occurrences out in the world as the experience of the stimuli chain would take longer than the stimuli chain itself. During the course of one block of trials in a psychophysics lab, many such oddball stimuli would be presented, and the lag in perception would accumulate, ultimately leaving our perception of the world far behind the actual occurrences in the world. Of course this is an unacceptable conclusion, and somehow the extensionalist to maintain their position has to deny that this sort of consequence follows from their view.

One approach that extensionalists have taken to block this argument is to argue that corresponding to every subjective time dilation there is a corresponding contraction of other stimuli (Phillips 2014b).¹²⁷ If there is a corresponding time contraction, or at least if the total dilation and total

¹²⁶ Other cases of time dilation include time dilation due to saccades (D. Burr, Tozzi, & Morrone, 2007) and time dilations due to *threatening* stimuli, e.g. stimuli that seem to be moving directly towards the subject as opposed to moving away from the subject (Eagleman 2008; Eagleman & Pariyadath 2009).

¹²⁷ Eagleman & Pariyadath (2009) give an analysis of the oddball illusion that is along these lines. Instead of there being a dilation of the oddball stimulus, Eagleman and colleagues argue that the repeated standard stimuli are

contraction nearly even out, then there won't be any accumulating lag between perception and the world.¹²⁸ Furthermore, the extensionalist could adopt the same approach that they made in response to the existence of postdictive effects. That is, they can say that experience does not drive action, and therefore, we wouldn't find any untoward behavioral consequences of a time lag like this.

In either case, the extensionalist is capable of avoiding the problematic conclusions that the atomist accuses their positions of having to accept. The important point to notice here is that the maneuvers that the extensionalist is appealing to here are not *wild* ones that can be adopted in order to salvage their theory *come what may*. That is, the fact that the extensionalist has plausible responses to the arguments given by the atomist is not merely a result that all theories are underdetermined by the data. Rather, the alternative explanations that the extensionalist adopts have independent plausibility due to other explanatory projects within the cognitive sciences. So, as it stands, the appeal to temporal illusions simply isn't enough to establish the truth of atomism or extensionalism.

The trace integration argument

Geoffrey Lee (2014b) has recently attempted to provide positive arguments for atomism based on general considerations about information processing in the brain. According to him, if the contents of experience are to have any effects on downstream processes (e.g. motor systems, introspection, language, memory, etc), then the mirroring constraint *must* be violated.

actually contracted. However, since Eagleman is not an extensionalist his contraction analysis doesn't engender the parallel problem that arose from time dilations and extensionalism.

¹²⁸ Phillips in fact cites Eagleman and Pariyadath (2009) in support of this proposal. However, Eagleman doesn't claim that the standard stimuli are contracted to make up for the dilation of the oddball, rather, Eagleman argues that just the standard stimuli are contracted, so relative to those stimuli, the oddball appears longer. Phillips cannot simply accept Eagleman's analysis, since it would engender the corresponding problem of perception speeding up when the explanation is tied to the mirroring constraint.

Let's lay out the argument as Lee himself does and then we can go over the steps in a little more detail. The argument as Lee puts it goes as follows (directly quoted from Lee 2014b):

1. In order to be accessible to post-perceptual processes like verbal report, high-level motor control, and domain-general reasoning, temporal information has to be trace-integrated.
2. Unless we have strong contrary evidence, we should assume that the contents of experience are accessible to these post-perceptual processes.
3. Therefore, we should assume that temporal information in consciousness is realized by the *output* of trace-integration.
4. The components of the content of a trace-integrated representation are represented at the same time (or over the same interval).
5. Experiences that have the same realizer timing have the same timing (the temporal correlation principle).
6. Therefore, the different experiential parts of a temporal experience have the same timing.
7. Therefore, temporal experiences are atomic, not process-like.

Perhaps the most central point that Lee rests his argument on is that the contents of experience, unless we have some strong reason to doubt this, are accessible to downstream processes (i.e. they are access-conscious, to use Block's terminology (Block 1995)).¹²⁹ Once we have this on the table, the rest of Lee's argument is supposed to follow from very general principles about *how* information is made accessible in the brain.

Consider, once again, the experience of the flash of lightning as preceding the crash of thunder. If we trace the responses of the sensory system to this sort of worldly event what we find is that at the periphery, say just looking at sensory transducer responses or perhaps at very early levels of cortical

¹²⁹ This allows that there might still be some conscious states that are phenomenally conscious but not accessible. The main point, though, is that if the state is access conscious, then we have reasons to think it is part of the content of experience.

processing, we find that there will be a temporally extended response to these worldly events. In a sense, we can look at this sensory response to the lightning and thunder as *implicitly* representing the temporal relation between the events, as the pattern of activity over time would give an informed and suitably capable observer the information needed to infer the temporal order of the lightning and thunder.

However, it is unclear that the mere existence of this temporally extended pattern of processing *by itself* is enough to make this temporal content available to downstream processes. For downstream processes to be able to make use of the information *the lightning preceded the thunder*, those perceptual processes that are now past, say the experience of the lightning, must leave behind *traces* that can influence later processing. It's only by integrating the traces of the past experiential events with the current events, i.e. the trace of the lightning with the current thunder, that the system can form the representation *the lightning preceded the thunder* which forms the content of our experience.¹³⁰

Since the representation that makes the temporal content of succession available to downstream processes *integrates* the traces of past perceptual processes into a single representation, then that representation will be representing a temporal interval *all at once*. But if it represents a temporal interval *all at once*, then it will not be divisible into discrete temporal parts that correspond to the parts of the represented interval. Therefore, mirroring fails, and with it so does extensionalism.

¹³⁰ In essence, the worry Lee has seems to be a metaphysical one. How can a truly past event, one that has left no trace, come to have any causal influence on current or future events? This is very much the same sort of worry that arose out of the literature coming from Shoemaker's *freezing worlds* thought experiment (Shoemaker 1969). Action at a temporal distance seems to be problematic, or at least, something that many metaphysicians are uncomfortable admitting into their metaphysical picture of the world. However, the metaphysics of causation becomes tricky at this point. For quite some time, analogous claims were made about causation *at a distance*. It was argued that in order for any causal influence to proceed from a cause to its effect there must be a continuous spatial trajectory by which this causal influence could travel. Yet, the Bell experiments are thought to have shown that this sort of spatial constraint needn't be maintained and in fact that there can be causation at a spatial distance. If we take space and time to be one and the same sort of entity, i.e. spacetime, then if we accept that causation at a spatial distance is coherent, then it seems as though we might have to accept that causation at a temporal distance is equally plausible. No matter how odd it might seem.

Lee is correct to notice that by *all at once* he cannot mean *all at one instantaneous moment*. Instead, following Mauk & Buonomano (2004), Lee describes how the brain must provide a *spatial coding* for temporal information. At the very periphery of the sensory system, at the transducers for instance, temporal information is *smeared* across the activity of the system over time (in this way, the transducer activity over time provides an implicit representation of temporal information). However, in order for the brain to make use of this information, the traces of activity must be converted from an implicit coding that is stretched out over time, to an explicit spatial coding that is realized over a relatively short interval (sufficiently short so that the activity of the system can be processed within a single processing bin). This spatial code (which Lee admits is a little bit of misnomer) can be realized by the activity of a spatially distributed set of neurons (see the subsequent discussion of state-dependent timing models) or through the activity of a single neuron over a very short interval. Nevertheless, Lee argues that when the traces of past experiences are integrated, then their integrated form will be such that it cannot be broken down into distinct temporal parts that each correspond to distinct temporal stages of the represented interval.

Lee is correct to insist that whatever type of property of a neural system is encoding some information or content it must be the case that that type of property is capable of causally influencing downstream system. For this reason, we've rejected the mere possession of temporal properties as being viable bearers of content since they seem to be incapable to influencing downstream processes (similarly, we shouldn't think that the color or shape of a neural system is by itself semantically significant for these very same reasons). If some experiential or perceptual event occurred in the past, no matter how near or distant, is to be made available to (or influence) downstream processes, then those past experiential and perceptual events must leave *traces* (whatever those might in fact be) that can directly influence future processing. One way to think of this is to imagine the following. If the experience of the lightning were to have left no trace at all, then to the downstream processes operating at the time of the thunder it would be as if the experience of the lightning never happened.

While Lee is correct about this general point about how information must be encoded there is one significant problem with his argument. It is unclear that we should identify individual experiences with the end state of some process that ultimately influences downstream processes. Lee's focus on cases of local motion (not discussed in this summary) and of succession might actually have clouded this issue. Intuitively, we often speak of *the experience of a sound*, perhaps the sound of a guitarist holding a note. We also speak of the *experience of the sound as having a duration of 1.5 seconds*, the *experience of the sound as having a duration of 1 second*, the *experience of the sound as having a duration of 0.5 seconds*, etc. In a way, we can think of the *experience of the sound as having a duration of .5 seconds* as being a part of the *experience of the sound as having a duration of 1 second*, and those experiences being a part of *the experience of the sound as having a duration of 1.5 seconds*, etc.

Call these experiences of longer durations *containing experiences* as they *contain* the experiences of the shorter durations. For any of these experiences to be accessible to downstream processes, their leading edge (i.e. the largest containing experience) must be capable of exerting all of the causal influence on downstream processes needed for these contents to be accessed. That is, if the containing experience is one of a sound having a duration of 1.5 seconds, then the contained experiences, e.g. of the sound having a duration of 1 second, will not exert any causal influence on downstream processes over and above those causal influences that the containing experience can exert. This, in essence, is just the trace integration argument as it would apply to the experience of duration.

Now, when we are evaluating the mirroring constraint we need to ask, what experience do we want to analyze? The experience of the tone? The experience of the tone as lasting 1.5 seconds? Or do we have to analyze the experience of the shortest duration that we may be capable of discerning, which in the case of audition is around 10ms (Rammsayer et al, 2015)? We obviously can't restrict ourselves to analyzing experiences that are carved down to representations of instantaneous moments, since we are analyzing *temporal experience* and these involve the experience of temporal extended intervals, so we

should be focusing on one of the options we just mentioned. We also likely shouldn't consider "experiences" that are too long, like the experience of growing up in the 80's. While we might call this *an experience* there really isn't any *single* content that we can comprehend at a moment to analyze – that is, this sort of experience is in a way not about anything specific, but rather is an agglomeration of a great many experiences or sometimes just picks out a judgment or belief about what the 80's were like.

The same worries that arose for considering the experience of growing up in the 80's also arises for something like considering the experience of the tone. The phrase 'the experience of the tone' merely helps us pick out the experience that we're interested in, but to properly analyze whether the mirroring constraint is satisfied we need to make this experience precise by specifying the content. That is, if we experienced the tone veridically and the tone lasted 1.5 seconds, then we should analyze the specific experience *of the tone lasting 1.5 seconds*. However, we should be reluctant to identify this experience with the final instant (or state) of the process that brought about its realization. As mentioned by Lee, a common criterion for a content being part of experience is that the content be accessible to downstream processes. However, instantaneous states aren't consciously accessible (as Lee also points out) since neural systems are not sensitive to instantaneous rates, but firing patterns over temporally extended windows. So, we need to analyze the process over a short window of time in virtue of which some experiential content becomes available. In other words, we shouldn't look at the state, but rather at the process by which information is made available.¹³¹

Once we are looking at the process by which some content is made available, then we can ask the question of whether that process satisfies the mirroring constraint. In the case of duration, when we analyze a particular duration experience, like one of *the tone lasting 1.5 seconds*, we need to ask whether

¹³¹ In a sense, this isn't much of a departure for the state-based analysis of experience. A state of a system is just a state of affairs in which the system instantiates some property. Moving to the process view, in the way that I am advocating here, is simply more of the same. The particular process that a system is engaged in is just another state of affairs in which the particular system is instantiating some property (or properties).

the process that brings this about has a temporal structure that mirrors the temporal structure of the experiences contained within the target experience. That is, does the temporal structure of the realizer of the 1.5 second experience have temporal stages that correspond to the duration experiences contained within the experience of 1.5 seconds (i.e. experience of 1.2 seconds, of 1 second, etc)? If the system represents a duration through some process that itself can be broken down into distinct temporal stages that correspond to the contained experiences, then we will have found a counterexample to Lee's argument. The result would simply be that his argument fails to be valid. There would be experiences, the leading edge of which combines any causal influence from previous stages, and thus involves trace integration, but the actual process by which the experience's realizer came to be in this state will have had a temporal structure that mirrors its content.

The argument against Lee as I've put it here is really only so far an argument that *something like this could be the case*. However, in the next section, we will begin to sketch out the empirical evidence that shows that something like this is actually the case. The process by which an experience of duration is brought about has temporal stages that are themselves semantically significant in that they represent the temporal stages that correspond to the contained duration experiences. Once again, as was the case with the other arguments for atomism and extensionalism we see that in order to assess how the temporal structure of experience relates to the temporal content of experience we need to appeal to empirical work on the very mechanisms that carry temporal information in perception. Nothing about introspection will do.

4.5. The fragmentary model of temporal perception

We've seen the arguments that people have put forward in favor of atomism and extensionalism and have seen that none of the arguments succeed in establishing their intended conclusions. As it stands, we're still in a position where we do not know how the temporal structure of experience relates to the

temporal content of experience, and in so far as figuring out this relation was important for understanding how experience comes to represent the temporal structure of our world, we are still in a position where we do not know how experience comes to have its temporal content.

In what follows we'll begin to see further that the question of whether the mirroring constraint is satisfied or not truly has very little bearing on how experience comes to have its temporal content. The primary reason for this is that temporal experience is not a single unified phenomenon, but is instead, a motley assortment of different capacities underpinned by different mechanisms that all come to carry information about and represent distinct aspects of the temporal structure of our world. According to this *fragmentary model of temporal perception*, we shouldn't expect there to be any single answer as to whether or not the representation of some temporal property in perception is mirrored by the temporal structure of the system doing the representing since there is a multitude of different representational systems that employ distinct representational strategies to keep track of and represent time.

In this section, I will argue for the fragmentary model of temporal perception. The first step in arguing for this view is to establish that the perception of time is due to a multitude of different mechanisms – i.e. *the multiple mechanisms claim*. Once the multiple mechanisms claim is established the next step is to look at the specific details of some of the mechanisms included in the fragmentary model. It is by looking at these mechanistic details that we will see that in some cases the mechanisms involved will satisfy the mirroring constraint while other mechanisms will not satisfy the mirroring constraint. As a result of the fragmentary model we can see that much of the philosophical literature on temporal experience, that has by and large been focused on debates over the mirroring constraint, are simply misguided. They took the debate to be an either / or debate (one or the other but not both), but in fact, the situation is closer to an either / or / or both sort of debate.¹³²

¹³² This way of putting the point is thanks to Kenneth Williford who commented on a conference presentation of this material at the 2015 meeting of the Southern Society for Philosophy and Psychology.

The multiple mechanisms claim

If any word accurately describes the literature on temporal perception it would be ‘complicated’. Researchers from a number of different areas in the cognitive sciences have been studying temporal perception for decades and each field brings with it a host of different explanatory aims, methodological tools, and historical baggage. That so many different areas of the cognitive sciences have shown interest in temporal perception shouldn’t be surprising. The temporality of our conscious life is such a striking aspect of our consciousness that any field that wants to understand what experience is like (or what it is or what explains consciousness) will have to take on temporal experience as a phenomenon to be explained. But even beyond just how things appear to us in experience, the perception of time, and the mental representation of time more generally, is vital to a wide range of tasks. From phoneme detection, word parsing, sensori-motor interactions, memory, narrative comprehension, musical appreciation, and so on, the mental representation of time is central to the phenomena that areas of the cognitive sciences study.

With all of this research one thing that might come as a surprise is the general *lack of consensus* on *what* brings our capacity for temporal perception about. Importantly, this disagreement isn’t over the more general question of how we might have any experience whatsoever, and therefore isn’t simply the result of the long standing puzzlement facing the *hard problem of consciousness*¹³³. Rather, the disagreement in the literature is simply over what sorts of mechanisms and processes are involved in the way we keep track of time. Without any consensus we find that theorists are positing a wide range of

¹³³ David Chalmers (1996) notoriously distinguished between two types of problems that arise when we study consciousness. The *easy problems of consciousness* are those involved in explaining the functional characteristics of consciousness such as how do manipulations of attention influence perception, how do we parse the visual field into objects and their features, etc. The *hard problem*, on the other hand, arises when we want to understand *why there is anything at all that it is like in having a mental life*. That is, the hard problem concerns the very fact that we are conscious beings with a phenomenal mental life as opposed to being mere zombies, creatures that perform complex information processing tasks without any conscious experience.

different models for how we perceive time, and these theorists are taking themselves to be positing models that are at odds with one another.

While many of these cases of seeming conflict are in fact cases of theorists proposing models that are genuinely in conflict with one another, many of these cases involve the positing of mechanisms that only *appear* to be in conflict but are actually perfectly compatible. The conflict generally arises when theorists take themselves to be looking for *the mechanism by which we perceive and experience time*. As Hartcher-O'Brien and colleagues put it (2016), the search for a single time keeping mechanism is motivated by the fact that in the world and in our experience the various temporal properties of the world (e.g. durations and rates) are intimately tied up with one another. The temporal structure of the world that we perceive through various sensory systems and across various time scales all seems to fit nicely together in way that structured the rest of the non-temporal features of the world we perceive.¹³⁴

Perhaps the most influential single mechanism model of temporal perception that's been developed has been the accumulator-pacemaker model originally developed by Treisman (1963) and subsequently developed by a number of authors, most notably with Gibbon's *scalar expectancy theory* (SET) (1977). According to SET, the perception of time is brought about by a three-part system. There is a central clock mechanism that operates through a pacemaker-accumulator process (see chapter 2 for a discussion), a memory store where reference memories are stored (i.e. memories for the durations of particular types of stimuli), and a decision mechanism that compares the output of the clock to temporal reference memories stored in memory to determine the duration of the stimuli being perceived.

While Gibbon's SET has been highly influential, like many influential theories in the cognitive sciences, many researchers have begun to reject parts or the entirety of the view.¹³⁵ Perhaps most salient is

¹³⁴ On this point, see the discussion of Matthen's *integration argument* (2014) in chapter 3 and the discussion of temporal integration in chapter 5.

¹³⁵ In some way, Gibbon's SET is to temporal representation as Anne Treisman's *feature integration theory* is to visual search (although the influence and departure from FIT is more obvious and clear cut than it is for SET). Both

the fact that SET makes a prediction about there being a single clock-like mechanism involved in perception, but no neurological evidence has been found to back this up (Eagleman & Pariyadath 2009). Similarly, given the wide range of time scales that we perceive time over – from as short as 10ms in audition to much longer time scales involved in event perception – it’s unlikely that one pacemaker-accumulator system would be capable of providing the temporal resolution needed for these various time scales (Butts et al 2007; Panzeri et al 2014). Furthermore, as we’ll discuss in more detail below, a number of pharmacological and mechanical interventions show that the abilities to perceive time within certain time scales and within certain modalities are independent of one another. Now, the original SET that Gibbon developed is capable of being modified to account for this various data, by modifying the sub-components of the model, however, as Wearden (1999) has pointed out, this flexibility is both a blessing and a curse. The blessing arises in that SET seems to have significant explanatory power, but this comes at the cost of becoming seemingly unfalsifiable as the parts are capable of being changed too much.

While many in the literature have rejected SET, and its offspring, the alternative proposals still seem to adopt a *one or the other* approach to finding *the* mechanism or *the* explanation of our capacity for temporal perception. Some examples of *single mechanisms* for the perception of time include revisions of SET (Hinton & Meck 1997; Matthews et al 2011), the *striatal-beat-frequency model* (Matell & Meck 2000, 2004), and activity in the default mode network (Carvalho et al 2016; Lloyd 2012). Some theories that have emerged after SET, however, do not simply appeal to a *single token mechanism* but instead appeal to a *single type of explanation*. According to these researchers, timing isn’t the product of a single mechanism but is rather the product of a single property of neural processing that can be found throughout the brain. Some examples of this include appeals to *efficient coding strategies* (Eagleman &

theories provided a concrete framework for the study of some psychological capacity that made clear predictions about the neurological and psychological processes involved in perception (more so in the case of FIT). And yet, researchers have found that the original models have problems, yet many of the resulting models retain quite a significant amount of the original view.

Pariyadath 2009) and *state-dependent network properties* (Buonomano 2000; Finnerty et al 2015; Ivry & Schlerf 2008; Karmarkar & Buonomano 2007).

It's here where we confront a division in the empirical literature between appeals to *extrinsic timing mechanisms* and *intrinsic timing mechanisms*. Extrinsic timing models appeal to dedicated time keeping mechanisms. There are those systems in the brain that process (or predominantly process) non-temporal aspects of our world, and then there are those systems that process temporal information. With these extrinsic models, the temporal features of stimuli must be actively bound to the non-temporal objects in the world. Intrinsic timing models, on the other hand, explain timing through the operation of some feature of neural systems or neural processes that are ubiquitous in the brain. These models have no need to appeal to any dedicating timing mechanisms.

It's important to not take this division to be one between single mechanism accounts and multiple mechanism accounts, as an extrinsic timing model could in principle appeal to multiple dedicated timing mechanisms (perhaps one for each sensory system). But, nevertheless, when we look at the debate between these two different approaches to timing models, we still find a *one or the other* approach.¹³⁶

One practical difficulty in adjudicating between these different models is that often they are developed on the basis of, sometimes slightly and sometimes largely, different experimental protocols that involve different tasks, time scales, and different sorts of temporal properties (i.e. durations versus orderings versus rhythms etc). Research that points to the role of the cerebellum in temporal perception often involves tasks that require the fine temporally tuning of non-cyclical action (e.g. the cerebellum is not involved in the reproduction of rhythmic behaviors (Breska & Ivry, 2016)). Research that implicates the default mode network often involves judgments about time over a longer time scale, several seconds in length, and across multiple sensory modalities (Lloyd 2012). The efficiency coding models (Eagleman

¹³⁶ One notable exception to this *one or another* approach to the debate over intrinsic and extrinsic timing models is found in (Maniadas & Trahanias 2016) who propose a computational model that incorporates components of both sorts of models.

& Pariyadath 2009) and state-dependent network models (Buonomano 2000; Finnerty et al 2015; Ivry & Schlerf 2008; Karmarkar & Buonomano 2007) look at the perceptual judgments made about the duration and sequence of events in the millisecond time range (30ms – 750ms) within individual modalities. Lastly, the revised versions of SET often employ experimental conditions that extend well-beyond 1-second in length (e.g. Zakay & Block, 1997). Similarly, in all of these approaches, it is not rare to find differences between prospective and retrospective timing tasks in which subjects are either told to keep track of the timing of some event before the trial begins or to report on the timing of some event after the trial has occurred.

One rather clean explanation of this disagreement and lack of consensus in the literature is that there simply is a fact to the matter of how we come to perceive time and these different proposals are in fact in genuine disagreement with one another. In other words, there is a single phenomenon to be explained, and these models disagree on how to explain that phenomenon. Another explanation of the disagreement, however, is to simply deny that we have anything like a single unified phenomenon on our hands that needs explanation. Instead, we perceive time at a number of different time scales, through various sensory modalities, and even the very temporal properties that we perceive differ (e.g. durations as opposed to sequences and orderings as opposed to rhythms and motion and so on). Each of these different aspects can be dealt with on their own and how we explain one aspect of our capacity to perceive time needn't prohibit any restrict how we explain other aspects of our capacity to perceive time.

In fact, some in the empirical literature have started recently to adopt this sort of approach. As Wittman and van Wassenhove (2009) put it:

In recent years, new ideas have emerged regarding the neurobiological mechanisms underlying the experience of time. In particular, major progress has been made with the realization that time perception may engage distinct brain mechanisms (and areas) depending on the time scale at which events occur...

(Wittmann & Wassenhove 2009, p. 1809)

The main divide that Wittman and van Wassenhove point to exists between those mechanisms involved in sub-second temporal representation and supra-second temporal representation. Similarly, in philosophy, recently Rick Grush (2016) has argued that there is a division in the semantics of temporal representations below and above 1 second in duration. Below one second, temporal representations represent B-theoretic properties (e.g. *earlier than*, *later than*, and *simultaneous with*), while above one second, temporal representations represent A-theoretic properties (e.g. *past*, *present*, and *future*).

Part of the argument that Wittman and van Wassenhove give for this divide involves noting the ways that different mechanisms contribute to sub vs supra second timing. Suprasecond timing is influenced by such factors as attention, emotion, and bodily movement. In contrast to the promiscuity that is found for suprasecond timing, subsecond timing seems to be much more recalcitrant to the influence of these systems (Wittmann 2009).

