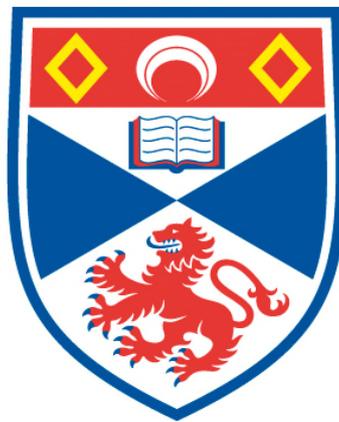


A FRAGMENTED WORLD

Martin A. Lipman

**A Thesis Submitted for the Degree of PhD
at the
University of St Andrews**



2015

**Full metadata for this item is available in
St Andrews Research Repository
at:**

<http://research-repository.st-andrews.ac.uk/>

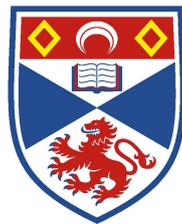
Please use this identifier to cite or link to this item:

<http://hdl.handle.net/10023/9498>

This item is protected by original copyright

A Fragmented World

Martin A. Lipman



University of
St Andrews

This thesis is submitted for the degree of PhD

at the University of St Andrews

2 March 2015

Declarations

1. Candidate's declarations:

I, Martin A. Lipman hereby certify that this thesis, which is approximately 78.000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in September 2010 and as a candidate for the degree of PhD in September 2011; the higher study for which this is a record was carried out in the University of St Andrews between 2011 and 2015.

Date 25/2/2015

Signature of candidate

2. Supervisor's declaration:

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date 25/2/2015

Signature of supervisor

3. Permission for publication: *(to be signed by both candidate and supervisor)*

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

PRINTED COPY

a) No embargo on print copy

ELECTRONIC COPY

a) No embargo on electronic copy

Date 25/2/2015

Signature of candidate Signature of supervisor

Please note initial embargos can be requested for a maximum of five years. An embargo on a thesis submitted to the Faculty of Science or Medicine is rarely granted for more than two years in the first instance, without good justification. The Library will not lift an embargo before confirming with the student and supervisor that they do not intend to request a continuation. In the absence of an agreed response from both student and supervisor, the Head of School will be consulted. Please note that the total period of an embargo, including any continuation, is not expected to exceed ten years. Where part of a thesis is to be embargoed, please specify the part and the reason.

Abstract

Objects often manifest themselves in incompatible ways across perspectives that are on a par. Phenomena of this kind have been responsible for crucial revisions to our conception of the world, both philosophical and scientific. The standard response to them is to deny that the way things appear from different perspectives are ways things really are out there, a response that is based on an implicit metaphysical assumption that the world is a unified whole. This dissertation explores the possibility that this assumption is false, that the world is fragmented instead of unified. On the proposed understanding of such worldly fragmentation, there is a notion of co-obtainment according to which two facts may obtain without co-obtaining. Since not every fact that obtains also co-obtains with every other fact, two incompatible facts may both obtain, as long as they do not co-obtain in the introduced sense. The possibility of such fragmentation sheds new light on a range of phenomena. It allows us to explore a view of time that takes the notion of passage as its defining primitive. It bolsters a no-subject view of experience against the objection that it leads to solipsism. It allows a realist view about colours to withstand the objection from conflicting appearances. And, it makes room for a view on which things really have the properties that are attributed to objects and events across different frames of reference, such as length, mass, duration and simultaneity. Overall, fragmentalism changes the way in which the manifest image feeds into an objective conception of the world: what is manifest to us is not misleading in what sort of properties it shows the world to have, it's only misleading in making it seem more unified than it really is.

Acknowledgements

First and foremost I would like to express my deepest gratitude to my supervisor, Katherine Hawley. She has been absolutely amazing. She has read and reread so much material, always with an open mind, no matter how unfamiliar the ideas explored in it, and always offering helpful comments. She taught me basic philosophical skills, about where explanations may end, about what standards we can or cannot impose on views, and about how to express matters more clearly. She was always there for personal support and advice. If I may ever supervise students myself, I can only hope that they will think as highly of me as I think of Katherine.