This sort of approach, however, in pointing to the differential contribution of systems in suprasecond timing as opposed to subsecond timing has its limitations. One could explain the contribution of various systems in suprasecond timing by simply noting that in these sorts of timing tasks there is more time for these auxiliary systems to make contributions. Whereas in the subsecond timing these other mechanisms do not have the time to properly engage. Note, however, that this sort of explanation of the differential contribution of these various systems does not require that we posit distinct timing mechanisms. It's only that we find behavioral / functional differences between sub-second and supra-second timing because these different sorts of tasks allow for a single mechanism to make use of different contributing systems.

Instead, the primary argument for the multiple mechanisms claim of the fragmentary model comes from the existence of certain dissociations between timing capacities and the lack of any unilateral timing deficits. Let's take these in turn.

First, it is well known that various psychopharmacological drugs can selectively impair select timing capacities. Rammsayer (1999) showed that haloperidol, a dopamine receptor antagonist used in the treatment of schizophrenia, and midazolam, a benzodiazepine used in the treatment of seizures and insomnia, both impair the discrimination of durations of approximately 1 second. However, only haloperidol impaired the discrimination of durations at much shorter time scales of 50ms. Similar effects have been found with the hallucinogen psilocybin (Wittmann et al 2007) in that the drug seems to leave the discrimination of temporal properties below 2 seconds untouched but severely impairs the perception of duration at durations of 2 seconds and longer.

In addition to the psychopharmacological interventions on the capacity to perceive time at various time scales, mechanical interventions provide similar results. The application of *rapid transcranial magnetic stimulation* (rTMS) to different specific brain regions has differential effects on specific timing capacities. Koch et al (2007) showed that the application of rTMS to dorsolateral frontal cortex impaired timing the seconds range whereas Jones et al (Jones et al, 2004) showed that application of rTMS to the cerebellum impaired timing in the millisecond range.

To further point out the influence of multiple mechanisms in temporal perception, various medical disorders seem to come hand in hand with selective impairments in abilities to perceive time. Schizophrenia is known to be correlated with difficulties in the integration of temporal information across sensory modalities. In normal subjects, slight temporal asynchronies in the onset of stimuli can often be discounted in the perception of simultaneity. For instance, our perceptual systems are capable of tolerating an asynchrony of 250ms in the timing of audio and visual speech signals that we still nevertheless perceive as occurring simultaneously (Dixon & Spitz 1980). This window has been called *the temporal window of integration* (TWI). Schizophrenics are known to have an extended TWI in that they are capable of tolerating larger asynchronies in the perception of simultaneity (Martin et al 2013).

People with autism spectrum disorder actually seem to have an opposite problem in that they have difficulty binding the multimodal temporal information into a unified perception (Stevenson et al., 2014).

The flipside of all of these specific and selective deficits in some aspect of our capacity to perceive time is that there seems to be no global deficits in temporal processing. That is, there is no known finding of a person that completely lacks the ability to perceive time. Even patients with severe anterograde amnesia, while they may be unable to retain any information that is consciously accessible from 1 second to the next, nevertheless retain the capacity to understand that they cannot remember what just happened, and in this understanding they exhibit some ability to mentally represent time (see the discussion of Mr. B in Craver et al (2014)) Furthermore, these patients with severe anterograde amnesia are still capable of completing sensori-motor tasks that require the capacity to properly time their behaviors (ibid).

Another purely empirical point in favor of the multiple mechanisms model is the existence of modality specific temporal illusions. Burr et al (2007) showed that saccades compress the perception of time for visual stimuli at the location of the saccade while leaving untouched the perception of duration for stimuli at different locations in the visual field and for stimuli in different modalities.

Direct empirical evidence aside, I believe we can give some theoretical reasons for doubting that there would be a single mechanism that brings about our capacity for temporal perception. In particular, I think there are evolutionary reasons for doubting this. Consider the importance of properly coordinating one's actions with the temporal structure of the events in the environment. The ability to synchronize one's behavior in this way is crucial to almost all (if not all) actions that an organism engages in. In other words, without the ability to properly coordinate action with the temporal structure of events in the environment, animal behavior would fail to be adaptive, and the result would be that the animal would quickly die. The ability to keep track of time is not an ability that can simply be painted on top of prior existing sensori-motor abilities that direct the organism's behaviors in the environment.

This marks an important difference between the capacity to perceive time and the capacity, for instance, to perceive color. Color vision developed as an *additional* visual capacity that provided animals with the hardware to use spectral information to detect contrasts and thereby detect objects in the world that could not be detected reliably on the basis of achromatic vision (Jacobs 2009). But, and this is important to emphasize, the development of color vision came on top of a prior working visual system. Color vision simply extended the system's capacities to provide organisms with more tools to properly navigate their environments.

The perception of time cannot be like this. The ability to keep track of time is not something that can simply be tacked onto a pre-existing cognitive architecture. The ancestral animals that lacked any ability to keep track of time would have simply died off. Instead, timing is vital to all tasks that an animal may perform, and as such, we should expect that with the development of each sensori-motor capacity there would be some timing capacity that is built into that system. Now, this isn't to deny the usefulness of some dedicated and centralized time keeping mechanism, in fact much of this initial temporal sensitivity could possibly be explained through cued synchronization models, but the fact that timing is so ubiquitous in our behaviors pushes us to think that timing capacities are likely ubiquitous throughout the brain.

These various cases of selective impairments of aspects of our perception of time strongly point to a model according to which temporal perception (and temporal experience) are brought about by distinct mechanisms. That is, the idea that temporal perception is due to some unified mechanism, and it is this unified mechanism that explains the seeming unity of time as it is perceived, simply cannot be maintained. Furthermore, the evolutionary importance of timing makes it implausible that timing could have developed independently of other sensori-motor capacities. Now, all this shows is that there are multiple mechanisms involved in temporal perception. In this way, we are in a position to notice that *there may not be* a single unified response to whether or not *temporal experience* satisfied the mirroring

constraint. Temporal experience simply isn't the single unified phenomena that would compel us to expect such a unified response. Instead, the behavioral and neurological differences that appear in the various manipulations of temporal perception indicate that there isn't a unified phenomenon here (except in the minimal sense that we can consider any assortment of phenomena as a unified one by stipulating that a complex collection of distinct phenomena should be categorized under a single banner).¹³⁷ In the next section, we will look at specific models that people have employed to account for the various capacities for perceiving aspects of the temporal structure of our world and we will see that only some of these models satisfy the mirroring constraint while others do not.

Two cases: against the atomism / extensionalism debate

With the multiple mechanism claim on the table, we're now in a position to evaluate the complex relationship between the temporal content of experience and the temporal structure of experience itself. We can restrict our discussion to models of two distinct types of mechanisms that are involved in the perception of time at very short time scales, from between 50ms to 750ms, since by looking at just these types of mechanisms we'll have sufficient reasons for the conclusion that there is no single answer to whether or not the mirroring constraint is satisfied. The first type of mechanism is involved in the perception of duration and, to a lesser extent, succession within individual modalities. The second type of mechanism is involved in the perception of temporal order across different modalities. It is by looking at these two mechanisms that we'll find counterexamples to *both* atomism and extensionalism. In particular, the mechanisms involved in the unimodal perception of duration and succession at short time scales

¹³⁷ The situation is similar to the phenomenon of *the cause of people eating oatmeal on Tuesdays*. We wouldn't think that the eating of oatmeal on Tuesdays picks out a unified biologically significant, or psychologically significant, phenomenon. People might eat oatmeal on Tuesdays for a variety of reasons. But, of course, we can force the unity of the phenomena by taking a conjunction of all the various causes. Still in this case it seems as though there isn't any really compelling reason to think that we've discovered a unified phenomenon out in the world.

satisfy the mirroring constraint, while the mechanisms involved in the crossmodal perception of temporal order violate the mirroring constraint. As a result, neither atomism nor extensionalism provide an adequate characterization of our temporal experience in general, and the debate over the mirroring constraint is shown to not solve the original puzzle of how it is that we can experience the temporally structured world around us.

The interest of focusing on these two systems comes from two main considerations. First, the time scale. By restricting ourselves to very short time scales we make sure (or at least, make it more likely) that the processes that we are discussing are purely perceptual. As mentioned, attention is not involved in temporal perception at this time scale, and often pre-attentive processes are taken to be sensory and not cognitive. Second, considering these two types of mechanisms provide us with an interesting cross section of our overall temporal experience. Introspectively, when we notice the temporal structure of our world, it generally makes no difference which modality is detecting the temporal structure since it seems as though the temporal structure of our world as its presented to us through the various sensory systems is unified. By looking at the mechanisms that are involved in unimodal and crossmodal temporal perception we get a more complete sense of the various ways that we may represent time.

Before we actually begin looking at these different cases, let's lay out a quick methodological point. The debate that we are focusing on is one over the mirroring constraint and in order to decipher the temporal structure of an experience we have to look at the temporal structure of the realizer of that experience. In order, then, to decide whether any aspect of temporal experience satisfies the mirroring constraint we must first determine that the content of the representation is as well as what the temporal structure of the mechanism or process that realizes the experience. Only with those two points in hand can we then turn to look at how the contents of the representations are encoded by the representational vehicles – that is, what properties of the mechanisms are semantically significant. Once we have all of

this figured out, then we will finally be in a position to judge whether an experience satisfies or fails to satisfy the mirroring constraint.

Case #1: unimodal perception of duration and succession

The dominant theories of the perception of time, within and across modalities, have often appealed to extrinsic timing mechanisms. For example, the pacemaker-accumulator component of SET is supposed to operate independently of the individual sensory systems that process and represent the non-temporal aspects of the world. It is only by a process whereby the pacemaker-accumulator system tracks the dynamics of the first-order sensory areas that we come to represent the world around us as having particular temporal properties.

Recently, however, several researchers have started to point to the possibility of an alternative to extrinsic timing models. Instead, they claim that temporal information can be processed and represented by the *intrinsic* properties of the neural populations that encode for the non-temporal aspects of the environment (Buonomano, 2000, 2005; Burr et al., 2007; Finnerty et al., 2015; Ivry & Schlerf, 2008; Lebedev et al., 2008; Mauk & Buonomano, 2004; Reutimann et al., 2004). So, according to this approach to explaining temporal perception, time keeping is explained by appealing to the properties of neural systems that are engaged in the representation of some aspect of the external world and as a result there is no need for a dedicated time keeping mechanism. Typically, these time keeping properties appealed to in intrinsic timing models are a result of the time dependent behavior of neural systems in response to external stimulation.

Some intrinsic timing models see timing as being an intrinsic property of *particular* neural systems. For example, Lebedev et al (2008) see timing as the result of ramping activity in pre-motor areas. However, their experiments were restricted to action guidance in rats, so it's unclear from their

work whether they would extend their explanation to other timing capacities. However, other theorists see intrinsic timing mechanisms as a much more ubiquitous feature of the brain. It is in looking at these more ubiquitous models that we'll find our counterexample to atomism.¹³⁸

The specific sort of model that I want to discuss is the *state dependent model* developed most saliently by Karmarkar & Buonomano (2007). According to this model, timing is brought about by the intrinsic, time dependent and state dependent properties of neural networks. The basic idea that guides how these models explain temporal representation can be helpfully illustrated through an analogy (ibid). Imagine a still pool of water – perhaps a barrel full of water. Now, drop a pebble into that water. The water's surface will begin to show ripples that radiate out from the pebble's impact on the water's surface. Provided that you have information about the size and force of the pebble that hit the water, you (i.e. an idealized observer) can actually use the current state of the ripples on the water's surface to tell how much time has passed since the impact of the pebble on the water's surface. The reason why you can use the ripples to do this is that the way in which the ripples propagate across the surface of the water is a time-dependent process. Now, if we drop another pebble into the already perturbed water, the second pebble will cause a new pattern of ripples to propagate across the surface of the water. Again, though, the ripples after the dropping of this second pebble can also be used to tell how much time has passed since the second pebble, however, in this case the temporal information carried by the spatial distribution of ripples will be a function of not only the properties of water, the pebble, and how the pebble was dropped, but it will also be a function of the prior state of the water. In other words, the properties of the water's surface that carry information about time will be dependent on the prior state of the system. As Goel and Buonomano put it, “ripples thus establish a short-lasting and dynamic memory of the recent stimulus

¹³⁸ However, even by looking at the ramping activation models of Lebedev and colleagues the same argument against atomism would succeed as the realizers of temporal experiences according to this model can be divided into temporal parts that are semantically evaluable.

history of the liquid, and it is possible to estimate the amount of time elapsed based on the current state of the liquid.” (Goel & Buonomano 2014).

State-dependent intrinsic timing models appeal to a similar rippling effect that has been found in artificial neural networks designed with biologically plausible constraints and even in in vitro neural populations (Finnerty et al 2015). The “rippling” in these systems is understood as a spatial pattern of activation that evolves over time. To illustrate, the picture is something like the following. Take a neural network. In characterizing the state of a neural network we must appeal to both *active* and *hidden* states of the network (Buonomano & Maass 2009; Goel & Buonomano 2014). The active states of the network are those neurons in the network that are currently firing and their firing rates. The hidden states of the network, however, are states that are not directly observable through extra-cellular means (Goel & Buonomano 2014). These “hidden” states include slow synaptic currents, channel kinetics, and short-term plasticity (STP)¹³⁹. What these hidden states do is provide an explanation for the dynamic changes in the active responses to incoming sensory stimulation. In particular, these hidden states of neural networks provide a time-dependent means of altering the responsiveness of individual neurons. The effect of these local hidden network properties is that at any given moment, the responsiveness of the individual neurons in the network will be a time-dependent function of these hidden states. In this way, when an incoming signal excited the network, at any given moment only a select subset of the network’s neurons will be responsive to the incoming signal. In this way, the active states of the network will produce a sort of “rippling” of activation in that at any given moment after stimulation there will be a spatially distinct subset of neurons that are firing. In essence, what gets produced by these hidden states is a network with time sensitive cells.

¹³⁹ STP plays a large role in the explanation of these state-dependent networks. For a nice discussion of STP see Zucker (1989).

Goel & Buonomano (2014) survey a number of examples in which time selective neurons have been found in wide variety of animals as well as how neural networks can be shown to keep track of time in vitro. That is, since neural networks are capable of keeping track of time in vitro, i.e. in the absence of any dedicating timing mechanisms, it must be the case that the capacity for timing that these systems exhibit is due to the intrinsic time dependent properties of neural systems themselves. Now the fact that neural systems exhibit some time-dependent properties that as a matter of fact carry information about time does not by itself show that it is these properties that are involved in actually representing time to the organism.¹⁴⁰ Recall that for a neural system to be representing some state of the world it must be the case that the proposed properties that carry information about the world are such that they can causally influence down stream processes.

As was described in chapter 2, one means of determining whether some properties of a representation are the semantically significant properties is to observe the behavioral effects of that system in order to see whether or not there are any novel behavioral / functional consequences of the particular properties in question. In other words, test to see if there are any expected oddities that would arise because of the particular encoding properties.¹⁴¹ In this particular case, the state-dependent nature of these network properties has the consequence that the representation of temporal properties should be *context dependent* in that the operation of the system at any moment is influenced by the immediate history of the system's activity. This very sort of context dependency in temporal perception has in fact been confirmed for very short time scales of around 100ms (Karmarkar & Buonomano 2007). The

¹⁴⁰ The state of the brain may carry information about a whole host of different things in a person's history that we would not take to be part of the content of any mental states. For instance, certain environmental conditions leave somewhat telling neural damage (e.g. carbon monoxide poisoning, syphilis, lead poisoning, etc) however that the brain carries information about these past events or conditions is not something that can be found through a reading off of the mental content of the individual. Instead, this is something that the state of the brain carries information about but this information isn't accessed by any consumer systems.

¹⁴¹ This is exactly what has driven many cognitive scientists working on general magnitude representations (e.g. space, time, and number representations) to posit the existence of a *mental number line* (e.g. Carey (2009)). The existence of a magnitude representation with that sort of semantically significant structure is supposed to explain a whole host of behavioral data. See Beck (2014) and Carey (2009) for reviews.

accuracy of temporal judgments of a tone's duration was directly influenced by whether or not the target tone was presented within a fixed or variable context (i.e. whether the temporal delay between a target tone and a lead-in tone was fixed or whether it was varied). The findings of Karmarkar & Buonomano were further confirmed by Ivry et al (unpublished work cited in (Ivry & Schlerf 2008)). However, when the duration of the target tone was increased beyond 300ms, the context effects were lost, indicating that at this time scale, duration judgments may no longer be the product of intrinsic timing mechanisms but might rely on other systems.

Now, a model like SET *could* in principle explain these context dependent effects, but as noted earlier, this is a blessing and a curse for SET. The ability of SET to explain findings of this sort, and in fact, findings of most sorts, comes from the fact that very little is actually known about the various components that make up SET (i.e. the pacemaker-accumulator system, the gating mechanism, the role and sort of memory involved, and the decision processes that compare the online pacemaker-accumulator readings with the stored durations in memory). So, if the various components of SET aren't well understood, then we can simply change how they behave in various ways to accommodate the data. Perhaps the context effects at short time scales emerge from some problem in controlling the gating mechanism that mediates between the pacemaker and the accumulator in a way that the opening and closing of the gate no longer precisely corresponds to the onset / offset of stimulation. This is of course something that could be further studied, but the interesting thing about the state-dependent timing models is that they appeal to features of neural networks that we have independent reasons for believing in, and these features are enough to explain the presence of the context effects. In other words, the intrinsic timing models provide us with a parsimonious explanation of the context effects in that they do not require the positing, or changing, of mechanisms solely for the explanation of these effects.

Now, the interest of discussing this model is to ultimately say something about the mirroring constraint and the atomism / extensionalism debate. To do so, let's consider a concrete case. Consider the

experience of hearing a tone that lasts 80ms (a duration that is well above the threshold for temporal discrimination in audition and that also falls well within the time scale that state-dependent models apply to). The experience that we're analyzing is the experience of the tone. At the end of the 80ms tone, we will have an experience with the content *that tone as having a duration of 80ms*. Recall from our discussion of the trace integration argument, that an experience of a tone as having a duration of 80ms will be a containing experience provided that we can discriminate the shorter durations that are contained by the duration of 80ms. That is, our experience of the tone as lasting 80ms will also present the tone as having had a duration of 70ms, 60ms, etc. So we have to ask, does the process that brings about the experience of the tone as lasting 80ms have a temporal structure such that the temporal parts of the process, i.e. of the experience's realizer, correspond to the contained experiences?

If we look at the way that the proposed state-dependent time keeping mechanisms come to represent time, it is through a process that unfolds over time such that the state of the system at one moment crucially depends on the state of the system at a prior moment and the overall process that gives rise to the experience will have these earlier stages as proper temporal parts (this is just how the hidden states of the network are fixed in order to modulate the active states of the system). In this way, the states of the system leading up to the experience of the tone as having a duration of 80ms will themselves be part of the process that realizes that experience. That is, the immediately preceding stages of the process are part of the process that realizes the experience. If we take the temporal stage¹⁴² prior to the final stage in which the tone is represented as having a duration of 80ms, then we can ask whether that stage itself is semantically interpretable as representing a temporally distinct stage of the overall temporal structure represented by the experience – i.e. does mirroring hold? What we find is that yes it does. The realizer of the experience of the tone as lasting 80ms itself can be divided into distinct temporal stages (i.e. preceding states of the time-dependent evolution of the system), and these distinct temporal stages

¹⁴² How long each individual stage should be is an empirical matter. Whatever the answer, however, the point being made here would remain the same.

themselves represent the *contained* experiences of the tone lasting 70ms, 60ms, etc. In other words, as the system evolves into its state where it represents the tone as lasting 80ms, it will proceed through stages that represent the tone as having had a shorter duration. It is these earlier stages that correspond to the contained experiences within the overall experience of the tone lasting 80ms.

Here we have a case in which mirroring holds, but not because of any necessary conditions about the nature of experience in general or due to some adoption of a resemblance theory of content. Instead, mirroring holds and extensionalism is vindicated for experiences of this sort due to the way in which the mechanisms of the brain operate. Given that we have some reasons to accept this sort of model for temporal representation within individual modalities at very short time scales, then we have reasons to believe that our experiences of this aspect of our temporal world will satisfy the mirroring constraint.

Case #2: multimodal perception of order and simultaneity

By considering how the brain comes to perceive durations at very short time scales within individual modalities we uncovered a counterexample to atomism. The mechanisms involved in that aspect of our temporal experience satisfy the mirroring constraint. It is time to consider another aspect of our temporal experience, the perception of temporal order across modalities at short time scales, and it is by looking at this aspect of our temporal experience that we'll see that here we have a violation of mirroring, and thereby a violation of extensionalism. As a result of these two cases the idea that we could give an univocal answer to whether or not temporal experience satisfies the mirroring constraint will be shown to be untenable. There simply is no single means by which the temporal content of experience relates to the temporal structure of experience itself.

A striking aspect of our temporal experience is the fact that the temporality of our world seems to unify the deliverances of our various sensory systems. For instance, when we experience the flash of

lightning as preceding the crash of thunder we are perceiving a temporal relation between events that simply cannot be captured by a mechanism that operates within an individual modality. As with the discussion of the previous case, the empirical investigation into the mechanisms that underpin this capacity is still in its infancy. However, we know enough about the perception of temporal order across modalities to make a claim that suffices to show that we can find a counter example to the mirroring constraint.

In the next chapter, we'll discuss in more detail the demands that are required to produce the unified perception of time across modalities and across time scales, but in this chapter we can restrict our discussion to those mechanisms involved in forming *temporal order judgments* (TOJ) at time scales of less than one second. Let's begin by surveying some data about TOJ in general.

The first thing to notice is that the perception of simultaneity / non-simultaneity comes apart from the perception of temporal order (i.e. perceiving one event as being earlier than or later than another event). The *simultaneity threshold* (ST), the temporal window within which stimuli must be presented to be perceived as simultaneous, differs between each sensory system. Audition has by far the lowest ST of any sensory system at around 2-5ms. Tones presented with an offset of less than this threshold are fused together and perceived as being either one tone or as simultaneous (depending on the pitch of the two tones). However, if stimuli are presented at this threshold, then they will be perceived as not being simultaneous but their temporal order will not be detected. Touch comes in next with a ST at around 10ms. Finally, vision comes in much slower with an ST of about 20-30ms.¹⁴³

While it is known that there is some variation in ST both across individuals and within individuals as a result of age or stimulus properties (e.g. the loudness of an auditory stimulus can influence the ST) (Poppel 1988), the important thing to note is that a pair of stimuli that just exceed the

¹⁴³ These times come from Poppel (1988). Why haven't I listed ST's for taste and olfaction? As Poppel explains, testing ST's in these sensory systems is difficult considering how stimuli in these modalities are presented.

ST will be perceived as occurring non-simultaneously, but their temporal order will not be accurately represented. Subjects will be at chance in reporting which of two stimuli was presented first, but they will be confident that one of them did in fact come first.¹⁴⁴ In order to perceive the temporal order of events, stimuli must be separated by enough time to pass the *ordering threshold* (OT). Interestingly, the OT is the same, about 30ms, across all sensory modalities for both unimodal and crossmodal TOJs (Kanabus et al 2002).

This is actually a rather interesting point. The experience of temporal order is not itself a simple matter of laying out events along some timeline or simply having experiences that themselves have some temporal order to them. If, as the extensionalist would insist, we experience two sounds as being non-simultaneous, then we will have two non-simultaneous experiences of the sounds and these experiences will have some definite temporal order to them (i.e. they will either be simultaneous or one will precede the other), yet this temporal order of experiences doesn't determine that we perceive those events as being in a certain order. Beyond the perception of simultaneity, the perception of temporal order requires something further than the mere ordering of events.¹⁴⁵

¹⁴⁴ Lee (2014a) takes this phenomenon as a strike against extensionalism by pointing out that accepting the mirroring constraint causes problems in understanding how an experience can accurately represent two stimuli as not happening simultaneously but not as being in any order. If we accept mirroring, then if two events that are experienced as being non-simultaneous, then the experiences of those events will themselves be non-simultaneous. If the experiences themselves are non-simultaneous, then it follows that one or the other experience must occur prior to the other. However, the temporal order of the experiences isn't apparent in our experience, so Lee argues that mirroring fails, as the temporal structure of experience isn't mirrored by the temporal content of experience. Phillips (2014b) responds to Lee by claiming that the mirroring constraint only imposes a *content* → *structure* dependency relation. If some temporal content is present in an experience, then that experience's structure will mirror that content. However, if some temporal property of the experience does not appear in experience, then there is no threat to mirroring as that would require a *structure* → *content* dependency. While Phillips might have saved extensionalism from an outright contradiction, I do think Lee is on to something here. The mirroring constraint was meant to play some role in explaining *how experience comes to have its content*, but if the mirroring constraint only imposes a *content* → *structure* dependency, then the mirroring constraint won't help us understand how experience comes to have its temporal content. As a result, the mirroring constraint becomes even further unmotivated.

¹⁴⁵ What this something further is is a matter of empirical study. It might simply be that the interval between the ST and OT thresholds indicates some source of noise in the representation of temporal order. Perhaps the placement of events on a mental timeline is *fuzzy* and approximate and at intervals closer to the ST the approximate nature of the placement results in enough certainty that the timing of the events is different but not enough to distinguish which precedes which. Or perhaps there are simply different mechanisms at play. However the empirical matter turns out, the points to be made in this section will still hold up.

Any model of the mechanism involved in TOJ must accommodate the data showing that TOJ are sensitive to *on the fly* recalibration. Consider the amount of time it takes for the auditory and visual signals to arrive to the brain from a single event in the world – for example the sound of the ball hitting the first baseman’s glove and the sight of the runner’s foot hitting the base. Umpires often will say that its by noting the temporal order of these events that they judge whether the runner is out or safe. Yet, the perceptual system cannot simply use the onset of sensory stimulation as a means of determining whether the auditory or visual event is occurring first. Basic physics tells us that the light signals travel much faster than sound signals. So, if the perceptual system were simply using time of sensory onset as a means of determining the temporal order of events, then there would be a systematic bias towards judging that visual stimuli occur sooner than auditory stimuli (although, the onset difference in the case of an umpire making a call is far to short to be of any significance). A bias of this sort would wreak havoc on our perception of the world given the fact that once we are about 100ft of remove from an event the arrival difference between auditory and visual signals does begin to exceed what we can discriminate. However, it is also known that the sensory transducers in audition (the receptors in the cochlea) respond faster than the photoreceptors in the retina. So, perhaps here we have some means of counteracting the fact that light travels faster than sound. In fact, it’s been calculated that at 10 meters, the quickness of the response of the auditory transducers are capable of overcoming the speed at which light travels, so that the signals will arrive at the brain at the same time. However, once we deviate from 10 meters, the signals no longer reach the brain at the same time. Somehow, the brain must be capable of discounting for these differences in signal travel time. In fact, there is quite a large literature on this topic that come to be known as the *temporal window of integration*. As long as signals arrive at the sensory transducers within a certain window, they *may* be perceived as simultaneous. A large number of factors influence the size of the TWI such as perceived distance and whether or not the signals concern audiovisual speech.¹⁴⁶

¹⁴⁶ See the discussion of this literature in (Callender 2008).

Similar issues arise when we consider other sensory pairings. Tap your foot and you will see, hear, and feel your foot tapping the floor. And, if all is working well, the sound, the feel, and the sight of the tap will all seem to occur simultaneously. However, the speed at which all of the signals reach the brain differ. Furthermore, the way in which their timing differs changes as you age. As your limbs grow, the time it takes for tactile signals to reach the brain increases.

Now, one particularly interesting feature of the way in which the perceptual system produces TOJs is that the mechanism for representing temporal orderings is sensitive to adaptation effects (Cai et al 2012; Stetson et al 2006). By consistently inserting a short delay between a motor action (e.g. pressing a button) and a visual event (e.g. a flash of light) the two events will come to be perceived as occurring closer in time than they actually are. There is a dragging whereby through repeated exposure the stimuli are perceived as occurring closer in time. Once this initial recalibration has occurred, and the initial delay between the motor action and the visual stimulation is removed, the effect is a drastic change in the perceived temporal relation of the motor action and the visual event. The effect can be so strong that a subsequent trial in which the motor action precedes the flash of light will be perceived as though the flash of light precedes the motor action (the perceived effect will appear to precede its cause!). This sort of adaptive effect is found not only in interactions between motor activity and vision but also found in audiovisual perception as well (Vroomen & Keetels 2010; Vroomen et al 2004). In many ways, the adaptive effects that are found in temporal order judgments parallel the adaptive effects that are found in other perceptual capacities (e.g. the perception of contrast, color, etc).

What sort of mechanism, then, do explanations of TOJ appeal to? Whatever mechanism we adopt would seem to initially have trouble satisfying the mirroring constraint. The initial sensory processing of the motor action and the flash of light will have a temporal structure to it, as the processes are simply physical processes like any other and must have a temporal location, but this temporal structure cannot be mirrored in the ultimate experience of temporal order, since that experience might present the events as

being in a reversed order. However, it is by looking at specific models of the relevant mechanism that we get a better sense as to *why* the mirroring constraint is violated in these cases.

Cai et al (2012) propose to account for the perception of temporal order by appealing to an opponency process very much like the opponency processes that are appealed to explain the adaptation effects in color perception. In their model, the outputs of the individual sensori-motor systems feed into an ensemble of delay-tuned neurons each of which exhibits a selective tuning for a particular temporal delay that partially overlaps with the temporal tuning of other delay-tuned neurons in the ensemble. The responses of these delay-tuned neurons are then pooled by *earlier than* and *later than* neurons. When everything works well, if there is a temporal difference of 50ms between a motor and a visual event, then this will differentially activate the delay-tuned neurons that have response ranges that include 50ms (with peak activation in those neurons whose receptive range centers around 50ms). This differential activation of the delay-tuned neurons would then be summed by the *earlier than* and *later than* neurons and it is the activation of these pooling neurons that then encodes the temporal order between the sensed events. Recalibration occurs through the repeated occurrence of certain stimuli changing the response gain of the *earlier than* and *later than* neurons. When the delay is then removed, the now scaled pooling neurons will respond as though a different group of delay-tuned neurons were activated, thereby shifting the perceived temporal order away from the actual order in which the stimuli were initially processed.

The interest of this model is that the representation of temporal order is independent of the order of the activation of the modality specific regions that process the non-ordered stimuli, and therefore the perceived order of events is independent of the order in which the individual events are perceived. Take the case in which after recalibration the objectively simultaneous stimuli are instead perceived as having distinct timing – i.e. the objectively simultaneous flash and button press are perceived as though the flash precedes the button press. The model of Cai and colleagues explains how the temporal relation of *the flash being earlier than the button press* is represented by a system that violates the mirroring constraint.

While the process by which the delay-tuned neurons influence the pooling neurons is a temporally extended one, there are no temporal stages of that process whose temporal structure mirrors the represented event structure of the flash being prior to the button press. At the level of the pooling neurons and the delay-tuned neurons while all of this processing takes time, it appears as though the flash and the button press are being processed simultaneously. Similarly, if we push the process back and consider more peripheral sensory stimulations, still, the processing of the flash and the button press will be simultaneous. There simply is no mirroring at all between the temporal content of the experience and the temporal structure of the experience itself.

Now, recall, the extensionalist response to the temporal illusions that the atomist appealed to involved a process by which the initial sensory stimulation was delayed by some interval so that the content could be played back to the subject in a manner that respected the mirroring constraint. However, in this case we find that such a playback (an internal movie theater) is not needed. The mechanism described provides all of the content needed to account for the temporal relations we experience. In this way this scenario is very different than those earlier illusions. In those cases, the playback of experience was posited to account for a discrepancy between what we actually experience and some processes that seem to lack that content (provided you accept the mirroring constraint) and lack that temporal structure. However, in this case the mechanism we have discussed fixes the content, and as a result, we have no need for the additional step to “get the content right”.

With the perception of temporal order at these time scales and across modalities we have a clear counterexample to mirroring and thereby a counterexample to extensionalism. Both atomism and extensionalism fail to provide the general characterization of temporal experience as their proponents have intended them to. It’s not the case that the temporal contents of experience bear a single relationship to the temporal structure of experience itself.

4.6. Conclusion

Much of the current philosophical literature on temporal perception is focused around the mirroring constraint. The debate, though, rests on some faulty assumptions about the nature of temporal perception. First, temporal perception is not the unified phenomenon that many of these arguments over the mirroring constraint seem to assume. Instead, temporal perception is the product of a number of different mechanisms that are all specialized for detecting specific aspects of the temporal world around us. On the basis of the fragmentary model of temporal perception we shouldn't have any a priori expectations that we could give a single answer to whether or not our temporal experience satisfies the mirroring constraint. Furthermore, once we start to delve deeper into the actual mechanisms that seem to underpin the various aspects of our temporal experience we find that some of these mechanisms employ representational strategies that satisfy the mirroring constraint while other mechanisms employ representational strategies that violate the mirroring constraint.

The upshot for this all is that if we are trying to explain how it is that we can have experiences with temporal content at all, then looking towards the mirroring constraint, or more generally looking to the temporal structure of experience, simply won't give us that answer. Rather, in order to understand how it is that experience can come to represent time, what we need is a theory of content that could explain how the perceptual systems that are present throughout the animal kingdom, and the ones that are unique to humans, can come to represent time. How experiences are temporally structured, or how perceptual processes unfold in time, might play a big role in explaining *how* experiences come to have their content (as was the case with the temporal dynamics of the SCN) but importantly it won't be the mirroring aspect of these systems that by themselves explains how they acquire their content. The particular difficulty, however, is that time is unlike many of the other aspects of our world that the perceptual system is capable of tracking and representing. Specifically, it's the odd way in which time can (or cannot) causally influence the operation of the perceptual system. Time, or temporal properties of

events, are not capable of causally influencing the operation of the perceptual system in the same way that something like the surface reflectance properties of objects can. For these reasons a simple causal theory of content seems to make it difficult to understand how perception could represent time. It's answering this question about theories of content that is needed to solve the puzzle of temporal experience.

Chapter 5: The Units of Temporal Perception and the Unity of Time

In the last chapter I argued for what I called the *fragmentary model of temporal perception*. According to which, ‘temporal perception’ does not pick out a single unitary phenomenon, but instead, picks out a cluster of distinct capacities for the representation of time. These capacities are underpinned by various distinct types of mechanisms. Some of these specialize carrying temporal information for events detected by a single modality, others are specialized for temporal information that falls within a specific time range, while others are specialized for temporal information that relates events detected by multiple modalities. Even some mechanisms might be specialized for distinct types of temporal properties, i.e. durations as opposed to temporal ordering. Importantly, and central to the fragmentary model, is that these various mechanisms often employ radically different representational strategies in virtue of which they come to carry their temporal information. As a result, there is no single story to be told about how the temporal structure of the processes that underpin our perception of time relates to the temporal content of experience.

While in the last chapter the emphasis was placed on the fragmentary nature of our temporal perception, in this chapter our focus will be to emphasize the way in which time, as it is presented to us in perception and cognition, *is unified*. The events that we perceive, whether they be within a single modality or across multiple modalities, or whether they be at very short time scales or much longer time scales, are all presented as being located within a single and seamless temporal framework. Furthermore, it’s not only that time is unified to us in introspection, or in our thought about time, but also our ability to perform coherent actions in the world requires that we be able to locate the deliverances of the individual sensory systems within a single temporal order. However, given the fragmentary model of temporal perception, the temporal structure of these events are all represented by distinct mechanisms. Somehow, then, it must be the case that the temporal information, encoded in the various mechanisms that underpin

our temporal perception, must be integrated into a coherent whole. It is this integration that is needed to explain *the unity of time* as it is presented to us that will be the focus of this chapter.

Unfortunately, an explanation of how this integration is accomplished won't be found in the empirical literature. The science isn't there yet. In this chapter I propose an account that will help us understand how the brain might overcome this integration problem, but in order to do so, we first need to turn away from the mental representation of time and look at a topic that is much more familiar to those working within the history and philosophy of science. I'll argue that the integration problem that the brain faces parallels in a number of important respects the problem that society faced that drove the establishment of a standardized system of units for the measurement of time. By noting these parallels, and how we as a society unitized time to overcome the analogous integration problem in our social time keeping practices, we will be in a position to realize that the brain could overcome its integration problem through a process of unitizing time. And ultimately, I will argue that the perceptual system already makes use of all the resources that would be required for it to do so. Perception, I argue, *unitizes time*, albeit not in terms of the standardized units like seconds, minutes, or hours.

The chapter will go as follows: In section 1, I will describe two arguments, one by Christopher Peacocke (2015) and the other by Jacob Beck (2014), that attempt to show that perception is *unit-free*. They argue that while it is the case that we readily perceive a number of distinct magnitudes, including duration, mass, length, etc., we do not perceive those magnitudes in terms of any units. In this way, they argue that the perception of time differs in a significant way from how we represent time within our social time keeping practices. In laying out their arguments I will provide some initial reasons for thinking that their arguments fail, however, it will not be until section 3 that we will see exactly how their arguments fail. In section 2, I will describe how we as a society have come to unitize time. In section 3, we quickly return to the arguments by Beck and Peacocke and show how they ultimately fail. With the arguments for the unit-free characterization of temporal perception out of the way, we can turn towards the positive

argument for the unitization of time in perception. In section 4, I recap some of the main arguments for the fragmentary model of perception and emphasize the way in which the integration problem faced by the fragmentary model parallels the problem that drove the adoption of a system of standardized units. In section 5, I go over some of the empirical literature on event perception, and finally in section 6, I argue that the perceptual system unitizes time in the construction of event models.

5.1. The (purportedly) unit-free nature of temporal perception

Central to our cultural practices of measuring and representing time is a system of standardized temporal units. We do not merely measure time, but we measure and represent time in terms of seconds, minutes, hours, etc., and it's in terms of these units that we report time to each other. However, in a series of papers, both Christopher Peacocke (Peacocke 1986, 2015) and Jacob Beck (Beck 2014), have argued that while we readily perceive temporal magnitudes, like durations, we do not do so in terms of any units.

Before we discuss the specifics of either of their arguments, we need to lay out a common core that both authors use to motivate their arguments.¹⁴⁷ Consider the following sort of scenario¹⁴⁸: You're tasked with re-arranging the furniture in a room. There is a piano in the middle of the room and you wonder whether the piano will fit along the empty stretch of wall between the door frame and the bookcase. You look at both and the piano appears to you to have a certain length, as does the empty stretch of wall, yet you can't tell whether they are of the same length or if one is slightly longer than the other. You say out loud to yourself, "Will it fit? Is *that* width," said while pointing at the piano, "the same as *that* width?" said while pointing at the wall. From the other room you hear someone yell out, "The piano is 62 inches wide!" Unfortunately, this simply doesn't help you. Learning the width of the piano in

¹⁴⁷ That there is this common core shouldn't be of any surprise since Beck takes himself to be grounding the unit-free nature of perception that Peacocke first developed in his (Peacocke 1986).

¹⁴⁸ The scenario described here runs together and condenses a few different scenarios that Peacocke uses in his (Peacocke 1986).

terms of inches simply isn't giving you the needed information in a way that would allow you to make the relevant decision. The person from the other room comes into the room, and looks at the piano, and quickly states that the empty stretch of wall is 58 inches wide. Luckily, your friend is an expert at eyeballing lengths and she accurately tells you that the piano will not fit along the wall.

What we're supposed to take from this sort of scenario is ultimately that our perception of length, as well as magnitudes in general since the scenario generalizes, is unit-free. In his *Analogue Content* (Peacocke, 1986), Peacocke argues that if we were to perceive distances in terms of inches, feet, or some other culturally defined unit, then we would expect that upon being given the width of the piano in inches that we would be able to settle the question of whether the piano fits along the wall. Furthermore, the scenario is supposed to give us reasons for thinking that we do not perceive distances in terms of any units whatsoever since someone who is proficient in how to use inches, feet, etc., should be capable of translating the units in virtue of which they perceive distance into the culturally defined ones. Since none of this seems to occur when you're told the length of the empty stretch of the wall, then it is concluded that we simply don't perceive length in terms of any units.

The skill that your friend has in being able to accurately eyeball the length of the wall is explained as not being genuinely perceptual but is the result of a post-perceptual process. According to Peacocke, the lengths of the objects in the room look the same to both you and your expert friend. The only difference is that your expert friend has a cognitive mechanism that allows her to interpret what she perceives in terms of standardized units.

So, if we do not perceive time in terms of any units, then how do we perceive time? Well, neither Peacocke or Beck are perfectly clear about this, however, they do agree on one significant point. Since we can represent magnitudes in both unit-free and unit-laden formats (e.g. *that duration* and *200ms*), this

suggests that the difference lies (at least in part) at the level of something like a *mode of presentation*.¹⁴⁹

Unit-laden representations utilize the standard unit in the representation of a magnitude, whereas unit-free representations do not.

Perhaps the most straightforward analysis of a unit-laden representation, say of a particular length, presents the magnitude of some object or event as a ratio value of a particular unit-standard. For instance, we may define the unit-standard for length in such a way that we assign a value of 1 to a particular object (e.g. the standard meter). Lengths can then be represented in a *unit-laden manner* as ratio values of the magnitude of the standard meter.¹⁵⁰ So, an analysis of the claim that an object is *5 meters long* is can be given as the ratio of the length of the object to the standard meter is 5.¹⁵¹

An alternative analysis of unit-laden representations, given by Suppes and Zinnes (1963 quoted in Peacocke, 2015), appeals to the notion of an *extensive system*. An extensive system includes three

¹⁴⁹ Peacocke hesitates to call the difference making aspect of unit-laden and unit-free representations a *mode of presentation*. His hesitation comes from a very specific, and historical, perspective. According to the original Fregean usage (Frege 1948), modes of presentation, or *senses*, played a special epistemic role in determining the referent (or extension) of a term (or representation). If two co-extensional terms differed in their modes of presentation, then it could be rational for a speaker to doubt whether or not the two terms were in fact co-extensional (as was the case prior to the discovery that Hesperus was the same planet as Phosphorus). However, if two terms ever shared their mode of presentation, then there could be no rational doubt that the two terms were co-extensional, as the modes of presentation determine the extension of a term. Peacocke doubts that this final restriction is satisfied by the magnitude representations in perception. Take the case of the piano and the empty stretch of wall. They both appear to have the same length, and in this sense the actual lengths of the two objects (allowing that the empty stretch of the wall is an object) will be presented to the subject in an identical manner. However, we can still doubt whether they are in fact the same length. Therefore, we cannot identify the manner in which the lengths of the objects appear to us in perception with Fregean senses, since if they were in fact senses, then there would be no doubt. Instead, Peacocke chooses to call the aspect of the representation that provides us with *how things look as the manner of perception* of the magnitude. However, in recent literature, the notion of a mode of presentation has strayed considerably from the original notion proposed by Frege. In this more recent literature, ‘modes of presentation’ is a descriptive term that picks out *whatever it might be* that makes the relevant difference between co-extensional representations. For ease, I will simply adopt this more lenient usage, ignoring the fact that it does not do justice to the original intended use.

¹⁵⁰ The details of the numerical assignment to magnitude values goes beyond our discussion here. Some systems of measurement have an arbitrary zero value and allow for negative magnitudes (e.g. the Celsius temperature scale) while other measurement systems have a zero value that is fixed in some less than arbitrary way and do not allow for negative values (e.g. the Kelvin scale). Temporal measurements are typically of the latter sort.

¹⁵¹ Not only does this analysis strike me as a common-sense analysis (if one could call an analysis of a unit-laden magnitude representation *common-sense*) but it is also one that is given by Suppes and Zinnes (1963) however, they ultimately provide the more complex analysis given in the following paragraph in the main text.

elements $\langle A, R, o \rangle$, where A is the domain consisting of non-numerical objects (e.g. objects with particular lengths), R is a binary relation that holds between elements of A , and o is a function whose domain and range are the elements of A . Supposing that A is a set of lengths, then aRb means that a is at least as long as b , and $a o b$ is interpreted as giving the length c that is equivalent to adding a to b . The non-numerical domain of this extensive system, A , is then mapped to a numerical system that preserves the relations of the extensive system. That is, the numerical assignment of values to the magnitudes in A is given by an isomorphism between the non-numerical extensive system and a numerical extensive system. The mapping is constrained through the choice of a unit standard, that is, some element u in A that is assigned a numerical value of 1. Unit-laden representations, say of a length of 5m, are given by stating that some object has a particular location in an extensive system, and that extensive system is mapped to the numerical extensive system in such a way that a privileged object, u , is assigned a value of 1.

In either analysis the unit-laden representation makes an appeal to the existence of a *chosen object*. Some entity that is taken to define the standard and which is used to establish the manner in which numerical values are assigned to the magnitudes in question. In practice, as we'll see this unit standard is chosen since it is assumed to have a regular magnitude. In the case of length, the current unit-standard that is used by the *International Bureau of Weights and Measures* is the meter which is defined as “the length of the path travelled by light in a vacuum during a time interval of $1/299,792,458$ of a second” (*The International System of Units (Brochure)*, 2006). In the case of the unit standard of mass, the chosen object is an actual object made of a platinum-iridium alloy located outside of Paris. It is through the use of an extensive system, and its unit standard, that absolute measurements of some magnitude can be given.

Unit-free representations are of a different sort. They do not appeal to a single chosen object that is assigned a numerical value of 1 in whatever unit terms are adopted. Instead, we can think of unit-free representations as demonstrative representation. When we perceive an object as having a particular length

we may simply perceive the length of the object as *that length* where the demonstrative picks out the length of the specific object. In fact, it seems as though this is how Peacocke understand the magnitude content of perception and action when he says, “Similar points to these [having just discussed the magnitude content of perception] apply equally to the mental states and events involved in action... You may intend to stretch your arm *that* far, level with the end of the bookcase you see” (Peacocke, 2015).¹⁵²

So, this is what is meant by the unit-free nature of magnitude perception. We perceive magnitudes, but we do so without any appeal to a unit standard. That is, in our representation of magnitudes we do not have to also keep in mind some standard object (or event) that we use as the basis for the representation of other magnitudes of that same type. The claim is not restricted to the (somewhat plausible) claim that we do not perceive magnitudes in terms of the very same units that our cultural measurement practices are structured around. Rather, the claim about the unit-free nature of perception says that there are absolutely no units in terms of which perception represents any magnitudes.¹⁵³

Peacocke on the unit-free character of temporal perception

Peacocke adopts a self-described *metaphysics first approach* to understanding the contents of perception. Begin by giving an account of the nature of the properties in the world around us, and how

¹⁵² There is the additional possibility that we could represent magnitudes as relational properties, e.g. that object is longer than that object. Yet this is simply a more complicated case of the demonstrative approach to representing magnitudes just mentioned.

¹⁵³ Both Peacocke and Beck take this sort of scenario to show more than just that perception is unit-free. In particular, both take this sort of scenario, with a few additions, to show that perceptual representations of magnitudes are *analog* in nature and that they do not pick out particular magnitudes but rather pick out magnitude ranges. The details of this aspect of their analysis of perception aren't vital, however, once you take on board the idea that perceptual representations of magnitudes pick out magnitude ranges, then problems arise for their characterization of the unit-free nature of perception. If perceptual magnitude representations pick out ranges, and not specific magnitude values, then being told that the piano is 62" wide might fail to give you the information to decide whether the piano will fit along the wall because the width range that the gap along the wall is represented as falling within might straddle 62". If that is the case, then being told the piano's width won't help as you won't know where in the represented magnitude range the piano falls. In fact, you can get this conclusion even on the assumption that the magnitude ranges are given in terms of some standardized units, including inches.

those properties impact the perceptual system, and then you will have a story about what the contents of perception are. Taking this approach with magnitude perception, Peacocke gives the following argument for the unit-free perception of magnitudes:

1. The magnitude properties *out in the world* only exist in a unit-free manner [*realism about magnitudes*]
2. The contents of perception, in good cases, are determined by the causes of those perceptual states. [*causal theory of perception*]
3. Therefore, our perception of magnitudes, in good cases, are caused by those very magnitudes we perceive. [from 2]
4. Since, only unit-free magnitudes exist out in the world, then the unit-free magnitudes are what cause our perception of magnitudes. [from 1 & 3]
5. Therefore, our perception of magnitudes is unit-free. [from 5]

As we can see, once Peacocke takes himself to have established a claim about the metaphysical nature of magnitudes he concludes that our perception of magnitudes must be unit-free provided we accept a very specific form of the causal theory of perception. So, let's take this argument a little more slowly.

In defense of (1), consider the following sort of statement:

(M) The car has a mass of 1000kg.