I'm also hugely grateful to my second supervisor, Aaron Cotnoir. He too has been absolutely amazing, pouring time and energy into me and my project. He personally taught me the technical skills I needed to develop the logical aspect of fragmentalism. His always insightful comments and caring support have been tremendously important to me during the past years. I really could not have wished for better supervision.

I owe special thanks to two friends. I'm thankful to Sander Werkhoven, for his listening ear, for being such a steady voice of common sense, and for his detailed comments on my work, which always pushed me for plausibility and clarity. And I want to thank Bruno Jacinto, for our countless hours of discussion and setting me straight on so many philosophical issues. Our discussions had a great influence on me, and my project.

I owe thanks to many others. Adrian Haddock's seminar on Moore's book *Points of View* (1997), lay the seeds for the interests I have pursued here, and I'm thankful that Adrian encouraged me to pursue them further. A long discussion with Colin Johnston, after a talk in Stirling, caused a

significant turn in my views on tense. Kit Fine's 'Tense and reality' (2005) was a major source of inspiration for this whole thesis, and his insightful comments on two talks steered me away from unfruitful avenues. Two long discussions with Gabriel Uzquiano helped me become much clearer about the motivations in play. Sarah Broadie, Derek Ball, Patrick Greenough, JC Beall, Graham Priest, Peter Sullivan, Laurie Paul, Andy Egan - they all took out time from their busy schedules to read drafts, or discuss fragmentalism. Finally, I benefited immensely from interactions with the many people based in or passing through Arché.

Elements from *Chapter 1* and *Chapter 2* are forthcoming in *Philosophical Studies* as 'On Fine's Fragmentalism'. I'm grateful for the helpful suggestions made by an anonymous referee of this journal. Parts of *Chapter 1* and *Chapter 7* are forthcoming in the *Australasian Journal of Philosophy* as 'Perspectival Variance and Worldly Fragmentation'. Two anonymous referees for this journal offered detailed and very helpful comments.

I would like to thank the St Andrews/Stirling Graduate Programme (SASP) for funding me the past three years, and offering such a wonderful intellectual environment.

On a more personal note, I want to thank my parents, Sigrid and Leo, for their unwavering support. And I want to thank my partner Robyn, for the sacrifices she made so that I could pursue this research, and for her understanding and encouragement throughout these years.

To Sophie and Nora.

Contents

Introduction	1
1 Introducing Fragmentalism	7
1.1 The possibility of worldly fragmentation	7
1.2 A basic semantics for co-obtainment	18
1.3 Summary and concluding remarks	26
1.4 Appendix: incorporating modality	27
1.5 Appendix: incorporating quantification	33
2 Fine on Fragmentalism	40
2.1 Fine’s conception of fragmentalism	40
2.2 Critical discussion I: on tense and truth	46
2.3 Critical discussion II: on (in)coherence	54
2.4 A fragmentalist A-theory	60
2.5 Summary and concluding remarks	65
3 Passage in a Fragmented World	66
3.1 Passage and A-theories	67
3.2 A Passage theory of time	72
3.3 The problem of change	81
3.4 Summary and concluding remarks	88
3.5 Appendix: truth-conditions	91
3.6 Appendix: translations into the A-theory	97