On its surface, (M) is a very simple claim. It predicates a magnitude property of some object. In this case, (M) predicates a particular mass of the car. Claims like (M) are not only commonly used to describe the world, but importantly, they also appear in causal explanations. If a car breaks through the lake ice after an unexpectedly warm day, we might ask the question of *why did the car break through the ice at that specific location? Why didn't it break through closer to shore or further out?* A perfectly respectable

answer would make an appeal to something like (M). The car broke through exactly where it did because it had a mass of 1000kg, and the ice was only so thick, and so on. If it had been more massive it would have broken through sooner, but if it had been less massive then it would have made it further out into the lake. At least in some cases, claims like (M) in attributing magnitude properties to objects also attribute properties that explain the causal powers of the objects those sentences are about. It is because the car has a particular magnitude property, *a mass of 1000kg*, that it has the causal powers that it does.¹⁵⁴

So far, nothing should seem overtly controversial here (however see footnote 154). What Peacocke takes issue with is a tendency in the literature (that he diagnoses) for analyzing statements like (M) in a manner that does away with predication of genuine and real magnitude properties in exchange for analyses of magnitude properties in terms of the systems by which we measure those predicates.¹⁵⁵ According to these *anti-realist* analyses, the property of *having a mass of 1000kg* should be analyzed in the following way: *the ratio of the mass of x* (where x is the object being predicated of) *to the unit standard is 1000*. Peacocke calls these analyses *unit-dependent*.

Peacocke argues that unit-dependent analyses of magnitude properties render magnitude properties causally inefficacious and as a result we would not be able to use statements like (M) in our causal explanations. According to a unit-dependent analysis, (M) could be made false simply by manipulating the unit-standard kilogram. If the actual property of having a mass of 1000kg was defined in

¹⁵⁴ Peacocke begins his argument concerning realism about magnitudes by using an argument of this sort where the example magnitude is mass (in his paper he discusses an avalanche) and then he expects that the argument should generalize to all magnitude types. However, this ability to generalize to all magnitude types seems unlikely, or at least, requires significant argument to establish. Recall, from discussion in chapter 3 and 4, it is actually unclear within the literature on the metaphysics of time whether time itself, and thereby temporal magnitudes like durations, possesses genuine causal powers. In that literature (Lewis 1973; Maudlin 2002; Newton-Smith 1980) it is often argued that temporal analogues of (M), such as *that event took 5 minutes*, should be translated in such a way that the actual temporal magnitude no longer has any causal import. However, as was the case earlier in the thesis, we will put aside questions of whether time is causally efficacious for the purposes of understanding how the mind comes to represent time.

¹⁵⁵ Peacocke attributes these anti-realist views to Suppes and Zines (1963). Earlier we discussed these sort of views as providing an analysis of magnitude statements. Here Peacocke is taking them as analyses of magnitudes themselves.

terms of a relation to the unit-standard, then filing off a little material from the unit-standard would render (M) false by effectively changing the mass of the car, yet, this seems to have little impact on the causal properties of the car. (M), and claims like it, would simply fail to track the actual causal powers of objects. Instead, a more common-sense, and realist approach to analyzing (M) would simply have it that the mass of the car is a genuine property of the car itself. Any unit-dependent analysis, claims Peacocke, will fail to accommodate the use of magnitude claims in causal explanations since they invariably introduce *extraneous objects* into the analysis. If we are then to hold out any hope that magnitude properties are genuinely causally efficacious, then we cannot analyze the metaphysical nature of magnitude properties in a unit-dependent manner. That is, magnitude exist in a unit-free manner.¹⁵⁶

With his particular realism about magnitude properties on the table, Peacocke gives an endorsement of the causal theory of perception when he states:

[W]hen all is working well, it seems that the perception of an object's having a certain magnitude (or magnitude in a range) is explained by the object's having that magnitude (or magnitude in a range). More precisely, a magnitude relative to a certain frame of reference causally explains the subject's perception of that magnitude, relative to that same frame.

(Peacocke, 2015, p. 25)

Given that magnitudes exist in a unit-free manner, and the contents of perception are determined by the causes of perceptual states, then it should follow that magnitude perception is unit-free.

¹⁵⁶ What then of the use of units in reporting the possession of a unit-free property? One way of understanding this is to think of unit-dependent language as providing a means of picking out the unit-free properties. One might try and analyze the fact in some Fregean fashion. Unit-statements are something like senses whereas the actual unit-free magnitudes are the referents. However, we needn't go that far and introduce two levels of semantic content (if senses are even genuinely semantic objects). Rather, we can think of the system of units as providing us with a system of predicates that pick out or attribute (depending on the sentence's syntax) the unit-free properties in the world.

Ultimately, the criticism of both Peacocke and Beck's arguments will rely on what comes in the next section, but we can briefly mention the strategy to be taken there. The process by which any clock, or mechanism for the measurement of time, comes to carry information about time is independent of whether and how we unitize time. The manner in which a system comes to carry information about time is a matter of production of the states of the representational system. Yet, the unitization of time is not on the production side of how a representation comes to carry information about time, but is rather a matter of how that information encoded by the clock is interpreted. That is, we need to distinguish between the way in which the states (or markings) on a measurement device relate to magnitude values out in the world from how those magnitudes in the world are presented to us. A thermometer, for instance, may come to represent the temperature through a number of different means (e.g. expanding fluids, expanding coils, transmission rates of some medium, etc.), but none of these means fixes whether or not we take the information carried by the thermometer as being in Celsius, Fahrenheit, Kelvin, or no units at all. Similarly, the same holds for clocks. A clock will have states that carry information about time independent of any system of units, since the operation of the clock, and its information gathering capacities, will be guaranteed by the dynamics of some aspect of the clock (recall from chapter 3, Smith called this property of clocks *their participatory nature*. Clocks do not merely represent time, but they participate in time (Smith 1988)). Peacocke is simply barking up the wrong tree.

However, independently of the arguments to come, we have reasons to doubt that Peacocke's argument succeeds. The worries for his argument come from his appeal to an implausibly strong form of the causal theory of perception. We know from studying perception more generally that the distal causes of perceptual states do not fully determine the content of the resulting perceptual states. For most perceptual states, especially visual ones, the operation of the perceptual system is *ampliative* in that it adds richness to the incoming sensory signal.¹⁵⁷ Take, for instance, color perception. The retinal response

¹⁵⁷ That perception is ampliative is not in conflict with an information-theoretic analysis of perceptual content. The ampliative nature of perception simply implies that the initial sensory signal does not contain the entirety of the

to incoming light is not a result of any one dimension of the incoming light. That is, the response of any retinal photoreceptor is not the result of luminance, wavelength, or (for some species) polarization. Rather, the retinal responses conflate these various aspects of the incoming signal to the point that one could elicit the very same photoreceptor response by manipulating both the luminance and the wavelength properties of the visual stimuli (Akins 2014; Akins & Hahn 2014; Hardin 1988). Instead, when we perceive the surface color of an object, our perceptual system employs a number of different strategies for singling out the spectral reflectance properties of the object itself, as opposed to retinal responses due to luminance influences or the spectral properties of the illuminant. The ultimate determinate of which surface color we perceive as being out in the world is not merely a product of the distal cause(s) of the perceptual state, but rather it is a complicated causal product of distal stimuli, spectral properties of the surrounding visual elements, properties of the intervening medium, illuminant, and a number of internal states from both perceptual and cognitive systems (Akins & Hahn 2014; Hansen et al 2006; Hardin 1988). The distal cause, the thing that we take perception to be representing, is only one piece in the overall causal story of what brings about a particular perceptual state. A defender of the strong version of the causal theory of perception could respond that the cause of the perceptual state is the sum total of those influences that brought about the perceptual state. However, this would not help their cause as either they would have to be committed to the claim that the resulting perceptual state represents a complex state of affairs spanning both internal and external states or they would have to insist that their resulting perceptual state is about the distal colors of objects but at the cost of making their theory vacuous. The

information that the perceptual system will eventually come to carry about the environment. Rather, the incoming sensory signal carries information that subsequent perceptual processes can use to form representations that carry information about other aspects of the environment. For example, the incoming sensory signal produced by retinal stimulation will carry information that is useful in the representation of surface color, but the sensory signal will not itself carry the information about what color the objects in the world are. It is only through the utilization of this initial information along with certain assumption that the visual system makes that we can have later perceptual processes that bear that appropriate informational relation with aspects of the environment.

complex cause of the perceptual state would simply fail to explain why one particular aspect of that causal story comes to be the content as opposed to any other aspect of the causal story.¹⁵⁸

The causal theory of perception cannot be understood as fully determining the content of perception. However, if we suppose that the distal causes of perceptual states only play a role in determining the content of perception, then we are open to the possibility that while a unit-free magnitude in the world is the distal cause of a perceptual state, perception may nevertheless represent that distal cause in a unit-dependent form.

Let's now turn to Beck's argument, as it will ultimately have the same downfall of focusing on the production side of temporal representations as a means of trying to establish the unit-free nature of perception.

Beck on the unit-free character of temporal perception

While Peacocke adopted a metaphysics first approach to establishing the unit-free nature of temporal perception, Beck (2014) takes a decidedly empirical approach to the same phenomena that adopts a very specific empirical literature as a backdrop. It is becoming increasingly well documented that a range of animals across the animal kingdom have the ability to represent temporal, spatial, and numerical magnitudes. Interestingly, alongside the evidence that the abilities to represent these domains are commonplace in the animal kingdom there are a number of computational / functional similarities between these representational capacities that some have taken as evidence for there being an underlying similarity or shared mechanism involved in the representation of all three domains.

¹⁵⁸ Similar concerns arise in the perception of non-temporal magnitudes. Consider the perception of length in the Ames's Room Illusion. Objects of different lengths seem to be of the same length, and objects of the same length seem to be of incredibly different lengths. In both cases, the actual length of the objects being perceived is not determining the content of length perception. Once again, the distal causal source of the perceptual state does not determine the content of perception.

Some of these similarities include similar adherence to *Weber's Law* in that precision of the representational capacity, in either of these three domains, is inversely proportional to the magnitude of the properties being represented. For instance, for short durations, rats can readily discriminate between a 1 second and a 2 second tone. But, at longer durations, the same difference in duration is not noticeable, as rats fail to distinguish between a 5 second and a 6 second tone (Church & Meck 1984). A larger difference between the stimuli is needed for the animal to reliably distinguish the stimuli. This scalar property of magnitude representation continues as one increases the absolute magnitude of the stimuli being discriminated.

Further oddities arise in considering how the representation of different domains interact with each other. The famous SNARC effect (*Spatial Numerical Association of Response Codes*) (Dehaene et al 1993) shows that there are interesting associations between the magnitude values represented in these various domains. For instance, in one of the earliest studies of the SNARC effect it was shown that the speed at which subjects made an even / odd judgment of a numerical value was influenced by the combination of the size of the number being judged and the spatial location of the numerical stimuli. For large numbers, responses were quicker if the numerals were presented to the right-hand side of the visual field while quicker responses to smaller numbers arose when the stimuli were presented to the left-hand side of space. Now, these associations would seem like pure coincidence unless we were to consider that humans may represent numbers along a mental number line that *is represented as* stretching from left to right. Larger magnitudes would utilize the right hand side of that mental number line thereby facilitating even / odd judgments for larger numerical values when presented to that region of space (Dehaene 2011).¹⁵⁹

¹⁵⁹ SNARC(like) effects have been found for associated responses with a wide range of different numerical, spatial, and temporal combinations. See Hubbard et al (2005) for a review of the large literature on these interactions.

Finally, along the lines of the SNARC effect, there is evidence that perceptual illusions exist that equally influence the perception of space, time and number. For instance, (Burr et al 2011) have shown that immediately prior to a saccade the perception of time, space, and number are all compressed.

Now, a very nice and parsimonious account that would explain the seemingly coincidental commonalities between the representations of space, time, and number would be to say that these are not coincidences at all but arise from there being a single common underlying mechanism for all three domains. *A theory of magnitude* (ATOM) (Walsh 2003) explains the representation of all three domains in this very way. There is a single common mechanism involved in the representation of all three domains.^{160, 161} The computational / functional similarities between the representation of the different domains merely is the result of there being a single underlying mechanism.

For the purposes of understanding Beck's argument for the unit-free nature of perception all we need to do for the moment is suppose that ATOM is true. There is a single token mechanism that underlies the general capacity for magnitude representation that can be used for the representation of magnitudes within specific domains. With ATOM in the background, we can now turn to Beck's argument in which he argues for the unit-free nature of perception on the basis of models of the underlying representational machinery.¹⁶²

The original proposed mechanism for magnitude representations was the pacemaker-accumulator model (as described in chapter 2). This model, originally developed for temporal representation and then

¹⁶⁰ A similar, and earlier, approach to magnitude representation came from Meck and Church (Church & Meck 1984; Meck & Church 1983) when they argued that there was a single pacemaker-accumulator system for the representation of time and number.

¹⁶¹ Variants of ATOM exist. The form discussed in the main text here is one in which there is a single token mechanism underlying all three representational capacities. However, much of the same benefit could be gotten if instead of there being a single token mechanism underlying all three capacities there could be multiple tokens of the same *type* of underlying mechanism. Deciding between these two version of ATOM isn't necessary as Beck's argument fairs equally well (or equally bad) on either version of the theory.

¹⁶² Oddly enough, in his (2014) Beck argues that these mechanisms tell us something about the content of perception, however, in his (2012) Beck argues that these very same mechanisms are cognitive and not perceptual.

extended to cover other domains, includes a pacemaker that produces *clicks* at a regular rate¹⁶³. A gating system controls whether the clicks produced by the pacemaker enter into the accumulator system where the clicks are summed. The total number of accumulated clicks are then taken to constitute a representation of the magnitude being measured. For instance, in the measurement of a stimulus duration, the gate will open at the beginning of the stimulus and close at the end of the stimulus. Since the pacemaker produces clicks at a regular rate, then the total number of clicks accumulated will carry information about the duration of the stimuli. Since according to ATOM, there is only one mechanism underlying the representation of space, time, and number, this proposed pacemaker-accumulator model would apply to the representation of all three domains.

Beck, however, notes a problem with the pacemaker-accumulator model that's been raised in the literature. If the pacemaker-accumulator model is intended to underpin the representation of space and number, in addition to time, then we would expect that it would take the system longer to represent larger spatial and numerical magnitudes than it would take to represent smaller magnitudes in those domains. The explanation of this expectation is due to the fact that the pacemaker produces clicks at a regular rate, and the representation of larger magnitudes requires the accumulation of a greater number of clicks. Therefore, the representation of larger spatial and numerical magnitudes would require a greater accumulation of clicks which would require a greater amount of time for the gate to be open. But, we know that the representation of larger spatial and numerical magnitudes does not take significantly longer to produce than the representation of smaller magnitudes in those domains (Carey 2009). So, the pacemaker-accumulator model simply does not account for the data that ATOM is meant to explain.

Instead of appealing to a pacemaker-accumulator system at the heart of ATOM, Beck argues that the *Church-Broadbent model* (see chapter 2 for a more detailed description of this model) is a better fit. While the pacemaker-accumulator model operated through the regular production of discrete clicks, the

¹⁶³ The clicks are most often understood as being neural spikes.

Church-Broadbent model operates through the continual *motion* of an array of neural oscillators with differing periods. Specific magnitudes are represented as vectors defined over the states (or phases) of the individual oscillators. According to Beck, this sort of model avoids the pitfalls of the pacemaker-accumulator model as there is no need for the clicks to build up over time, and therefore, there is no need for the representation of larger non-temporal magnitudes to take longer than the representation of smaller magnitudes.¹⁶⁴

Putting aside the success of ATOM, what does all of this discussion about these two types of mechanism tell us about whether the perception of time is unitized or not? Perhaps it is best to simply use Beck's own words:

Early models of [analog magnitude representations] were iterative and appealed to pulses of energy generated by a pacemaker and stored in an accumulator. According to the accumulator model of Meck and Church (1983), the rat's pacemaker releases a pulse every 200 milliseconds. Thus, *according to their model duration is measured in units that are extensionally equivalent to 200 milliseconds*, such that a three-second duration (say) would be equivalent to fifteen of these units... More recent models, which operate in parallel, *are not obviously committed to any units*. For example, according to Church and Broadbent's (1990) model, durations are associated with the phases of a set of neural oscillators of differing periods. Thus, a given duration will typically be associated with several oscillators whose individual periods total to the duration being measured (ignoring error). *As a result, there is no unit in this model by which a total duration is*

¹⁶⁴ While this is what Beck claims, it's likely false that the Church-Broadbent mechanism is able to avoid the pitfalls of the pacemaker-accumulator mechanism. In the representation of larger temporal magnitudes, the Church-Broadbent model makes use of the longer period oscillators, that is, it takes time for the array of oscillators to arrange themselves to encode a longer duration. Similar time constraints would be in place for the representation of larger non-temporal magnitudes as well. However, if that's the case, then it would follow that the representation of larger non-temporal magnitudes would also take longer, since the system will only arrive in the appropriate state after a (relatively) lengthy time-dependent process has completed.

measured... To borrow a helpful phrase from Christopher Peacocke (1986), [analog magnitude representations] might thus be characterized as *unit free*.

(Beck, 2014, p. 27, emphasis added)

Beck continues to say:

I hypothesize that the phenomena [the unit-free perception of magnitudes] that Peacocke [1986] isolates is [sic] grounded in [analog magnitude representation] – i.e., that our conscious experiences have a unit-free character because those experiences are generated by [analog magnitude representations] which are themselves unit free.

(Beck, 2012, p. 27)

So, Beck is making the following inference – because the Church-Broadbent model does not operate through any discrete *clicks* it does not represent magnitudes in terms of any units, however, if the underlying mechanism were to operate through some discrete clicks, then it *would* represent magnitudes in terms of some units. Very clearly, Beck is making an inference from the structure of the representational vehicle to the content of the representation. In general, of course, this sort of inference is always suspect as it is easy to make a vehicle / content confusion, especially when temporal representation is at issue.

But of course, there is just the threat of making a confusion. For all the cases in which such an inference leads us to a falsehood, there may nevertheless be cases in which the inference leads us to some truth about mental representation. To see why this inference of Beck's is in fact committing a vehicle / content confusion, let's turn to the following section where we look at how we as a society have come to unitize time. It is here that we'll see that both Beck and Peacocke are making essentially the same mistake. Their arguments concern the mechanisms, or explanations, for how the initial time keeping mechanisms in perception come to have their temporal content. But, this is simply mistaken. The

unitization of time is not a matter of how clocks come to carry information about time, but is rather a matter of how we use and interpret those clocks. The standardization of temporal measurement through the adoption of a system of standardized units for time provides us with a framework, or common-code, by which we can represent and communicate the temporal information encoded in individual clocks, thereby allowing us to reliably compare and integrate the readings made by various types of time keeping devices.

5.2. The cultural unitization of time

The topic of measurement is one that's received quite a bit of attention from historians and philosophers of science. Most of that literature, though, approaches the topic of measurement from the perspective of very general issues in the philosophy of science concerning the relationship between theory and observation and the role of convention and operationalism in scientific practice. While those topics are surely worthwhile topics to pursue, our interest in this chapter is actually something that often gets mentioned in passing in the first few pages of works that discuss these other aspects of measurement. Invariably, in setting up the topic to be discussed, both historians and philosophers of science will acknowledge the central role that the adoption of a system of standardized units for measurement plays in explaining scientific progress.

In order for science to progress it must be the case that the measurements that our theories rely on are reliable. This reliability comes in two forms. First, there is the within device reliability (i.e. within device precision). It had better be the case that if a measurement tool I use in one context produces a particular output (e.g. displays particular numerals in a digital display or has an indicator arrow pointing in a particular direction) for the measurement of a particular magnitude, then the next time that the tool produces that particular output that it will indicate the same magnitude property out in the world (or conversely, that when confronted with the same magnitude the measurement tool will produce the same

response). The second form of reliability appears when we consider the regularity of measurement across measurement devices (i.e. across device precision)). It better be the case that when I measure some event as having a duration of 200ms today in Vancouver, that a measurement in New York tomorrow that indicates 200ms will pick out the same temporal property as my measurement in Vancouver. There is a sense in which a measurement device can get things *right* or *wrong*.¹⁶⁵ These forms of reliability have to do with the ability to properly calibrate our measuring tools. The adoption of a system of standardized units, i.e. the adoption of a *chosen object*, allows for the possibility of this sort of calibration.

Closely related to these forms of reliability, the adoption of a system of standardized units provides us with a common-code for the representing of magnitudes measured by, what are often, radically different types of measurement devices. It is through the adoption of a single system of measurement units that we can compare the readings of a stopwatch, hourglass, and analog clock. Not only does this common-code allow us to calibrate these devices to achieve some degree of interdevice reliability, but it allows us to use different types of devices that operate over very different scales. To help us understand this point, we can look to the history of the development of standardized practices for the measurement of temperature. In his book, Hasok Chang (2004) has a wonderful quote from William Halley (the discoverer of the comet). Halley, voicing his frustration with the state of the practices for the measurement of temperature in the 17th Century said,

I cannot learn that any of them [different measurement devices]... were ever made or adjusted so as it might be concluded that what the Degrees or Divisions...did mean; neither were they ever otherwise graduated, but by Standards kept by each particular Workman, without any agreement or reference to one another.

¹⁶⁵ *Getting things right*, or accuracy, with a measurement device involves quite a bit of work of stipulating just what properties out in the world should go with what markings there are on a measurement device. For instance, for a thermometer to get the temperature right when it reads 100 degrees Celsius, it must be the case that the reading of 100 degrees Celsius be in response to a property that is chosen according to our temperature measuring practices as being what should be denoted by a reading of 100 degrees Celsius.

(Halley, 1693, p. 655, quoted in Chang 2004, p. 9)

You can think of the situation in this way: In 17th Century Europe various people had developed a variety of different tools for the measurement of temperature. Imagine, if you will, that scattered across the laboratories of Europe were a variety of different canisters, of all shapes and sizes, filled with all sorts of different fluids, some mercury, some alcohol, some water (of various purity), and located at a variety of different elevations, all being used to gather information about the ambient temperature of the medium in which they were placed. Each “Workman” within his own lab would employ one of these devices, with its own peculiar behaviors, and with its own idiosyncratic system of markings, and they would produce measurements of phenomena that were likely reproducible within their own labs. Further, each “Workman” could produce reliable measurements of the relative temperatures of phenomena in their own lab, as they could compare the indicator values for their own devices in different contexts.

However, what they could not do, and what spurred Halley to voice his frustration, is reliably compare the measurements made in one lab with the measurements made in another lab. If one person in Vienna noted that a particular chemical reaction raised the temperature of a quantity of water by 3 markings on their own measurement scale, then someone working in Oslo, using a different measurement device, would face significant troubles in interpreting the Viennese findings. The only way that they could make sense of the findings of the Viennese lab would be if they possessed the special knowledge needed to translate the readings of the Viennese thermometer into the marking found on their own thermometers. Otherwise, the Viennese report would be almost meaningless.

What was needed was not simply a *chosen object* to calibrate the various measurement devices (i.e. that the boiling point of water should be used to calibrate thermometers), but also some common-code or scale that would allow for the common interpretation of the various results produced by various different types of measuring devices (i.e. how to define an extensive / numerical system on the basis of the chosen object). In this way, the movement of a fluid by three marks in own person’s device could be

compared with the movement of another fluid by 2 marks in another person's device through the adoption of a single common scale. It isn't required that each thermometer therefore be made of the exact same material. Instead, all that is required is a common representational framework within which to interpret the various measurement devices. If each "Workman" were to learn how their own device related to the common scale, then there would be no need to learn any idiosyncratic translation rules between measurement devices. Thus the communication of various researchers would be far simpler and the measurements that each researcher produces could be integrated to produce a more coherent and unified picture of the world that is being studied.

In the case of temperature, the confusion that Halley was concerned with was overcome by the adoption of a standardized system of measurement that included the adoption of a single unit of measurement and a common scale.¹⁶⁶ Now, if we turn away from temperature and focus on time, we find that many of the same concerns arise in the measurement of the temporal domain.

Looking through history, and even around the world today, we find that people have used a strikingly diverse assortment of devices for keeping track of and measuring time – e.g. sundials, hourglasses, digital clocks, analog clocks of different sorts, 24-hour clocks, 12-hour clocks, water clocks, atomic clocks, etc. Any one of these devices could be used by a person to keep track of time, since like the thermometers, there is a lawful operation of the underlying mechanism of the device that allows the device to carry information about the domain it is measuring. However, just like the various thermometers of the various Workmen that Halley despised (the devices, not the people) the markings on these various time keeping devices would be unintelligible to those using other types of time keeping devices, unless someone were to have specific information about how to translate between the various markings.

¹⁶⁶ Another significant advance in the measurement of temperature arose through the adoption of mercury as the reliable fluid for the construction of thermometers. How it was determined that mercury, as opposed to water, reacted most reliably to changes in temperature, without appeal to a reliable thermometer to determine this fact is a fascinating question in the literature on the epistemology of measurement. For discussions see (Chang 2004; Fraassen 2008; Tal 2013, 2016).

Now, of course, if the number of types of time keeping devices is relatively small, then one could overcome the calibration and integration problems by learning the translation rules between every pair of measurement devices (or at least between one's own measurement tool and those of everyone else). However, this is clearly not how we as a society overcame this problem. Instead, each person learns how their own device (or whichever devices they use) relate to a single common-scale defined by the standardized units of *seconds*, *minutes*, *hours*, etc. By adopting this system of standardized temporal units, we now have a means for the calibration of the various time keeping devices, and a means of communicating the temporal information gathered by the various time keeping devices used around the world.

In order to ultimately draw the parallel between the operations by which we as a society have come to unitize time and how our brains might overcome the temporal integration problem to bring about the unity of time as we perceive and cognize it, we need to look in a little more detail as to how we as a society have unitized time. It is by looking at this process that we'll see the particular resources that we as a society have used for the establishment of a system of standardized temporal units.

Since 1967 the second has been defined as “the duration of 9,192,631,770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the caesium-133 atom” (*The International System of Units (Brochure)* 2006). This physical process, the energy transitions of the caesium-133 atom are the *chosen object* for our system of temporal measurement and is mapped to a value of 1 in the numerical extensive system that we employ. With this unit standard in hand we are put in a position to calibrate the various time keeping devices that we use. While this standard definition is simply stated, the actual process by which this definition is implemented in the service of standardizing our time keeping practices is rather complex.¹⁶⁷

¹⁶⁷ Much of the story that is about to be given can be found in more detail in a wonderful paper *Making Time* by Eran Tal (2016).

The process begins with a collection of 12 *primary clocks* scattered around the globe. These primary clocks are the devices that most closely *realize* the definition of the standard second in that they are the most reliable devices that we can construct that allow for the monitoring of energy levels of the caesium atom. While these primary clocks are incredibly reliable, they nevertheless differ slightly amongst themselves in the speed of their operations and due to a number of technical reasons are only reliable over the relatively short scale of several weeks. These short-lived primary clocks are used to calibrate a larger collection of *secondary clocks*. These secondary clocks, the more familiar atomic clocks, while not quite as reliable as the primary clocks are nevertheless capable of sufficiently reliable behavior over the course of months and even years.

However, the process does not stop with the calibration of the secondary clocks. Each of the several hundred secondary clocks differ in the speed of their operations. That is, for any given interval, there is no guarantee that any two secondary clocks will give the exact same reading. Since there is no standard time keeping device apart from these primary and secondary clocks we cannot appeal to some more authoritative clock to determine which clocks are correct. Instead, a reliable temporal standard is bootstrapped from the operation of these various secondary clocks.

At regular intervals a weighted average of a collection of 420 secondary clocks is used to produce a standard measurement of the time that elapsed over the just finished interval. Any clock whose reading matches (or approximates within a certain margin) the weighted average is upgraded and given more weight in subsequent calculations (with a maximum weighting of 0.7). If a clock's reading differs significantly from the calculated weighted average, then that clock is downgraded and given less weight in (and ultimately removed from) subsequent calculations. This weighted average is the source of a number of different time standards that are ultimately what we coordinate our time keeping practices with. Perhaps the most important of these temporal standards is *Coordinated Universal Time* (UTC) that is used for most civilian, governmental, and scientific practices.

Once UTC is calculated, clocks around the globe are (re)calibrated to UTC at five-day intervals. At each of these intervals, the measurements of the clocks around the globe are compared to the retrospective weighted average calculated in the above manner and they are corrected to accommodate for any disparity between the calculated interval given by UTC and the interval noted by the individual clocks.

There might be a temptation to describe the process by which UTC is calculated as employing one large, complex, and distributed clock mechanism. The individual primary and secondary clocks would simply be understood as sub-components of a larger clock that they compose, and it is to this *universal clock* that we calibrate our time keeping practices. There would be some ‘tick’ (or group of ticks) of this clock that we could identify as the implementation of the unit-standard second. However, to think of the calculation of UTC as being the product of one large distributed clock would simply be to mistake the exact manner in which UTC is produced.¹⁶⁸

Clocks operate by exploiting some process that is taken to have a regular duration. It is because the swings of the pendulum have a certain period, or because the sand in the hourglass fall at a specific rate, that these types of clocks come to carry information about time. However, there is no such assumption in the production of UTC. Rather, the calculation of UTC *imposes* structure on the underlying time keeping mechanisms that are used in the production of UTC. There is no updating of UTC that is taken to have a regular duration, rather, UTC provides us with the means to impose some regular or common scale to time as measured by actual clocks. It is not even assumed that the updating of UTC will occur at a regular interval, as the ultimate duration of the interval between any calculation of UTC is calculated in retrospect. It is this retrospective aspect of the calculation of UTC that shows that the system as a whole cannot be interpreted as a single large clock.

¹⁶⁸ Thank you to Eric Margolis for making me see that I needed to clarify this point.

As a result of this distributed process, Eran Tal has claimed that,

UTC is not a clock; it does not actually ‘tick’ and cannot be continuously read off the display of any instrument. Instead, UTC is an abstract measure of time: a set of numbers calculated monthly in retrospect, based on the readings of participating clocks.

(Tal, 2016, p. 4)

It is through the calculation and distribution of UTC that the individual clocks from around the world are calibrated to a single temporal standard. Furthermore, this temporal standard provides us with a single common-code for the interpretation of the various time keeping devices around the globe. Given the unit standard we can interpret the marking along a clock face in terms of *seconds*, *minutes*, and *hours*. There were of course other means by which time was standardized prior to the adoption of UTC, but the story remains the same. It is a result of this standardization, and the adoption of a system of units, that allows for the ready integration of temporal information across measurement contexts.

There are a few important points to take home from this. First, there is no actual physical thing that we can point to that is the standard second. There is no physical clock against which we can compare and calibrate all other clocks. Instead, the process of unitization provides us with a representational framework that gives us the means by which to *interpret* the operation of various particular clocks. By taking a chosen event (in the case of time this is an idealized event abstracted from the operation of a great many clocks) as defining the unit standard we thereby produce a representational system for representing temporal magnitudes regardless of how they are measured. Once we know how to interpret a clock’s readout, whether it’s the direction of hands on the clock face, numerals on a display, or levels of sand, in relation to the standardized time given by UTC, then we have a common-code by which we can

compare the readings of various time keeping devices. All we need to know is how the clock's readout relates to temporal standards.¹⁶⁹

Second, and this is rather important for our purposes, is the fact that the operations by which we come to interpret a clock in terms of some system of units is independent of how that clock comes to carry information about time. Take for instance an analog clock that works through the decompression of a spring that turns a series of internal gears. It is the regular rate at which the compression spring releases energy and how this translates into movement of the internal gears and the hands of the clock that establishes the information gathering properties of the clock. Similarly, it is the regular rate at which sand falls through an hour glass that establishes the information gathering properties of an hourglass. In both cases, clocks come to carry information about the temporal world in virtue of the way that they themselves operate through time (e.g. since the dynamics of the analog clock are what they are, if all is going well when the hands move this or that way across the face of the clock, then X amount of time will have passed). Importantly, nothing about units of time is required for clocks to have this information gathering capacity. Our social conventions simply do not enter into the picture.

Yet, for us to make use of this information we must encode this information in a manner that is appropriate for how we will use that information. In the case of our social time keeping practices, the need to coordinate our time-sensitive behaviors regardless of the time of time keeping device being used requires that we encode this temporal information in some common-code that can be applied to all time keeping devices. That is, we need to interpret the information encoded in clocks in terms of some standardized system of units. In this way, we come to represent the temporal properties recorded by our

¹⁶⁹ As Tal (2016) puts it, the unitization of time consists in the establishment of application conditions for a measurement concept. There is not physical thing, the standard second, by which we represent time in relation to. Rather, the unitization of measurement provides us with a common framework for the application of some representational mechanisms. I by and large agree with Tal, however, for our purposes here, since we want to apply this literature to the study of the mind, I would shy away from the use of the term 'concept'. Concepts as they appear in the philosophy of science are rather different sorts of creatures (or at least may be very different) from concepts as they are appealed to in the cognitive sciences as an object of study.

time keeping devices under a specific mode of presentation. And in fact, this accords nicely with the original Fregean notion of a mode of presentation, where modes of presentation were given a heavy epistemic weight to them. It is because of the particular mode of presentation of some worldly object or property that we come to acquire certain inferential capacities having to do with that way of specifying the object.¹⁷⁰

Our interpretation of a clock in terms of some system of units is independent of whether that clock in anyway *ticks* or *makes use of discrete processes* in order to carry information about time. In other words, the interpretation of clocks in terms of a common representational code given to us by a system of units is just that. It is *a way of interpreting information encoded in clocks* and *not a manner by which clocks come to gather this information*.¹⁷¹ We interpret the information encoded in clocks in relation to some representation of a physical process (real or idealized) that we have taken to be reliable in its duration. If we want to understand how some measurement system unitizes time (or whether it unitizes time at all), then we should not look to how the actual measurement devices employed by that system come to carry information about time. Rather, we need to look at how consumers of those time keeping devices come to use and represent the information encoded by the initial measuring devices.

¹⁷⁰ This way of understanding modes of presentation can be traced back to even earlier work of Frege, namely his *Begriffsschrift* (1879) in which we find the seeds of his later sense / reference distinction.

¹⁷¹ There's a sense that many people might simply be utilizing an overly simple idea of measurement when they think that units are involved in the actual gathering of information by our measurement devices. As young children in school we are often taught how to measure distances by being given a piece of paper or stick that is one foot in length. We then measure longer distances by lining up end to end these one-foot-long sticks. We then simply count how many sticks we used and there we have a measurement of distance in terms of feet. Perhaps units might be employed in a more direct way for the measurement of space (or other domains), but I hope the story of how temporal measurements are standardized that we just saw shows that this sort of story cannot be used for the temporal case.

5.3. Returning to Beck and Peacocke

We can now quickly point out why the arguments of both Beck and Peacocke are unable to show that perception is unit-free. As we saw by looking at the process by which we as a society have come to unitize time, the means by which any clock comes to carry information about time is distinct from the process by which we come to interpret the information carried by those clocks in terms of some standard system of temporal units. By focusing on the discreteness or lack of discreteness of the underlying mechanisms involved in magnitude representation, Beck is entirely focusing on the means by which the mechanisms come to carry information about time¹⁷². Similarly, Peacocke by focusing on the causal origins of our magnitude perceptions also focuses on the process by which our perceptual states come to have their contents. That is, Peacocke focuses on how it is that perception could come to be about (or carry information about) the magnitudes exemplified in the world around us.

If we want to know whether perception unitizes time, then we need to know not how the peripheral time keeping mechanisms of perception come to carry information about time, but whether there is a system employed by the perceptual system that imposes a unitized structure on the temporal information encoded by these time keeping mechanisms. In the same way that UTC is used to provide a structure or common-code for the interpretation of the individual clocks around the world, we need to see whether there is such an analogous downstream representational system. But, as it stands, the arguments for the unit-free characterization of perception seem to fail, and as a result, we are open to provide positive reasons for thinking that perception does in fact unitize time. It's to this task that we turn now.

¹⁷² The mistake that Beck is making is similar to a related content / vehicle confusion made by (Gallistel & Gelman 2000) in regards to whether the discreteness of the underlying magnitude representations bears any significance for whether our numerical representations represent integers or real numbers. See Laurence and Margolis (2005) for an articulation of this criticism.

5.4. The fragmentary model of temporal perception¹⁷³

In the previous chapter I argued for the fragmentary model of temporal perception. Due to a variety of different reasons it was argued that temporal perception is not the product of a single unitary mechanism, and therefore, temporal perception is not a single unified phenomenon or capacity. The mechanisms by which we perceive time are specialized and operate over specific time scales. Some mechanisms operate within individual modalities while others operate and produce representations of time that apply to events detected by multiple modalities.

Not only are there multiple token-distinct mechanisms involved in temporal perception, but these mechanisms are of radically different *types*. That is, the mechanisms that gather information about the temporal structure of our world, what we can call *the peripheral clock-like mechanisms¹⁷⁴*, utilize radically different representational strategies. These mechanisms differ from one another in much the same way that a sundial, hourglass, analog clock, and digital stopwatch all come to carry information about time by employing radically different representational mechanisms.

To briefly recap the arguments for this claim recall that in the last chapter important differences were pointed out in the neural models of intra- and intermodal temporal representations at various time scales. At short time scales, between 30ms to 500ms, recently popular models exploit the intrinsic state-dependent properties of neural populations that encode non-temporal stimulus properties within individual modalities. For instance, Buonomano and colleagues (Buonomano 2000; Buonomano & Maass 2009; Goel & Buonomano 2014; Karmarkar & Buonomano 2007) have argued that state-dependent properties of neural populations, specifically the time-dependent evolution of the spatial distribution of activity within the population, can be used to carry information about stimulus duration. Furthermore, the

¹⁷³ To avoid copious amounts of redundancy between this section and the argument for the fragmentary model from the previous chapter I will only be providing a restricted defense of the view here. See chapter 4 for a more complete argument.

¹⁷⁴ These mechanisms are peripheral in the sense that they are the *initial* mechanisms that gather specific types of temporal information in perception.

state-dependent properties can be generated in biologically plausible neural networks as well as with neural assemblies maintained in vitro (Finnerty et al 2015). Finally, these state-dependent models nicely accommodate the manner in which the perception of time at these short time scales seems to be sensitive to the immediately preceding perceptual stimuli in a way that other models fail to explain (Ivry & Schlerf 2008).

At longer time scales, state-dependent models fail (Ivry & Schlerf 2008). But, other models seem to do nicely here. In particular, at the scale of several seconds, the classic pacemaker-accumulator models seem to do nicely. In particular, the *scalar expectancy model* (Gibbon, 1977) nicely accounts for how attention and memory both influence the perception of time at this time scale. Yet, these models do less well at the very short time scales that the state-dependent models cover.

Finally, as argued in chapter 4, cross-modal temporal perception, and specifically the cross-modal perception of temporal ordering at short time scales, seems to employ yet further different mechanisms. Cai et al (2012) attempt to explain the manner in which cross-modal ordering stimuli can be adapted on the fly by appealing to an opponency-system that mirrors the opponency systems found elsewhere in perception.¹⁷⁵

Information about the temporal structure of the world is gathered by a diverse collection of different time keeping mechanisms. Furthermore, how these diverse types mechanisms represent time varies depending on which type of mechanism we focus on. To put the point slightly differently, these various mechanisms employ radically different codes for the representation of time. Yet somehow the temporal information gathered by these different types of mechanisms must be capable of being integrated to produce the apparent unity of time that was mentioned at the outset of this chapter. Somehow temporal information from different modalities and at different time scales comes to be

¹⁷⁵ For details see page *case #2* in chapter 4.

integrated to drive coherent behavior and the apparent unity to time that we confront upon introspection and thought.

Notice the parallels here between the integration problem that brain faces and the sort of confusion that worried William Halley. Halley was concerned with the fact that each “Workman” employed a distinct type of device for measuring temperature and for each device there was a distinct standard to which the device was calibrated each device would report its measurements in terms of its own idiosyncratic measurement scale (i.e. it’s own idiosyncratic marking system). In order to be able to compare the readings made between devices, and to use multiple devices to form a coherent representation of the world, some means of translating between these systems was necessary. That translation process was facilitated by the adoption of a standardized system of temporal units. By adopting a system, like UTC, society was capable of using the standard second to coordinate our time keeping practices.

The brain faces a similar problem. According to the fragmentary model of temporal perception, our perception of time is the product of a great many different types of time keeping mechanisms that employ radically different representational strategies. You can equate these representational strategies with the different markings on the different measurement devices used for temperature. Each time keeping mechanism has its own (type specific) idiosyncratic means of representing and encoding time. Yet, somehow, as we’ve discussed the temporal information stored in these various devices needs to be integrated to underpin the unity of time in how it is presented to us. The question, then, is how does the brain overcome this task? Is there a system that translates the temporal information encoded in the various time keeping mechanisms of the brain into a single common code? Or, is there some other explanation for why upon introspection and through action temporal information, gathered through various modalities and at various time scales, seems to be integrated into a single seamless temporal order?

At this point the science of temporal perception doesn't have an answer for how the brain overcomes this integration problem. Furthermore, it's not even clear that what we're looking for is *the* way that the brain overcomes this problem or whether there are going to be multiple different ways in which the temporal information encoded in the various mechanisms in perception gets integrated. So, we can't appeal to the empirical literature for a direct answer to how the unity of time as its presented to us in perception and cognition gets explained. In the section that follows I will lay out some empirical findings that give us reason to think that we can explain the integration of temporal information through a process where perception itself unitizes time.

5.5. Event perception

Let's begin by making a terminological clarification. In what follows I will be using the term 'peripheral time keeping mechanism', or varieties of the same, to pick out the particular time keeping mechanisms that gather information in perception. These are the sorts of clock-like mechanisms that were described in the discussion of the fragmentary model of temporal perception. The task before us is to explain how the temporal information encoded in these peripheral mechanisms comes to be integrated into a coherent picture of the temporal structure of the environment. Those systems that utilize the temporal information encoded in these peripheral mechanisms will be what I call 'consumers' of this information.

Now, given that there are quite a few consumers of the temporal information encoded in peripheral time keeping mechanisms (e.g. motor systems, learning mechanisms, complex perceptual systems, etc), it's rather unlikely that there will be one explanation for how temporal information gets integrated in the brain. In what follows, we'll focus on one specific sort of consumer system, the system involved in *event perception*.

When we, adults with properly functioning perceptual systems, perceive the world around us we do not merely perceive objects and their features. We do not merely perceive a static world. Rather we see objects that exemplify features that figure into dynamic events.¹⁷⁶ Our perceptual systems are capable of stitching together the various deliverances of the individual senses into a complex and dynamic representation of the world around us. Events in the world can be short, as in *the impact of the ball against the baseball bat*, slightly longer, as in *the swinging of the baseball bat*, slightly longer still, as in *the pitcher's motion and the batter's swing in response*, and even longer still, as in *the entire at bat*, *the entire half of the inning*, and even *the entire game*. We come to grasp these events in our world (at least partly) through our perception of these events. On the one hand, perception by itself is likely capable of providing us access to shorter lived events, e.g. *the impact of the ball on the bat*, *the batter's swing*, etc. On the other hand, our grasp of longer scale events, e.g. *the entire half of the inning*, is likely a product of some combination of perceptual and cognitive mechanisms, as the representation of these longer scale events require many more cognitive resources (perhaps most obviously memory processes play a significant role in these cases). While it is arguably the case that at both the perceptual and cognitive stages at which events are represented memory systems are involved, as the events we are representing become longer this memory component takes an increasingly central role (as more and more of the event details cannot be accounted for by online sensory processes). Interestingly, though, there seems to be no clear cut introspective boundary between the *perception* and the *cognition* of event structure - there isn't a clear boundary as to where our perceptual grasp of events transitions into a cognitive grasp of events. While the literature on event models in perception acknowledges that there isn't a clear upper limit to where our perception of events ends (see Zacks et al (2007) and Zacks & Tversky (2001)), we can restrict ourselves to talking about shorter events that span several seconds, as this is sufficient for the discussion to follow.

¹⁷⁶ The discussion in this section has to be restricted to a discussion of the adult perceptual system. What infants are capable of perceiving is a difficult question that requires its own extended discussion.

Event perception is an ideal case for evaluating how temporal information comes to be integrated in perception as event perception requires the integration of information from various modalities and across various time scales. Consider what it takes to perceive a pitcher's windup and pitch at a major league baseball stadium. We receive information from all the senses, we see the windup, we feel the vibration in the stands as the crowd cheers, we hear the chants of the crowd and the crack of the ball as it hits the bat. All of the senses contribute some temporal information. Furthermore, the event can take up to several seconds depending on the particular style of the pitcher, so we also have to integrate temporal information from the millisecond scale, used to perceive the crack of the ball on the bat, the movement of the pitcher's arm, the movement of the batter, etc, to the scale of several seconds by watching the entire event unfold or the entire movement of the pitcher. Furthermore, the distances in a baseball stadium are often sufficiently large that the transmission differences between auditory and visual signals are capable of being detected and must be compensated for. All of the problematic interactions that will interest us are present in just this sort of perceptual episode. Since the event is a multimodal one that spans several time scales we know, given the fragmentary model of temporal perception, that the perception of this event requires that we integrate the temporal information encoded in various peripheral time keeping mechanisms to construct a representation of the complex event as occurring over a seamless temporal interval.

Many descriptions of event perception describe the process as one in which the perceptual system constructs an *event model* (Zacks et al., 2007; Zacks & Tversky, 2001). The models not only consume the temporal information encoded in the peripheral time keeping mechanisms, but they importantly often include a *predictive* and *retrodictive* elements. Consider what happens in cases of trajectory estimation. The spatiotemporal path that some object is perceived as taking can often be influenced by later perceptual events, as was the case with the postdictive effects found in the cutaneous rabbit illusion.¹⁷⁷ In

¹⁷⁷ For discussions of these cases of apparent motion see Grush (2007). In fact, Grush provides an account of temporal perception that relies heavily on evidence concerning both retrodictive and predictive effects found in

these cases later stimuli will cause a re-writing of the perceived trajectory that some object takes. In other cases, the process by which trajectory estimation occurs involves a *predictive* component whereby the perceptual system anticipates where an object *will be*. Take for instance cases of *representational momentum*. In these cases, subjects are shown an object tracing some predictable path along a computer screen, and then the object is removed. Subjects are then asked to point to the final location of the object on the screen. In most cases subjects will not point to the final location of the stimuli but will instead point to a location somewhat farther along the trajectory that the object had been following.¹⁷⁸ These cases of representational momentum are taken to show that when asked to point to the last location of the object, subjects are not acting on the basis of the last location at which there was sensory stimulation (i.e. the last location that corresponds to the last area of retinal activation) but are instead acting on the basis of an internal representation of the worldly event that goes beyond the activity of the peripheral sensory areas.¹⁷⁹

Importantly, these event models are not clock-like in how they represent the environment. Event models are not formed on the basis of some further clock-like mechanism that tracks either the worldly events or the more proximal states of the sensory system. That event models contain both retrodictive and predictive aspects makes it unlikely that the representation of temporal structure in event models is the product of any clock-like monitoring of internal sensory stimulation, as there simply would be sensory stimulation of the sort to track (especially in the predictive case). Instead, event models are distinct type of representational system than what we find in the peripheral time keeping mechanisms. While the peripheral time keeping mechanisms are clocklike in that they gather information about the temporal

trajectory estimation. In this way, we can actually understand his *emulation theory* (Grush 2004) as best describing the perception of events, and not simply the perception of time in general.

¹⁷⁸ For a selection of the literature on representational momentum see (Freyd & Jones 1994; Freyd & Finke 1984; Hubbard 1996; Hubbard & Bharucha 1988; Verfaillie & d'Ydewalle 1991).

¹⁷⁹ So, here in the construction of event models it appears as though once again we have a place where the temporal structure of the representational vehicle or process comes apart from the temporal structure of the events being represented.

structure of the world, the event models instead provide a representational framework within which the temporal information encoded by the peripheral time keeping mechanisms can be represented in an integrated form.

Beyond the fact that our event models provide a framework within which the temporal information encoded by various time keeping mechanisms can be integrated is the fact that the mechanisms involved in the construction of event models segment events into constituent chunks (or smaller events) at multiple time scales in parallel. In a series of experiments (reported in (Zacks et al 2007)), subjects were shown a complex action (e.g. someone cleaning a room) and were asked to press a button whenever they judged that a significant segment of the action was completed. In some trials subjects were asked to focus on *longer scale segments* while in other trials subjects were asked to focus on *shorter scale segments* (the directions were intentionally left vague like this in the experiments to leave subjects free to segment the scenes however they deemed appropriate). What was found was that there was a significant amount of within and across individual reliability in how the complex events were segmented into both long and short scale segments. Furthermore, it was also found that there was an alignment of segment boundaries across the short scale and long scale trials. That is, subjects who were only ever asked to make long scale segmentation judgments would choose segments the boundaries of which would coincide with a subset of the segment boundaries chosen by people who were only ever asked to make short scale segmentation judgments.¹⁸⁰ In this way it seems as though the construction of event models does not occur only at a single time scale at a time but rather involves the segmentation of events into discrete parts at multiple time scales in parallel.

¹⁸⁰ Now, of course one explanation, that takes much of the mystery away from these findings, is that the actual events in the world that the subjects were observing would themselves typically involve sub-events that would coincide with longer scale events. For instance, the action of placing the mop against the wall (a short scale segment) would naturally align itself with the end of the longer scale event of mopping the floor. So, since the stimuli were naturalistic (e.g. videos of research assistants performing mock tasks), the actual temporal structure and make up of the events in the stimuli may have forced this coincidence.

If we look at the specific theories of event perception in the literature what we find are varying accounts that typically focus on how non-temporal information from the various sensory systems comes to be integrated into a coherent and complex event. In a way, what we find is a complex form of intersensory binding. Yet in other areas of the cognitive sciences, apart from what is found in the literature on event models, we do find some limited attempts to explain how temporal information from multiple modalities comes to be integrated, however, this research tends to focus almost exclusively on two aspects of this integration. First, there is the cross-modal perception of synchrony or simultaneity. This literature focuses on what has been called *the temporal window of integration* (TWI). It is known that for two stimuli to be perceived as being simultaneous it needn't be the case that those stimuli impact the sensory receptors simultaneously, instead, they must arrive at the sensory receptors within a particular window within which they can be integrated. Interestingly, the size of the TWI is sensitive to a number of different stimulus factors, perhaps most notably are the influences of perceived distance on the TWI and whether or not the stimuli are *thematically coherent* (i.e. audiovisual speech where the mouth movements correspond with the sound are capable of being perceived as simultaneous over a much greater TWI than a random visual and auditory event)¹⁸¹. The second area of research, that is far less researched than the TWI, focuses on *duration summation*. In this literature the temporal information being integrated across modalities is not simultaneity but is rather the duration of a multimodal chain of stimuli (Takahashi & Watanabe, 2015).¹⁸²

¹⁸¹ See literature discussed in Callender (2008) and Montemayor (2013) for a review of the literature on the TWI.

¹⁸² In addition to these two areas of research that explicitly focus on the crossmodal integration of temporal information, we also have an enormous body of research that seems to implicitly tap into this very same capacity. One of the primary means of investigating people's ability for temporal perception comes from *reproduction tasks*. In these tasks subjects are expected to perceive some stimuli in some modality, and then they must *reproduce* the temporal pattern of the stimuli (e.g. duration or rhythm) through tapping a button. This entire experimental paradigm relies on the ability for the motor systems to accurately utilize the temporal information encoded in perception. However, this doesn't explicitly require that temporal information from across various time keeping devices be integrated. Rather it merely requires that the motor system have access to the perceptual information. It is only tasks that would require the temporal information in perception and motor activity to be integrated for some further task.

In both of these cases we have some research that is being conducted on how temporal information from various components of the perceptual system comes to be integrated for further use. Some authors have attempts to explain this integration by positing further clock-like mechanisms that track the temporal structure of more proximal sensory activity.¹⁸³ For example, the perception of simultaneity is not the result of simultaneous sensory receptor activity, but is instead the product of simultaneous sensory processing later in the perceptual system. Similarly, such a mechanism would explain the cross-modal perception of duration through a process of tracking the duration of sensory processing. However, it's not clear that this sort of approach would succeed. As we saw in the last chapter, there is no simple story to be told about how the temporal content of perception relates to the temporal structure of perception. Yet, this sort of *higher-level clock mechanism* would only be sensitive to the temporal structure of perception and not the temporal content of perception. But it is sensitivity to the temporal *content* of perception that is required. A simple higher-order clock mechanism, by itself, will not provide us with a means of accessing the temporal content of perception in general (e.g. noting the timing of a perceptual process would not allow us to access temporal information, such as the predictive elements in trajectory estimation, that is not mirrored by the temporal structure of experience itself). Some other story is needed and it has to be the case that the story is given in a way that the process of event model construction has access to the *temporal information encoded in perception and not just the temporal structure of perception itself*.

5.6. The units of temporal perception

How then might we explain this sort of integration? It is here where we finally get to see the relevance of the earlier discussion on the unitization of time. The proposal is that the perceptual system

¹⁸³ See the brief discussion of asynchrony detectors in (Navarra et al., 2012).

does something that parallels what we as a society did in order to construct our system of temporal units. Let's first sketch out what it would mean for the perceptual system to unitize time.

Unit-laden representations of time present time under a particular type of *mode of presentation* that represents time in terms of (or in relation to) a standardized unit, which is in turn defined by some *chosen object*. When I, in English, report that the duration of a tone is 500ms by using the terms 'that tone is 500ms long', I can be understood as reporting that *the ratio of the tone to the standard second is 0.5* or that *the tone has a duration that is assigned a value of 0.5 on the numerical extensive system that assigns the standard second a value of 1*. As a speaker of English it may be that neither of these statements strike the right chord as properly capturing the content or meaning of my utterance that the tone is 500ms long, but nevertheless these characterizations of the statement capture the correct truth conditions of what I have said.¹⁸⁴ This sort of unit-laden way of describing the length of the tone is rather different than how I might express the duration of the tone in some demonstrative way (e.g. by saying the tone has *that duration* while somehow ostending to some event that is 500ms long). While both speech acts will report the same facts, it is clear that they do so in rather different ways. Event models, I argue, represent time in a manner that is closer to the non-demonstrative account.

Central to there being a unit-laden representation of time is the appeal to some *chosen object* in establishing a unit standard. And further, on the basis of this chosen object a representational system is established in virtue of which other time keeping mechanisms are calibrated and interpreted in terms of. In what follows, I will argue that the construction of event models likely involves just this sort of appeal to units.¹⁸⁵

¹⁸⁴ In fact, however, these sorts of re-descriptions of the meaning of my statement aren't too far off. As anyone who has tried to teach elementary school children their units of measurement, one way to teach someone the meaning of 'a meter' is to tell them that *one meter is the distance that you would get by placing 100 1-cm cubes in a straight line*.

¹⁸⁵ The story to be told here is one that ultimately requires empirical vindication. The sketch being provided is hoped to not only provide some theoretical space to explain the sort of integration I have in mind, but also to ultimately spur direct empirical investigation.

Something that is known about both duration and event perception is that in both cases the perceptual system makes use of memories of similar events in the representation of the temporal structure of events. Perhaps the most salient appeal to memories in the perceptual representation of time comes as part of the scalar expectancy theory (SET). Recall that SET explains temporal perception by appealing to a system with three main components. First, there is the pacemaker-accumulator system. Second, there is a memory system that stores *temporal reference memories*, which are memories of the duration of past stimuli. Third, there is a decision process that compares the readout of the pacemaker to the stored temporal reference memories in order to form a judgment about the duration of some event. Let's focus on the second component, the temporal reference memories (TRM). TRM are posited to play two roles in the theory. First, they play a role in adjusting the scalar properties of the temporal representation over varying time scales (Jones & Wearden 2003). Discrimination is best for stimuli that are closest in duration to the temporal reference memories, or temporal *standards* as Wearden calls them. TRM also play a more direct role in many temporal perception tasks in that subjects are often asked to make comparative judgments concerning the duration of some standard stimuli and a target stimulus. This requires that the temporal properties of the standard stimuli are retained in working memory, otherwise, subjects would be unable to perform the relevant comparison task. Interestingly though, subjects do not retain the specific duration of stimulus standards in memory but instead store the average duration of similar reference standards (Levy et al 2015).

While the role of TRM just described was given within the framework of SET the data on which the theory is built is something that any theory of temporal perception must accommodate. If this data shows that perceptual judgments about duration are sensitive to the average duration of temporal standards in these experiments, then any theory of temporal perception must make space for memories of this sort. As Jones and Wearden put it, "the first [role of TRM] is to provide a temporal reference for behavioral stability by storing "important" times." (Jones & Wearden 2003, p. 3), and any account of

temporal perception will have to explain how behavioral stability in the temporal coordination of activity can be achieved.

Event perception is also sensitive to stored temporal memories, but in this case these memories are referred to as *event schemata* and they encode the structure of typical and familiar event types. For instance, more familiar event types (e.g. mopping the floor) are segmented in ways that correspond to known subevents for that event type (e.g. wringing the mop) (Zacks et al 2007; Zacks & Tversky 2001).

In both cases, the way in which we perceive the temporal structure of the world around us is sensitive to not just the incoming sensory stimulation but is also directly influenced by stored temporal memories for typical or familiar event or stimulus types. Since event models include information about event duration it's likely the case that insofar as TRM are involved in duration perception they are also involved in the construction of event models. At this point we actually have enough to build the initial case that the perceptual system unitizes time. The TRMs involved in duration perception serve as standards relative to which other events are measured. The TRMs provide a standard by which to compare the readings of the various peripheral time keeping mechanisms. In this way, event models provide a framework, and a code given by the TRM, for the interpretation of the various time keeping mechanisms. All that is required for the integration of temporal information is that the various peripheral time keeping mechanisms be comparable to the familiar duration that is stored in the form of a TRM.

However, we can go further in our defense of the unit-laden account of temporal perception. The next part of the story comes from the way in which the perceptual system optimizes its representations of the temporal structure of events to capture the most reliable sources of information. Consider a parallel case from spatial perception. When people receive conflicting cues as to the location of an event through multiple modalities (e.g. conflicting spatial cues from audition and vision), the perceptual system makes a decision to bias one source of information over the other in order to produce a single coherent estimation as to the location of some event. In most cases, when given conflicting information through multiple

modalities, there will be a bias towards relying more heavily on the spatial information given through vision (Choe et 1975).¹⁸⁶ However, it's been shown that typical biasing towards visual information is not because vision is always going to be given the most weight, but rather reflects an estimation of the reliability of the various information sources. If white noise is introduced into the visual signal, thereby degrading the reliability of vision as a source of spatial information, then the typical bias towards visual information vanishes (Alais & Burr, 2004). In typical cases, the bias for vision is explained as a result of vision having the highest spatial acuity of all of the sensory systems. The typical case is simply a limiting case of the reliability estimation explanation.

Almost exactly parallel findings have been found for temporal perception and audition. In typical cases, the temporal localization of events will resolve conflicting temporal cues by biasing auditory sources of information. However, if white noise is introduced to the auditory signal, then this bias is lost in favor of whatever information source is judged to be most reliable (Shi et al., 2013). The upshot of all of this is that the perceptual system in making crossmodal judgments of both temporal and spatial structure makes use of whatever source it takes to be the most reliable. In this way, the perceptual system calibrates the interpretation of the various time keeping mechanisms to be in line with whatever time keeping mechanism it deems to be most reliable.

Here we find another parallel between the perception of time and the way in which time has been unitized as part of our social time keeping practices. In both cases, our temporal measurements are biased in a way that favors the readings given by those mechanisms that we deem to be most reliable. In the construction of UTC, the calculation of time's passage was the product of a weighted mean of a great number of individual time keeping mechanisms in which the weighting was a response to the calculated reliability of the individual clocks. Similarly, in the perception of time there is a biasing of the readings of

¹⁸⁶ This is the reason why even when movie speakers are located behind your seat, you still perceive speech as coming from in front of you where the visual image is.

the various time keeping mechanisms in the favor of those clocks that are judged as being most reliable (i.e. the least noisy sensory signal).

In our social time keeping practices a chosen event (i.e. the particular energy transitions of the caesium-133 atom) is taken to establish a temporal standard relative to which other time keeping mechanisms are calibrated and which is used in the interpretation of various time keeping mechanisms. Similar resources seem to not only be available to the perceptual system but in fact seem to be used by the perceptual system in the construction of event models. The perceptual system takes the most reliable temporal representations of event types as standards by which the other peripheral time keeping mechanisms can be interpreted and which serve as a means of calibrating the readings of the other time keeping mechanisms. These reference memories provide us with a common-code by which to interpret the readings of the type distinct time keeping mechanisms employed in perception.

Even introspectively there is something to this. Consider a temporal analog of Peacocke's expert eye-baller (i.e. the person who upon looking at something could easily tell you its length). When we try and estimate the duration of some event, often, what we might do is look to see if there is another event the duration of which we are fairly certain of, and then use that event as a standard by which to judge the duration of the target event. This can happen at both longer and shorter time scales. I might judge that a pause in the conversation had a duration of 20 seconds by using the stored representation of the duration of some other event as a standard by which to make my comparison. For longer scale events, it's helpful to call up memories of similarly long scale events. For shorter scale events, the parallel is also true, in that you call up memories of shorter scale events. The proposal however is that something like this gets done automatically in order for the integration of temporal information.

Now, the parallels do not extend to cover every aspect of how temporal information is integrated in psychological and social time keeping but of course, who would have thought that the parallels would encompass everything? The chosen event in our cultural time keeping practices, the energy transitions of

the caesium-133 atom, were chosen through a complex and iterative process of scientific discovery. Temporal standards are chosen through a process that uses best scientific theories to choose an event type that ought to be highly reliable and which can be in principle measured. Nothing of this sort happens in perception. Rather, the perceptual system adopts some event (likely given some particular task demands or judgments about the regular behavior of events in the environment)¹⁸⁷ and produces a representation of that average duration of events of that type, through whatever means allows it to form the most reliable memory representations. In this way, the choice of chosen event for establishing a system a common code is less up to theoretical decisions and is likely driven by the individual's own history of experiences along with particular factors that influence the reliability of TRM.

Nevertheless, the proposal here is that by unitizing time in this way, the brain has the resources needed to overcome the integration problem that stood in the way of explaining the unity of time as it is presented to us in perception and cognition. Once the perceptual systems involved in event perception learn how to interpret the individual "markings" of the peripheral time keeping mechanisms by relating them to a single common code, then there is nothing to prevent the integration of temporal information from happening. And, once the common code is established, then we will have a unitization of time.

5.7. Conclusion

The emphasis of the last chapter was the fragmentary nature of temporal perception. This chapter, however, focuses on the way in which time is presented to us in perception and cognition as a unified and seamless framework within which the events around us take place. What makes this *unity of time* so interesting is that it is supposed to arise from a system that utilizes a variety of type-distinct time keeping

¹⁸⁷ It's possible that the particular *chosen object* in temporal perception isn't even a worldly event but is instead some internal event that is monitored, as there are a number of internal processes with regular rhythms, and this internal process is used as a standard. Perhaps it is here where the variety of cortical or bodily rhythms enters into the picture of temporal perception.

mechanisms to gather information about time. Somehow, the perceptual system must be capable of integrating all of this temporal information despite the fact that the peripheral time keeping devices encode time in “idiosyncratic” ways.

The question was then how does the brain do this. Interestingly, the means to understanding how the brain might accomplish some task took us through a literature that is more common in the history and philosophy science than it is in the philosophy of mind or neuroscience. We as a society faced a similar sort of integration problem in that we needed to find a way of standardizing out time keeping practices so that the various time keeping devices in use could be integrated to form a coherent picture of the temporal structure of our world. As a society we adopted a system whereby we unitize time to overcome this problem. It is through representing time under this particular unit-laden mode of presentation that we came to interpret the various time keeping devices we use in society through a single common code.

This strategy that we as a society adopted, I argued, is one that is available to the perceptual system in the construction of event models. Perception integrates the temporal information encoded in the peripheral time keeping mechanisms of perception through the representation of a unit-standard, a chosen event, that serves as a reference point for the interpretation of the other temporal information encoded in the brain. In this way, time is unified in the construction of our complex event models.

Chapter 6: Conclusion

Time is one of those aspects of the environment that has puzzled philosophers since antiquity. Even more so today with modern advances in physics time is becoming even more mysterious as the scientific picture of time is departing more and more from our common sense or intuitive understanding of time (Callender, 2008; Ismael, Forthcoming, 2010). Understanding how it is that we come to experience time is not only crucial to having a complete understanding of how minds fit into the natural world but more importantly it connects up with our overall understanding of the physical world altogether. It is for this reason, partly, that philosophers have for quite some time been interested in trying to understand how it is that we can have temporal experiences of a dynamic and changing world. However, as argued in the beginning of the dissertation, focusing on just the conscious experience of time is limiting. It's crucial for animals that they be able to keep track of time in a variety of ways, many of which never enter into our conscious experience of the world.

At the beginning of this dissertation I laid out a particular *explanatory* problem that I called *the temporal coordination problem* that was meant to broaden our horizons away from merely focusing on conscious experience. All animals need to coordinate their behaviors with the temporal structure of the events in the environment around them. The difficulty arises in trying to *explain how* this coordination happens. What we've covered in the preceding chapters is just a small portion of how animals, including humans, come to coordinate their behaviors with the daily temporal patterns in their environment as well as the temporal structure of their environment at the relatively short time scale of milliseconds to several seconds.

While the question of how minds relate to time has always captured the attention of philosophers and cognitive scientists, it was the aim of this thesis to show how the empirical and philosophical literature on this topic should be brought together to produce a richer and more complete philosophical and empirical understanding of how the mind relates to time. Central to this discussion was noticing the

two roles that time has in our mental lives. First, our mental lives exist and unfold in time and have their own complex temporal structure. Second, much of our mental lives is directed at time itself.

That time plays these two roles in our mind hasn't gone unnoticed by either philosophers or cognitive scientists. In fact, much of the literature on temporal representation has focused on how these two roles for time interact, with many researchers arguing that in some way the temporality of the mind itself will have a role in explaining how mental states come to be directed at or about time.

The interest, however, in focusing on how the mind relates to time isn't limited to just this very topic. It extends to give us some inroads to understanding much more sweeping questions about the mind. Time is so central to every animal that needs to navigate its environment, and time is also such a unique and *abstract* aspect of our environment, that by understanding how the ability to represent time may be so widespread in the animal kingdom we get a foothold in understanding how minds come to be about their environments in general. That is, if we can understand how the representation of time might be given a naturalistic understanding, then we are much closer to understanding how the mind in general fits within our naturalistic understanding of the world.

Let me more specifically summarize the main conclusions of chapters 3, 4, and 5.

In chapter 3 I argued that many animals, including humans, possess a genuine sense of time that provides them with information about the approximate time of day. One of the challenges in providing an account of how animals can coordinate their behavior with or be sensitive to time is that time is simply an odd aspect of our environment. It isn't something concrete that we can interact with, point to, or that has any causal influence on events in the world. The last one of these, that time isn't causally efficacious, is a feature of time that has been taken by a number of authors to cause problems with providing an explanation for how we mentally represent time. In particular, we saw that a number of authors have used this feature of time as a reason for why we cannot have a genuine sense of time, as it is only through causal interaction with the environment that our sensory systems supposedly come to track or gather

information about the world. Furthermore, the role that time has in our mental lives as a framework within which the individual deliverances of the sensory systems are integrated has also given people pause in accepting that we may genuinely have a sense of time. Time, in playing this role of providing a framework for the integration of the deliverances of the individual senses, would not provide a suitably direct means of representing the environment.

In chapter 3 I argued that neither of these concerns about the nature of time itself or the role of time as an integrative framework in the mind give us good reason for doubting that we have a genuine sense of time. By properly understanding the role that attributions of sensory systems play in our folk and scientific understanding of the mind we see that what is key to something's being a sensory system is that it *directly gathers information about* and *tracks in realtime* some aspect of the world. Importantly, though, it was argued that while in most cases a causal connection between the sensory systems and the world is needed for the sensory systems to keep track of the world, in the case of time, specifically in the case of the circadian system's ability to keep track of the approximate time of day, this causal connection isn't needed. The circadian systems keep track of time and provide us, as well as many other animals, with a genuine sense of time, despite the fact that time does not cause the circadian system to be in any of its states.

The key to understanding how the sense of time can come to keep track of the approximate time of day despite not having any states that are the casual effects of time is through noting how the circadian system, like clocks in general, not only represent time but also *participate in time*. It is through the way that the circadian oscillator develops through time, and possesses a nearly 24-hour period, that it comes to have its information gathering capacities. Yet, the role that this unfolding in time has in explaining how the circadian system comes to represent time does not presuppose any sort of resemblance theory. We have a sense of time, but it has only a limited role in explaining how we come to coordinate our behaviors with time.

In chapters 4 and 5 I turned our attention to what philosophers and cognitive scientists typically mean when they speak of “temporal perception”. We were concerned with how animals come to perceive through the various sensory modalities the fine-grained temporal structure of the events around them, at a scale spanning milliseconds through several seconds. In chapter 4 I argued for what I called *the fragmentary model of temporal perception*. Many researchers, philosophers and scientists alike, have typically attempted to explain the perception or experience of time through the operation of a single unified type of explanation. However, the evidence simply does not bear out this assumption. ‘Temporal perception’ does not pick out a single unified capacity, but rather invites of many polysemous interpretations. Each of these uses of ‘temporal perception’ picks out some capacity that animals possess to track and represent some aspect of the temporal world around them.

The interest of the fragmentary model, especially as it was discussed in chapter 4, had to do with how it changes many of the debates that have been central to philosophical discussions of temporal perception. Perhaps the central debate in the philosophical literature, since at least the early modern empiricists, has been over *the mirroring constraint*. Do the temporal contents of experience *mirror* the temporal structure of experience itself? An assumption in this debate has always been that it will either be the case that temporal experience does or does not satisfy the mirroring constraint. It’s an all or nothing affair. However, it is only through direct inspection of the mechanisms that underpin our temporal perception, and not through introspection or appeals to general principles about the brain’s operations, that we can settle the debate over mirroring. Taking this tactic, I argued that it turns out that while some of the mechanisms that underpin our temporal perception satisfy the mirroring constraint, others do not. In explaining how we come to perceive time, there is no single story to be told about how the temporal contents of experience relate to the temporal structure of experience itself. In this way, the current neuroscientific literature on temporal perception gives us insight into incredibly long lasting philosophical debates concerning how time relates to the mind.

While chapter 4 emphasized the fragmentary nature of temporal perception and how this impacted central philosophical debates, chapter 5 instead focused on the way in which time is presented to us in perception and cognition as a unified and seamless framework within which the particular events we grasp are located. There is a unity to time as it appears in perception and cognition. However, given the fragmentary model of temporal perception, this unity is actually rather difficult to explain. Somehow the temporal information, at various time scales and gathered through various modalities, must be integrated to underpin this unity to time. However, the empirical literature still doesn't have an explanation for how this might even occur. Given that the various mechanisms that gather this temporal information often employ radically different representational strategies from one another a straightforward integration of the information gathered by these various mechanisms is not possible. Further since these various mechanisms rely on there being different relations between the temporal contents and the temporality of the mechanisms themselves we cannot appeal to the temporality of the mental processes involved in this sort of integration to do the relevant work. Something else is needed.

It was in this chapter that I argued that looking at literature that is more familiar to historians and philosophers of science working on measurement can actually give us insight into how the brain might be performing some operation. Specifically, it is by drawing parallels between the temporal integration problem that the brain faces and how we as a society have come to adopt a system of standardized units in our time keeping practices that we can argue that the brain overcomes its temporal integration problem by unitizing time. While this proposal is ultimately one that can enjoy direct empirical investigation, the fact that the literature from the history and philosophy of science can propel research in the cognitive science of temporal representation is interesting in and of itself. That it also provides us with a viable framework for understanding how the unity of time might be produced by the brain is the ultimate pay off of this chapter.

In the introduction of this dissertation, and throughout the body of the text, many of the arguments centered around the philosophical literature on mental representation. Importantly, in order to provide a naturalistic understanding of the mind it is widely accepted that we must be able to provide a naturalistic explanation for how mental representations can come to have their content. In the introduction I argued that while many in the philosophical literature attempt to find *the theory* that explains how mental representations have their content, we needn't in fact think that there will only be one theory of content that accounts for all of the content attributions that we make in theorizing about the mind. That we have mental representations is one thing, and there is a general unity to what mental representations are (i.e. they are mental particulars with semantic properties), but there needn't be only one means by which mental representations come to have their contents.

What we've seen, though, is that our attributions of content and our acceptance of theories of content can be fruitfully informed by what we know about the underlying neural machinery that is involved in the representation of time. Repeatedly, I have been appealing to information-theoretic accounts of content. The reason for this is that the account provides a generalizable theoretical framework for understanding how a variety of different mechanisms may come to represent some aspect of the world. In particular, the theory provides us with the tools for understanding how such a wide range of animals can come to mentally represent such an *abstract* aspect of the environment like time. It is also sufficiently general that it allows for cases in which the temporal structure of our mental processes themselves plays a role in explaining how our mental states come to represent time and for those cases in which the temporality of our mental processes plays no role in explaining how we come to represent time. While it may be the case that we needn't accept any one theory of content in theorizing about mental representation, it is nevertheless the case that the information-theoretic account developed by Dretske (1981) provides us with an astonishingly powerful tool for understanding how animals can coordinate their behaviors with and represent the temporal structure of the world around them.

The thesis had to be limited in its scope. Timing is so important to almost every task any animal performs and the temporal structure of our mental processes and the world around us is incredibly complex. To have provided an exhaustive theory of how time in general impacts the mind would be an immense project. What we've seen in this dissertation, however, is an important slice of this overall explanatory project. How the mind comes to form its initial representations of time will surely be of importance for understanding how more sophisticated time keeping capacities (e.g. temporal concepts, temporal reasoning, memory) allow for animals to coordinate their behaviors with even more complex aspects of the temporal world around them. The remainder of the temporal coordination task remains to be explained. However, for the time being, we can content ourselves with having gained a more complete understanding of some portion of the temporal coordination problem, and seeing how our understanding of temporal representation is furthered by bringing together the philosophical and empirical literature on the topic.

Bibliography

- Akins, K. (1996). Of Sensory Systems and the “Aboutness” of Mental States. *Journal of Philosophy*, 93(7), 337–372.
- Akins, K. (2014). Black and White and Color. In *Consciousness Inside and Out: Phenomenology, Neuroscience and the Nature of Experience* (Vol. 6, pp. 173–223). Dordrecht.
- Akins, K. A., & Hahn, M. (2014). More Than Mere Colouring: The Role of Spectral Information in Human Vision. *British Journal for the Philosophy of Science*, 65(1), 125–171.
- Alais, D., & Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Current Biology: CB*, 14(3), 257–262.
- Alais, D., & Cass, J. (2010). Multisensory Perceptual Learning of Temporal Order: Audiovisual Learning Transfers to Vision but Not Audition. *PLOS ONE*, 5(6), e11283.
- Andersen, H. (2014). The Development of the “Specious Present” and James’ Views on Temporal Experience. In D. L. V. Arstila (Ed.), *Subjective Time: the philosophy, psychology, and neuroscience of temporality* (pp. 25–42). MIT Press.
- Andersen, H., & Grush, R. (2009). A Brief History of Time-Consciousness: Historical Precursors to James and Husserl. *Journal of the History of Philosophy*, 47(2), 277–307.
- Aristotle. (2004). *De Anima (On the Soul)*. Penguin UK.
- Baars, B. (2005). Global workspace theory of consciousness: toward a cognitive neuroscience of human experience. *Progress in Brain Research*, (150), 45–53.
- Barsalou, L. W. (1999). Perceptual Symbol Systems. *Behavioral and Brain Sciences*, 22(4), 577–660.
- Bechtel, W. (2005). Explanation: A Mechanist Alternative. *Studies in History and Philosophy of Biol and Biomed Sci*, 36(2), 421–441.
- Bechtel, W. (2011a). Mechanism and Biological Explanation. *Philosophy of Science*, 78(4), 533–557.

- Bechtel, W. (2011b). Representing Time of Day in Circadian Clocks. In A. Newen, A. Bartels, & E.-M. Jung (Eds.), *Knowledge and Representation*. CSLI Publications and Paderborn.
- Bechtel, W. P. (1998). Representations and Cognitive Explanations: Assessing the Dynamicist Challenge in Cognitive Science. *Cognitive Science*, 22(3), 295–317.
- Beck, J. (2012). The Generality Constraint and the Structure of Thought. *Mind*, 121(483), 563–600.
- Beck, J. (2014). Analogue Magnitude Representations: A Philosophical Introduction. *The British Journal for the Philosophy of Science*, axu014.
- Benovsky, J. (2012). The Causal Efficiency of the Passage of Time. *Philosophia*, 40(4), 763–769.
- Bernal, S. (2005). Object Lessons: Spelke Principles and Psychological Explanation. *Philosophical Psychology*, 18(3), 289–312.
- Biederman, I. (1987). Recognition-by-Components: A Theory of Human Image Understanding. *Psychological Review*, 94(2), 115–117.
- Block, N. (1995). On a Confusion About a Function of Consciousness. *Brain and Behavioral Sciences*, 18(2), 227–247.
- Bolles, R. C., & Moot, S. A. (1973). The rat's anticipation of two meals a day. *Journal of Comparative and Physiological Psychology*, 83(3), 510–514.
- Boroditsky, L. (2000). Metaphoric structuring: Understanding time through spatial metaphors. *Cognition*, 75(1), 1–28.
- Boroditsky, L. (2001). Does language shape thought?: Mandarin and English speakers' conceptions of time. *Cognitive Psychology*, 43(1), 1–22.
- Boroditsky, L. (2011). How Languages Construct Time. In S. Dehaene & E. Brannon (Eds.), *Space, Time, and Number in the Brain* (pp. 333–341). Elsevier.
- Boroditsky, L., & Prinz, J. (2008). What Thoughts Are Made Of. In G. R. Semin & E. R. Smith (Eds.), *Embodied Grounding: Social, Cognitive, Affective, and Neuroscientific Approaches* (pp. 98–115). Cambridge University Press.

- Boroditsky, L., & Ramscar, M. (2002). The roles of body and mind in abstract thought. *Psychological Science*, 13(2), 185–189.
- Both, C., G. Bijlsma, R., & E. Visser, M. (2005). Climatic effects on timing of spring migration and breeding in a long-distance migrant, the pied flycatcher *Ficedula hypoleuca*. *Journal of Avian Biology*, 36(5), 368–373.
- Braddick, O., & Atkinson, J. (2011). Development of human visual function. *Vision Research*, 51(13), 1588–1609.
- Bratzke, D., Seifried, T., & Ulrich, R. (2012). Perceptual learning in temporal discrimination: asymmetric cross-modal transfer from audition to vision. *Experimental Brain Research*, 221(2), 205–210.
- Breska, A., & Ivry, R. B. (2016). Taxonomies of timing: where does the cerebellum fit in? *Current Opinion in Behavioral Sciences*, 8, 282–288. <https://doi.org/10.1016/j.cobeha.2016.02.034>
- Broad, C. D. (1923). *Scientific Thought*. Routledge and Kegan Paul.
- Brown, A. M. (1990). Development of visual sensitivity to light and color vision in human infants: a critical review. *Vision Research*, 30(8), 1159–1188.
- Buhusi, C. V., & Meck, W. H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. *Nature Reviews. Neuroscience*, 6(10), 755–765.
- Bullock, T. H., & Cowles, R. B. (1952). Physiology of an Infrared Receptor: The Facial Pit of Pit Vipers. *Science*, 115(2994), 541–543. <https://doi.org/10.1126/science.115.2994.541-a>
- Buonomano, D. V. (2000). Decoding temporal information: A model based on short-term synaptic plasticity. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 20(3), 1129–1141.
- Buonomano, D. V. (2005). A Learning Rule for the Emergence of Stable Dynamics and Timing in Recurrent Networks. *ResearchGate*, 94(4), 2275–83.
- Buonomano, D. V., & Karmarkar, U. R. (2002). How do we tell time? *The Neuroscientist: A Review Journal Bringing Neurobiology, Neurology and Psychiatry*, 8(1), 42–51.

- Buonomano, D. V., & Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, 10(2), 113–125.
- Burr, D. C., Ross, J., Binda, P., & Morrone, M. C. (2011). Chapter 12 - Saccades Compress Space, Time, and Number. In *Space, Time and Number in the Brain* (pp. 175–186). San Diego: Academic Press.
- Burr, D., Tozzi, A., & Morrone, M. C. (2007). Neural mechanisms for timing visual events are spatially selective in real-world coordinates. *Nature Neuroscience*, 10(4), 423–425.
- Butts, D. A., Weng, C., Jin, J., Yeh, C.-I., Lesica, N. A., Alonso, J.-M., & Stanley, G. B. (2007). Temporal precision in the neural code and the timescales of natural vision. *Nature*, 449(7158), 92–95.
- Byrne, A. (2011). Transparency, Belief, Intention. *Aristotelian Society Supplementary Volume*, 85(1), 201–221.
- Cai, M., Stetson, C., & Eagleman, D. M. (2012). A Neural Model for Temporal Order Judgments and Their Active Recalibration: A Common Mechanism for Space and Time? *Frontiers in Psychology*, 3.
- Callender, C. (2008). The Common Now. *Philosophical Issues*, 18(1), 339–361.
- Carey, S. (2009). *The Origin of Concepts*. Oxford University Press.
- Carnap, R. (1937). *The Logical Syntax of Language*. London, K. Paul, Trench, Trubner & Co., Ltd.
- Carruthers, P. (2006). *The Architecture of the Mind: Massive Modularity and the Flexibility of Thought*. Oxford: Clarendon Press.
- Carvalho, F. M., Chaim, K. T., Sanchez, T. A., & Araujo, D. B. de. (2016). Time-Perception Network and Default Mode Network Are Associated with Temporal Prediction in a Periodic Motion Task. *Frontiers in Human Neuroscience*, 10.
- Casasanto, D. (2008). Who's afraid of the big bad Whorf? Crosslinguistic differences in temporal language and thought. *Language Learning*, 58, 63–79.

- Casasanto, D., & Boroditsky, L. (2008). Time in the Mind: Using Space to Think About Time. *Cognition*, 106(2): 579-593.
- Casasanto, D., Fotakopoulou, O., & Boroditsky, L. (2010). Space and Time in the Child's Mind: Evidence for a Cross-Dimensional Asymmetry. *Cognitive Science*, 34(3), 387–405.
- Cassone, V. M., Paulose, J. K., Whitfield-Rucker, M. G., & Peters, J. L. (2009). Time's arrow flies like a bird: two paradoxes for avian circadian biology. *General and Comparative Endocrinology*, 163(1–2), 109–116.
- Cassone, V. M., & Westneat, D. F. (2012). The bird of time: cognition and the avian biological clock. *Frontiers in Molecular Neuroscience*, 5, 32.
- Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
- Chang, H. (2004). *Inventing Temperature: Measurement and Scientific Progress*. OUP Usa.
- Chemero, A. (2000). Anti-Representationalism and the Dynamical Stance. *Philosophy of Science*, 67(4), 625.
- Chiao, C. C., Cronin, T. W., & Marshall, N. J. (2000). Eye design and color signaling in a stomatopod crustacean *Gonodactylus smithii*. *Brain, Behavior and Evolution*, 56(2), 107–122.
- Choe, C. S., Welch, R. B., Gilford, R. M., & Juola, J. F. (1975). The “ventriloquist effect”: Visual dominance or response bias? *Perception & Psychophysics*, 18(1), 55–60.
- Chuard, P. (2011). Temporal Experiences and Their Parts. *Philosophers' Imprint*, 11(11).
- Church, R., & Broadbent, H. (1990). Alternative Representations of Time, Number, and Rate. *Cognition*, 37(1–2), 55–81.
- Church, R. M., & Meck, W. H. (1984). The numerical attribute of stimuli. *ResearchGate*, 445–464.
- Clark, A. (2001). Visual Experience and Motor Action: Are the Bonds Too Tight? *Philosophical Review*, 110(4), 495–519.

- Clavagnier, S., Falchier, A., & Kennedy, H. (2004). Long-distance feedback projections to area V1: Implications for multisensory integration, spatial awareness, and visual consciousness. *Cognitive, Affective, & Behavioral Neuroscience*, 4(2), 117–126.
- Clayton, N. S., & Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, 395(6699), 272–274.
- Coull, J. (2011). Discrete Neuroanatomical Substrates for Generating and Updating Temporal Expectation. In S. Dehaene & E. Brannon (Eds.), *Space, Time, and Number in the Brain*. Elsevier.
- Craver, C. F. (2007). *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford University Press, Clarendon Press.
- Craver, C. F., Graham, B., & Rosenbaum, R. S. (2014). Remembering Mr B. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior*, 59, 153–184.
- Craver, C. F., Kwan, D., Steindam, C., & Rosenbaum, R. S. (2014). Individuals with episodic amnesia are not stuck in time. *Neuropsychologia*, 57, 191–195.
- Crick, F., & Koch, C. (1990). Toward a Neurobiological Theory of Consciousness. *Seminars in the Neurosciences*, 2, 263–275.
- Crick, F., & Koch, C. (2003). A framework for consciousness. *Nature Neuroscience*, 6(2), 119–126.
- Cronin, T. W., Caldwell, R. L., & Marshall, J. (2001). Sensory adaptation: Tunable colour vision in a mantis shrimp. *Nature*, 411(6837), 547–548.
- Dainton, B. (2008). The Experience of Time and Change. *Philosophy Compass*, 3(4), 619–638.
- Dainton, B. (2011). Time, Passage and Immediate Experience. In C. Callender (Ed.), *Oxford Handbook of Philosophy of Time* (p. 382). Oxford University Press.
- Dainton, B. (2014). Temporal Consciousness. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2014). Retrieved from <http://plato.stanford.edu/archives/spr2014/entries/consciousness-temporal/>, September 2016.

- D L Schacter, C Y P Chiu, & Ochsner, and K. N. (1993). Implicit Memory: A Selective Review. *Annual Review of Neuroscience*, 16(1), 159–182.
- Dally, J. M., Emery, N. J., & Clayton, N. S. (2006). Food-Caching Western Scrub-Jays Keep Track of Who Was Watching When. *Science*, 312(5780), 1662–1665.
- DeCoursey, P. J., Walker, J. K., & Smith, S. A. (2000). A circadian pacemaker in free-living chipmunks: essential for survival? *Journal of Comparative Physiology. A, Sensory, Neural, and Behavioral Physiology*, 186(2), 169–180.
- Dehaene, S. (1997). *The Number Sense: How the Mind Creates Mathematics*. Oxford University Press.
- Dehaene, S. (2003). The neural basis of the Weber-Fechner law: a logarithmic mental number line. *Trends in Cognitive Sciences*, 7(4), 145–147.
- Dehaene, S. (2011). Section III Shared Mechanisms for Space, Time and Number? In S. Dehaene & E. Brannon (Eds.), *Space, Time and Number in the Brain* (pp. 119–121). San Diego: Academic Press.
- Dehaene, S., Bossini, S., & Giraux, P. (1993). The mental representation of parity and number magnitude. *Journal of Experimental Psychology: General*, 122(3), 371–396.
- Dehaene, S., Kerszberg, M., & Changeux, J.-P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences*, 95(24), 14529–14534.
- Dennett, D. C. (1992). *Consciousness Explained* (1 edition). Boston: Back Bay Books.
- Dennett, D. C., & Kinsbourne, M. (1992). Time and the Observer: The Where and When of Consciousness in the Brain. *Behavioral and Brain Sciences*, 15(2), 183–201.
- Descartes, R. (1960). *Discourse on Method, and Meditations*. New York, Liberal Arts Press.
- Descartes, R., & Cress, D. A. (1993). *Meditations on First Philosophy (Third Edition)*. Hackett Publishing.

- Dixon, N. F., & Spitz, L. (1980). The detection of auditory visual desynchrony. *Perception*, 9(6), 719–721.
- Dobkins, K. R., & Anderson, C. M. (2002). Color-Based Motion Processing Is Stronger in Infants Than in Adults. *Psychological Science*, 13(1), 76–80.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. MIT Press.
- Dretske, F. (1995). *Naturalizing the Mind*. MIT Press.
- Eagleman, D. M. (2008). Human time perception and its illusions. *Current Opinion in Neurobiology*, 18(2), 131–136.
- Eagleman, D. M., & Pariyadath, V. (2009). Is subjective duration a signature of coding efficiency? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1525), 1841–1851.
- Eckert, M. A., Kamdar, N. V., Chang, C. E., Beckmann, C. F., Greicius, M. D., & Menon, V. (2008). A Cross-Modal System Linking Primary Auditory and Visual Cortices. *Human Brain Mapping*, 29(7), 848–857.
- Eliasmith, C. (1996). The Third Contender: A Critical Examination of the Dynamicist Theory of Cognition. *Philosophical Psychology*, 9(4), 441–63.
- Falchier, A., Clavagnier, S., Barone, P., & Kennedy, H. (2002). Anatomical evidence of multimodal integration in primate striate cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 22(13), 5749–5759.
- Farah, M. J. (1990). *Visual Agnosia: Disorders of Object Recognition and What They Tell Us About Normal Vision*. MIT Press.
- Field, H. (1978). Mental Representation. *Erkenntnis*, 13(July), 9–61.
- Finnerty, G. T., Shadlen, M. N., Jazayeri, M., Nobre, A. C., & Buonomano, D. V. (2015). Time in Cortical Circuits. *The Journal of Neuroscience*, 35(41), 13912–13916.
- Fodor, J. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. MIT Press.

- Fodor, J. (1991). Special Sciences Jerry Fodor. In R. Boyd, P. Gasper, & J. D. Trout (Eds.), *The Philosophy of Science* (p. 429). MIT Press.
- Fodor, J. A. (1975). *The Language of Thought*. Harvard University Press.
- Fodor, J. A. (1980). Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology. *Behavioral and Brain Sciences*, 3(1), 63.
- Fodor, J. A. (1994). *The Elm and the Expert*. MIT Press.
- Fodor, J. A. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford University Press.
- Fodor, J. A. (2008). *Lot 2: The Language of Thought Revisited*. Oxford University Press.
- Fraassen, B. C. van. (2008). *Scientific Representation: Paradoxes of Perspective*. OUP Oxford.
- Frege, G. (1879). *Begriffsschrift: Eine Der Arithmetische Nachgebildete Formelsprache des Reinen Denkens*. L. Nebert.
- Frege, G. (1948). Sense and Reference. *The Philosophical Review*, 57(3), 209–230.
- Freyd, J. J., & Finke, R. A. (1984). Representational momentum. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10(1), 126–132.
- Freyd, J. J., & Jones, K. T. (1994). Representational momentum for a spiral path. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 20(4), 968–976.
- Frischhut, A. M. (2014). Diachronic Unity and Temporal Transparency. *Journal of Consciousness Studies*, 21(7–8), 34–55.
- Friston, K. J., Frith, C. D., Dolan, R. J., Price, C. J., Zeki, S., Ashburner, J. T., & Penny, W. D. (2004). *Human Brain Function*. Academic Press.
- Gallistel, null, & Gelman, null. (2000). Non-verbal numerical cognition: from reals to integers. *Trends in Cognitive Sciences*, 4(2), 59–65.
- Gallistel, C. R. (1990). *The Organization of Learning*. MIT Press.
- Gallistel, C. R. (1996). The Perception of Time. In K. Akins (Ed.), *Perception: Vancouver Studies in Cognitive Science* (Vol. 5). Oxford University Press.

- Gallistel, C. R., & Gelman, R. (1992). Preverbal and verbal counting and computation. *Cognition*, 44(1–2), 43–74.
- Gehring, W. J. (2014). The evolution of vision. *Wiley Interdisciplinary Reviews: Developmental Biology*, 3(1), 1–40.
- Geldard, F. A., & Sherrick, C. E. (1972). The Cutaneous “Rabbit”: A Perceptual Illusion. *Science*, 178(4057), 178–179.
- Gelder, T. V., & Port, R. F. (1995). Explorations in the Dynamics of Cognition. In T. V. Gelder & R. Port (Eds.), *Mind as Motion*. MIT Press.
- Gibbon, J. (1977). Scalar expectancy theory and Weber’s law in animal timing. *Psychological Review*, 84(3), 279–325.
- Gibbon, J., Church, R. M., & Meck, W. H. (1984). Scalar Timing in Memory. *Annals of the New York Academy of Sciences*, 423(1 Timing and Ti), 52–77.
- Goel, A., & Buonomano, D. V. (2014). Timing as an intrinsic property of neural networks: evidence from in vivo and in vitro experiments. *Phil. Trans. R. Soc. B*, 369(1637), 20120460.
- Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, 15(1), 20–25.
- Gracheva, E. O., Ingolia, N. T., Kelly, Y. M., Cordero-Morales, J. F., Hollopeter, G., Chesler, A. T., ... Julius, D. (2010). Molecular basis of infrared detection by snakes. *Nature*, 464(7291), 1006–1011.
- Grice, H. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and Semantics 3: Speech acts* (pp. 41–58). Academic Press.
- Grice, H. P. (1962). Some Remarks About the Senses. In R. J. Butler (Ed.), *Analytical Philosophy, First Series*. Oxford University Press.
- Grush, R. (2004). The Emulation Theory of Representation: Motor Control, Imagery, and Perception. *Behavioral and Brain Sciences*, 27(3), 377–396.

- Grush, R. (2005). Internal models and the construction of time: generalizing from state estimation to trajectory estimation to address temporal features of perception, including temporal illusions. *Journal of Neural Engineering*, 2(3), S209-218.
- Grush, R. (2006). How to, and How Not to, Bridge Computational Cognitive Neuroscience and Husserlian Phenomenology of Time Consciousness. *Synthese*, 153(3), 417–450.
- Grush, R. (2007). Time and Experience. In T. Muller (Ed.), *Philosophie der Zeit*.
- Grush, R. (2016). On the Temporal Character of Temporal Experience, its Scale Non-Invariance, and its Small Scale Structure.
- Halley, E. (1693). An Account of Several Experiments Made to Examine the Nature of the Expansion and Contraction of Fluids by Heat and Cold, in order to ascertain the Divisions of the Thermometer, and to Make that Instrument, in all Places, without Adjusting by a Standard. *Philosophical Transactions of the Royal Society of London*, 17, 650–656.
- Halpern, M. (1987). The Organization and Function of the Vomeronasal System. *Annual Review of Neuroscience*, 10(1), 325–362.
- Hansen, T., Olkkonen, M., Walter, S., & Gegenfurtner, K. R. (2006). Memory modulates color appearance. *Nature Neuroscience*, 9(11), 1367–1368.
- Hardin, C. L. (1988). *Color for Philosophers*. Hackett.
- Hartcher-O'Brien, J., Brighthouse, C., & Levitan, C. A. (2016). A single mechanism account of duration and rate processing via the pacemaker–accumulator and beat frequency models. *Current Opinion in Behavioral Sciences*, 8, 268–275.
- Hinton, S. C., & Meck, W. H. (1997). The “internal clocks” of circadian and interval timing. *Endeavour*, 21(2), 82–87.
- Hoerl, C. (2009). Time and Tense in Perceptual Experience. *Philosophers' Imprint*, 9(12), 1–18.
- Hoerl, C., & McCormack, T. (2011). Time in Cognitive Development. In C. Callender (Ed.), *The Oxford Handbook of Philosophy of Time*. Oxford University Press.

- Hoerl, C., & McCormack, T. (2016). *Making Decisions About the Future : Regret and the Cognitive Function of Episodic Memory*.
- Holloway, F. A., & Wansley, R. A. (1973). Multiple retention deficits at periodic intervals after active and passive avoidance learning. *Behavioral Biology*, 9(1), 1–14.
- Hubbard, E. M., Piazza, M., Pinel, P., & Dehaene, S. (2005). Interactions between number and space in parietal cortex. *Nature Reviews. Neuroscience*, 6(6), 435–448.
- Hubbard, T. L. (1996). Representational momentum, centripetal force, and curvilinear impetus. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 22(4), 1049–1060.
- Hubbard, T. L., & Bharucha, J. J. (1988). Judged displacement in apparent vertical and horizontal motion. *Perception & Psychophysics*, 44(3), 211–221.
- Hume, D. (2000). *A Treatise of Human Nature*. Oxford University Press.
- Husserl, E. (2008). *On the Phenomenology of the Consciousness of Internal Time*. (J. B. Brough, Trans.) (Softcover reprint of the original 1st ed. 1991 edition). Springer.
- Hutto, D. D., & Myin, E. (2013). *Radicalizing Enactivism: Basic Minds Without Content*. MIT Press.
- Ismael, J. (Forthcoming). From Physical Time to Human Time. In Y. Abrams (Ed.), *Cosmological and Psychological Time*. Springer.
- Ismael, J. (2010). Temporal Experience. In C. Callender (Ed.), *Oxford Handbook on Time*. Oxford University Press.
- Ivry, R. B., & Schlerf, J. E. (2008). Dedicated and intrinsic models of time perception. *Trends in Cognitive Sciences*, 12(7), 273–280.
- Jacobs, G. H. (2009). Evolution of colour vision in mammals. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1531), 2957–2967.
- James, W. (1890). *The Principles of Psychology, Vol. 1* (Revised ed. edition). Dover Publications.

- Jones, C. R. G., Rosenkranz, K., Rothwell, J. C., & Jahanshahi, M. (2004). The right dorsolateral prefrontal cortex is essential in time reproduction: an investigation with repetitive transcranial magnetic stimulation. *Experimental Brain Research*, 158(3), 366–372.
- Jones, L. A., & Wearden, J. H. (2003). More is not necessarily better: Examining the nature of the temporal reference memory component in timing. *The Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology*, 56(4), 321–343.
- Jordan, G., Deeb, S. S., Bosten, J. M., & Mollon, J. D. (2010). The dimensionality of color vision in carriers of anomalous trichromacy. *Journal of Vision*, 10(8), 12–12.
- Jordan, G., & Mollon, J. D. (1993). A study of women heterozygous for colour deficiencies. *Vision Research*, 33(11), 1495–1508.
- Kalmijn, A. J. (1982). Electric and magnetic field detection in elasmobranch fishes. *Science*, 218(4575), 916–918.
- Kanabus, M., Szelag, E., Rojek, E., & Pöppel, E. (2002). Temporal order judgement for auditory and visual stimuli. *Acta Neurobiologiae Experimentalis*, 62(4), 263–270.
- Kant, I. (1998). *Critique of Pure Reason* (Translated and Edited by Paul Guyer & Allen W. Wood). Cambridge.
- Karmarkar, U. R., & Buonomano, D. V. (2007). Timing in the Absence of Clocks: Encoding Time in Neural Network States. *Neuron*, 53(3), 427–438.
- Keeley, B. L. (2002). Making Sense of the Senses: Individuating Modalities in Humans and Other Animals. *Journal of Philosophy*, 99(1), 5–28.
- Kelly, S. D. (2005). The Puzzle of Temporal Experience. In A. Brook (Ed.), *Cognition and the Brain: The Philosophy and Neuroscience Movement* (pp. 208–238). Cambridge: Cambridge University Press.
- Killeen, P. R., & Gregor, J. (1988). A behavioral theory of timing. *Psychological Review*, 95(2), 274–295.
- Kirsh, D. (2006). Implicit and Explicit Representation. In *Encyclopedia of Cognitive Science*. John Wiley & Sons, Ltd.

- Ko, C. H., McDonald, R. J., & Ralph, M. R. (2003). The Suprachiasmatic Nucleus is not Required for Temporal Gating of Performance on a Reward-based Learning and Memory Task. *Biological Rhythm Research*, 34(2), 177–192.
- Koch, C. (2004). *The Quest for Consciousness*. Roberts and Company.
- Koch, G., Oliveri, M., Torriero, S., Salerno, S., Lo Gerfo, E., & Caltagirone, C. (2007). Repetitive TMS of cerebellum interferes with millisecond time processing. *Experimental Brain Research*, 179(2), 291–299.
- Kozuch, B. (2015). Dislocation, Not Dissociation: The Neuroanatomical Argument Against Visual Experience Driving Motor Action. *Mind and Language*, 30(5), 572–602.
- Kulvicki, J. (2005). Perceptual Content, Information, and the Primary/Secondary Quality Distinction. *Philosophical Studies*, 122(2), 103–131.
- Lakoff, G., & Johnson, M. (1980). *Metaphors We Live By*. University of Chicago Press.
- Lamme, V. A. (2001). Blindsight: the role of feedforward and feedback corticocortical connections. *Acta Psychologica*, 107(1–3), 209–228.
- Lamme, V. A. F. (2006). Towards a True Neural Stance on Consciousness. *Trends in Cognitive Sciences*, 10(11), 494–501.
- Langacker, R. W. (1986). An Introduction to Cognitive Grammar. *Cognitive Science*, 10(1), 1–40.
- Lapid, E., Ulrich, R., & Rammsayer, T. (2009). Perceptual learning in auditory temporal discrimination: no evidence for a cross-modal transfer to the visual modality. *Psychonomic Bulletin & Review*, 16(2), 382–389.
- Laurence, S., & Margolis, E. (2005). Number and Natural Language. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The Innate Mind: Structure and Content* (pp. 216–235). Oxford University Press.
- Lebedev, M. A., O'Doherty, J. E., & Nicolelis, M. A. L. (2008). Decoding of temporal intervals from cortical ensemble activity. *Journal of Neurophysiology*, 99(1), 166–186.

- Lee, G. (2014a). Extensionalism, Atomism, and Continuity. In N. Oaklander (Ed.), *Debates in the metaphysics of time*. Bloomsbury.
- Lee, G. (2014b). Temporal Experience and the Temporal Structure of Experience. *Philosophers' Imprint*, 14(3).
- Leslie, A. M., Xu, F., Tremoulet, P. D., & Scholl, B. J. (1998). Indexing and the Object Concept: Developing 'what' and 'where' systems. *Trends in Cognitive Sciences*, 2(1), 10–18.
- Levy, J. M., Namboodiri, V. M. K., & Hussain Shuler, M. G. (2015). Memory bias in the temporal bisection point. *Frontiers in Integrative Neuroscience*, 44.
- Lewis, D. (1976). The Paradoxes of Time Travel. *American Philosophical Quarterly*, 13(2), 145–152.
- Lewis, D. K. (1973). *Counterfactuals*. Blackwell Publishers.
- Lewis, D. K. (1986). *On the Plurality of Worlds*. Blackwell Publishers.
- Lewis, P. A., Miall, R. C., Daan, S., & Kacelnik, A. (2003). Interval timing in mice does not rely upon the circadian pacemaker. *Neuroscience Letters*, 348(3), 131–134.
- Llinás, R., Ribary, U., Contreras, D., & Pedroarena, C. (1998). The neuronal basis for consciousness. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 353(1377), 1841–1849.
- Lloyd, D. (2012). Neural Correlates of Temporality: Default Mode Variability and Temporal Awareness. *Consciousness and Cognition*, 21(2), 695–703.
- Locke, J. (1689). *An Essay Concerning Human Understanding*. Hackett Publishing.
- Lopes, D. M. M. (2000). What is It Like to See with Your Ears? The Representational Theory of Mind. *Philosophy and Phenomenological Research*, 60(2), 439–453.
- Lycan, W. G. (1996). *Consciousness and Experience*. MIT Press.
- Machamer, P. K., Darden, L., & Craver, C. F. (2000). Thinking About Mechanisms. *Philosophy of Science*, 67(1), 1–25.
- Macpherson, F. (2011). Taxonomising the Senses. *Philosophical Studies*, 153(1), 123–142.

- Malapani, C., & Fairhurst, S. (2002). Scalar Timing in Animals and Humans. *Learning and Motivation*, 33(1), 156–176.
- Maniadakis, M., & Trahanias, P. (2016). When and How-Long: A Unified Approach for Time Perception. *Frontiers in Psychology*, 7.
- Marchant, E. G., & Mistlberger, R. E. (1997). Anticipation and entrainment to feeding time in intact and SCN-ablated C57BL/6j mice. *Brain Research*, 765(2), 273–282.
- Marcus, G. F. (2003). *The Algebraic Mind: Integrating Connectionism and Cognitive Science*. MIT Press.
- Marr, D. (1982). *Vision*. Freeman.
- Marshall, J., & Oberwinkler, J. (1999). Ultraviolet vision: The colourful world of the mantis shrimp. *Nature*, 401(6756), 873–874.
- Martin, B., Giersch, A., Huron, C., & van Wassenhove, V. (2013). Temporal event structure and timing in schizophrenia: preserved binding in a longer “now.” *Neuropsychologia*, 51(2), 358–371.
- Martin, M. G. F. (2002). The Transparency of Experience. *Mind and Language*, 4(4), 376–425.
- Matell, M. S., & Meck, W. H. (2000). Neuropsychological mechanisms of interval timing behavior. *BioEssays: News and Reviews in Molecular, Cellular and Developmental Biology*, 22(1), 94–103.
- Matell, M. S., & Meck, W. H. (2004). Cortico-striatal circuits and interval timing: coincidence detection of oscillatory processes. *Brain Research. Cognitive Brain Research*, 21(2), 139–170.
- Matthen, M. (2014). Active Perception and the Representation of Space. In D. Stokes, M. Matthen, & S. Biggs (Eds.), *Perception and Its Modalities* (pp. 44–72). Oxford University Press.
- Matthen, M. (2015). The Individuation of the Senses. In M. Matthen (Ed.), *Oxford Handbook of the Philosophy of Perception* (pp. 567–586). Oxford University Press.
- Matthews, W. J., Stewart, N., & Wearden, J. H. (2011). Stimulus intensity and the perception of duration. *Journal of Experimental Psychology. Human Perception and Performance*, 37(1), 303–313.
- Maudlin, T. (2002). Remarks on the Passing of Time. *Proceedings of the Aristotelian Society*, 102(3), 237–252.

- Mauk, M. D., & Buonomano, D. V. (2004). The neural basis of temporal processing. *Annual Review of Neuroscience*, 27, 307–340.
- McTaggart, J. E. (1908). The Unreality of Time. *Mind*, 17(68), 457–474.
- Meck, W. H., & Church, R. M. (1983). A mode control model of counting and timing processes. *Journal of Experimental Psychology. Animal Behavior Processes*, 9(3), 320–334.
- Mendelovici, A., & Bourget, D. (2014). Naturalizing Intentionality: Tracking Theories Versus Phenomenal Intentionality Theories. *Philosophy Compass*, 9(5), 325–337.
- Merritt, D. J., Casasanto, D., & Brannon, E. M. (2010). Do monkeys think in metaphors? Representations of space and time in monkeys and humans. *Cognition*, 117(2), 191–202.
- Meyer, U. (2013). *The Nature of Time*. Clarendon Press.
- Millikan, R. G. (1989). Biosemantics. *Journal of Philosophy*, 86(July), 281–97.
- Millikan, R. G. (1995). Pushmi-Pullyu Representations. *Philosophical Perspectives*, 9, 185–200.
- Montemayor, C. (2013). *Minding Time: A Philosophical and Theoretical Approach to the Psychology of Time*. Brill.
- Moore, R. Y. (1997). Circadian rhythms: basic neurobiology and clinical applications. *Annual Review of Medicine*, 48, 253–266.
- Navarra, J., García-Morera, J., & Spence, C. (2012). Temporal adaptation to audiovisual asynchrony generalizes across different sound frequencies. *Perception Science*, 3, 152.
- Newton-Smith, W. (1980). *The Structure of Time*. Routledge & Kegan Paul.
- Nudds, M. (2004). The Significance of the Senses. *Proceedings of the Aristotelian Society*, 104(1), 31–51.
- Nudds, M. (2011). The Senses as Psychological Kinds. In F. Macpherson (Ed.), *The Senses: Classic and Contemporary Philosophical Perspectives*. Oxford.
- O’Dea, J. (2011). A Proprioceptive Account of the Senses. In F. Macpherson (Ed.), *The Senses: Classical and Contemporary Philosophical Perspectives*. Oxford University Press.

- Panzeri, S., Ince, R. A. A., Diamond, M. E., & Kayser, C. (2014). Reading spike timing without a clock: intrinsic decoding of spike trains. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 369(1637), 20120467.
- Peacocke, C. (1986). Analogue Content. *Proceedings of the Aristotelian Society*, 60, 1–17.
- Peacocke, C. (2015). Magnitudes: Metaphysics, Explanation, and Perception. In A. Coliva, V. Munz, & D. Moyal-Sharrock (Eds.), *Mind, Language and Action: Proceedings of the 36th International Wittgenstein Symposium* (pp. 357–388). De Gruyter.
- Pearson, S. N., Halford, S., & Foster, R. G. (2009). The evolution of irradiance detection: melanopsin and the non-visual opsins. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1531), 2849–2865.
- Pelczar, M. (2015). *Sensorama: A Phenomenalist Analysis of Spacetime and Its Contents*. Oxford University Press.
- Perry, J. (1979). The Problem of the Essential Indexical. *Noûs*, 13(December), 3–21.
- Phillips, I. (2010). Perceiving Temporal Properties. *European Journal of Philosophy*, 18(2), 176–202.
- Phillips, I. (2012). Attention to the Passage of Time. *Philosophical Perspectives*, 26(1), 277–308.
- Phillips, I. (2014a). Breaking the Silence: Motion Silencing and Experience of Change. *Philosophical Studies*, 168(3), 693–707.
- Phillips, I. (2014b). Experience of and in Time. *Philosophy Compass*, 9(2), 131–144.
- Poppel, E. (1988). *Mindworks: Time and Conscious Experience*. Harcourt Brace Jovanovich.
- Prinz, J. J. (2000). A Neurofunctional Theory of Visual Consciousness. *Consciousness and Cognition*, 9(2), 243–59.
- Prinz, J. J. (2002). *Furnishing the Mind: Concepts and Their Perceptual Basis*. MIT Press.
- Prinz, J. J. (2012). *The Conscious Brain: How Attention Engenders Experience* (1 edition). Oxford ; New York: Oxford University Press.

- Proske, U., & Gandevia, S. C. (2012). The Proprioceptive Senses: Their Roles in Signaling Body Shape, Body Position and Movement, and Muscle Force. *Physiological Reviews*, 92(4), 1651–1697.
- Prosser, S. (2007). Could We Experience the Passage of Time? *Ratio*, 20(1), 75–90.
- Purves, D., Augustine, G. J., Fitzpatrick, D., Katz, L. C., LaMantia, A.-S., McNamara, J. O., & Williams, S. M. (2001). Summation of Synaptic Potentials. Retrieved from <http://www.ncbi.nlm.nih.gov/books/NBK11104/>, September 2016.
- Pylyshyn, Z. (1989). The Role of Location Indexes in Spatial Perception: A Sketch of the FINST Spatial-Index Model. *Cognition*, 32(1), 65–97.
- Pylyshyn, Z. W. (2001). Visual Indexes, Preconceptual Objects, and Situated Vision. *Cognition*, 80(1–2), 127–158.
- Pylyshyn, Z. W. (2007). *Things and Places: How the Mind Connects with the World*. The MIT Press.
- Quine, W. V. (1960). *Word and Object*. The MIT Press.
- Raby, C. R., Alexis, D. M., Dickinson, A., & Clayton, N. S. (2007). Planning for the future by western scrub-jays. *Nature*, 445(7130), 919–921.
- Ralph, M. R., Foster, R. G., Davis, F. C., & Menaker, M. (1990). Transplanted suprachiasmatic nucleus determines circadian period. *Science (New York, N.Y.)*, 247(4945), 975–978.
- Rammsayer, T. H. (1999). Neuropharmacological evidence for different timing mechanisms in humans. *The Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology*, 52(3), 273–286.
- Rammsayer, T. H., Bortner, N., & Troche, S. J. (2015). Visual-auditory differences in duration discrimination of intervals in the subsecond and second range. *Perception Science*, 1626.
- Ramsey, W. M. (2007). *Representation Reconsidered*. Cambridge University Press.
- Rashbrook, O. (2013). An Appearance of Succession Requires a Succession of Appearances. *Philosophy and Phenomenological Research*, 87(3), 584–610.

- Recanati, F. (2012). Immunity to Error Through Misidentification: What It is and Where It Comes From. In S. Prosser & F. Recanati (Eds.), *Immunity to Error Through Misidentification: New Essays* (pp. 180–201). Cambridge University Press.
- Reichenbach, H. (1956). *The Direction of Time*. Dover Publications.
- Reichenbach, H. (1957). *The Philosophy of Space and Time* (1st English Ed. edition). New York: Dover Publications.
- Reid, T. (1855). *Essays on the Intellectual Powers of Man*. Early Modern Texts Online.
- Reutimann, J., Yakovlev, V., Fusi, S., & Senn, W. (2004). Climbing Neuronal Activity as an Event-Based Cortical Representation of Time. *The Journal of Neuroscience*, 24(13), 3295–3303.
- Rey, G. (1997). *Contemporary Philosophy of Mind: A Contentiously Classical Approach*. Blackwell.
- Rosenthal, D. M. (2002). Explaining Consciousness. In D. J. Chalmers (Ed.), *Philosophy of Mind: Classical and Contemporary Readings* (pp. 109–131). Oxford University Press.
- Rosenwasser, A. M., Pelchat, R. J., & Adler, N. T. (1984). Memory for feeding time: possible dependence on coupled circadian oscillators. *Physiology & Behavior*, 32(1), 25–30.
- Ruby, N. F., Hwang, C. E., Wessells, C., Fernandez, F., Zhang, P., Sapolsky, R., & Heller, H. C. (2008). Hippocampal-dependent learning requires a functional circadian system. *Proceedings of the National Academy of Sciences*, 105(40), 15593–15598.
- Russell, B. (1905). On Denoting. *Mind*, 14(56), 479–493.
- Ryder, D. (2004). SINBAD Neurosemantics: A Theory of Mental Representation. *Mind & Language*, 19(2), 211–240.
- Schibler, U., Ripperger, J., & Brown, S. A. (2003). Peripheral Circadian Oscillators in Mammals: Time and Food. *Journal of Biological Rhythms*, 18(3), 250–260.
- Shams, L., Kamitani, Y., & Shimojo, S. (2000). Illusions: What you see is what you hear. *Nature*, 408(6814), 788–788.

- Shi, Z., Church, R. M., & Meck, W. H. (2013). Bayesian optimization of time perception. *Trends in Cognitive Sciences*, 17(11), 556–564.
- Shoemaker, S. (1969). Time Without Change. *Journal of Philosophy*, 66(12), 363–381.
- Sincich, L. C., Horton, J. C., & Sharpee, T. O. (2009). Preserving Information in Neural Transmission. *The Journal of Neuroscience*, 29(19), 6207–6216.
- Skyrms, B. (2010). *Signals*. Oxford University Press.
- Smith, B. (1988). On the Semantics of Clocks. In J. H. Fetzer (Ed.), *Aspects of AI*. Kluwer.
- Sober, E. (2006). Parsimony. In J. Pfeifer & S. Sarkar (Eds.), *The Philosophy of Science: An Encyclopedia*.
- Solomon, S. G., & Lennie, P. (2007). The machinery of colour vision. *Nature Reviews. Neuroscience*, 8(4), 276–286.
- Stampe, D. W. (1977). Towards a Causal Theory of Linguistic Representation. *Midwest Studies in Philosophy*, 2(1), 42–63.
- Stephan, F. K. (2002). The “Other” Circadian System: Food as a Zeitgeber. *Journal of Biological Rhythms*, 17(4), 284–292.
- Stetson, C., Cui, X., Montague, P. R., & Eagleman, D. M. (2006). Motor-Sensory Recalibration Leads to an Illusory Reversal of Action and Sensation. *Neuron*, 51(5), 651–659.
- Stevenson, R. A., Siemann, J. K., Schneider, B. C., Eberly, H. E., Woynaroski, T. G., Camarata, S. M., & Wallace, M. T. (2014). Multisensory Temporal Integration in Autism Spectrum Disorders. *The Journal of Neuroscience*, 34(3), 691–697.
- Stich, S. (1982). From Folk Psychology to Cognitive Science: The Case Against Belief. In a Woodfield (Ed.), *Philosophical Review* (pp. 418–421). MIT Press.
- Suppes, P., & Zinnes, J. (1963). Basic Measurement Theory. In D. Luce (Ed.), *Journal of Symbolic Logic* (pp. 1–1). John Wiley & Sons.

- Takahashi, K., & Watanabe, K. (2015). Mental Summation of Temporal Duration within and across Senses. *PLOS ONE*, 10(10), e0141466.
- Tal, E. (2013). Old and New Problems in Philosophy of Measurement. *Philosophy Compass*, 8(12), 1159–1173.
- Tal, E. (2016). Making Time: A Study in the Epistemology of Measurement. *British Journal for the Philosophy of Science*, 67(1), 297–335.
- Taylor, R. (1994). Brave New Nose: Sniffing out Human Sexual Chemistry. *Journal of NIH Research*, 6.
- The International System of Units (Brochure)*. (2006) (8th ed.). Bureau International des Poids et Mesures.
- Thom, J. M., & Clayton, N. S. (2013). Re-caching by Western Scrub-Jays (*Aphelocoma californica*) Cannot Be Attributed to Stress. *PLOS ONE*, 8(1), e52936.
- Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97–136.
- Treisman, M. (1963). Temporal discrimination and the indifference interval. Implications for a model of the “internal clock.” *Psychological Monographs*, 77(13), 1–31.
- Tse, P. U., Intriligator, J., Rivest, J., & Cavanagh, P. (2004). Attention and the subjective expansion of time. *Perception & Psychophysics*, 66(7), 1171–1189.
- Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. MIT Press.
- Tye, M. (2000). *Consciousness, Color, and Content*. MIT Press.
- Ungerleider, L., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *Analysis of Visual Behavior*. MIT Press.
- Usher, M. (2001). A Statistical Referential Theory of Content: Using Information Theory to Account for Misrepresentation. *Mind & Language*, 16(3), 311–334.

- Vasalou, C., & Henson, M. A. (2010). A multiscale model to investigate circadian rhythmicity of pacemaker neurons in the suprachiasmatic nucleus. *PLoS Computational Biology*, 6(3).
- Velarde, R. A., Sauer, C. D., O. Walden, K. K., Fahrbach, S. E., & Robertson, H. M. (2005). Pteropsin: A vertebrate-like non-visual opsin expressed in the honey bee brain. *Insect Biochemistry and Molecular Biology*, 35(12), 1367–1377.
- Verfaillie, K., & d’Ydewalle, G. (1991). Representational momentum and event course anticipation in the perception of implied periodical motions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17(2), 302–313.
- Vroomen, J., & Keetels, M. (2010). Perception of intersensory synchrony: A tutorial review. *Attention, Perception, & Psychophysics*, 72(4), 871–884.
- Vroomen, J., Keetels, M., de Gelder, B., & Bertelson, P. (2004). Recalibration of temporal order perception by exposure to audio-visual asynchrony. *Brain Research. Cognitive Brain Research*, 22(1), 32–35.
- Walsh, V. (2003). A Theory of Magnitude: Common Cortical Metrics of Time, Space and Quantity. *Trends in Cognitive Sciences*, 7(11), 483–488.
- Watanabe, T., Harner, A. M., Miyauchi, S., Sasaki, Y., Nielsen, M., Palomo, D., & Mukai, I. (1998). Task-dependent influences of attention on the activation of human primary visual cortex. *Proceedings of the National Academy of Sciences*, 95(19), 11489–11492.
- Wearden, J. H. (1999). “Beyond the fields we know...”: exploring and developing scalar timing theory. *Behavioural Processes*, 45(1–3), 3–21.
- Wearden, J. H. (2001). Internal Clocks and the Representation of Time. In C. Hoerl & T. McCormack (Eds.), *Time and Memory* (pp. 37–58). Oxford University Press.
- Wearden, J. H., & Doherty, M. (1995). Exploring and developing a connectionist model of animal timing: Peak procedure and fixed-interval simulations. *Journal of Experimental Psychology: Animal Behavior Processes*, 21(2), 99–115.

- Weiskopf, D. A. (2007). Concept Empiricism and the Vehicles of Thought. *Journal of Consciousness Studies*, 14(s 9-10), 156–183.
- WHORF, B. L. (1944). THE RELATION OF HABITUAL THOUGHT AND BEHAVIOR TO LANGUAGE. *ETC: A Review of General Semantics*, 1(4), 197–215.
- Wightman, F. L., & Kistler, D. J. (1992). The dominant role of low-frequency interaural time differences in sound localization. *The Journal of the Acoustical Society of America*, 91(3), 1648–1661.
- Wittmann, M., Carter, O., Hasler, F., Cahn, B. R., Grimberg, U., Spring, P., ... Vollenweider, F. X. (2007). Effects of psilocybin on time perception and temporal control of behaviour in humans. *Journal of Psychopharmacology (Oxford, England)*, 21(1), 50–64.
- Wittmann, M., & Wassenhove, V. van. (2009). The experience of time: neural mechanisms and the interplay of emotion, cognition and embodiment. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1525), 1809–1813.
- Wolfe, J. M., & Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, 5(6), 495–501.
- Woodward, J. (2005). *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press, USA.
- Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., & Reynolds, J. R. (2007). Event Perception: A Mind/Brain Perspective. *Psychological Bulletin*, 133(2), 273–293.
- Zacks, J. M., & Tversky, B. (2001). Event structure in perception and conception. *Psychological Bulletin*, 127(1), 3–21.
- Zakay, D., & Block, R. A. (1997). Temporal Cognition. *Current Directions in Psychological Science*, 6(1), 12–16.
- Zucker, R. S. (1989). Short-term synaptic plasticity. *Annual Review of Neuroscience*, 12, 13–31.