4	Phenomenality in a Fragmented World	100
4.1	Three considerations in favour	101
4.2	Phenomenal facts as appearance facts	108
4.3	The problem of conflicting perspectives	115
4.4	Fragmentation in the appearance facts	118
4.5	Co-consciousness and split brains	123
4.6	Summary and concluding remarks	133
5	Colours in a Fragmented World	135
5.1	Conflicting appearances	135
5.2	A fragmentalist account of colour	145
5.3	Summary and concluding remarks	150
6	Special Relativity in a Fragmented World	151
6.1	From Newtonian to Galilean spacetime	152
6.2	From Galilean to Minkowski spacetime	160
6.3	Unity, invariance, and antirealisms	166
6.4	A fragmentalist account of relativity	174
6.5	Relativity, tense and passage	179
6.6	Summary and concluding remarks	181
7	General Discussion: Perspectival Variance	183
7.1	Perspectival variance	184
7.2	Other instances of perspectival variance	194
7.3	Objectivity in a fragmented world	200
7.4	Summary and concluding remarks	207
7.5	Appendix: metaphysical relativism	207
	Conclusion	210

Introduction

This dissertation develops a general metaphysical view of the world. According to this view, the world is rich in properties and fragmented in structure. The dissertation makes a plea for the intelligibility and importance of the theoretical concept involved in this conception of the world, a notion of the co-obtainment of facts that allows for two facts to obtain without co-obtaining. If we acknowledge this possibility of worldly fragmentation, we see that the standard conception of a large range of phenomena depends on an implicit assumption about the metaphysical structure of the world, namely that the world is a thoroughly unified place, that all facts that obtain co-obtain with each other.

Much turns on the introduction of the new concept of co-obtainment. I know of no way to define this notion and yet everything is conditional on our grasp of it. This places a great burden on other means of elucidation. It will be a methodological assumption of this dissertation that we are able to make conceptual leaps and adopt concepts that are not defined in terms of concepts we already possess and that, by adopting new concepts as a result of informal elucidation, we can see things in a truly new light and improve our understanding of things. Rosen describes this process of informal elucidation as follows:

If we do not begin with a definition, we must offer some sort of informal elucidation. We all know roughly how this works in other parts of philosophy. The neophyte is presented with a battery of paradigms and foils, ordinary language paraphrases (with commentary), and bits and pieces of the inferential role of

the target notion, and then somehow as a result of this barrage he cottons on. (Rosen 2006: 14).

I will employ all these forms of elucidation.¹ I will offer ordinary language paraphrases and metaphors, and will construct models that represent the difference between facts that co-obtain and facts that fail to do so. Set theoretic models are the basis for semantic machinery from which we can read the inferential features of the proposed concept and, more generally, the logical structure of the fragmentalist's conception of the world. The elucidation then continues on through the application of the concept in concrete cases.

The possibility of obtaining facts that do not need to co-obtain renders the assumption that all facts actually co-obtain substantive. I will call this the unity assumption. Not only can we make sense of a world in which it's mistaken, there are actually many cases in which the assumption *seems* mistaken. The relevant cases can be loosely described as perspectival phenomena: cases where there are certain conditions under which one fact seems to obtain and a different set of conditions under which another incompatible fact seems to obtain, without there being any set of conditions under which these facts all seem to co-obtain. I will discuss four such cases:

Time. It may be that from the perspective of one time, the candle seems straight and from the perspective of a different time, the same candle seems bent.

Consciousness. It may be that, from the perspective of one subject, A, there seems something it's like when A undergoes experience, while from the perspective of another subject, B, there seems something it is like when B undergoes experience.

¹Missing from Rosen's list is the use of metaphors. Lakoff and Johnson (1980) make a convincing case that many of our ordinary concepts are just well-worn metaphors. They show in particular that many philosophical concepts have metaphorical origins (see Lakoff and Johnson 1999). None of this is to say that such concepts cannot latch onto anything in the world; it's just to say that the use of metaphors may be essentially involved in the creation of new concepts.

Colours. From the perspective of one observer, a surface may seem to be green and from the perspective of a different observer, the same surface may seem to be blue.

Relativity. From the perspective of an observer at rest, an object can seem to be moving with a constant velocity while from the perspective of an observer that moves, the same object can seem to be at rest.

These kinds of cases offer pre-theoretical support for views according to which things are as they appear to be from each of the different perspectives, and fragmentalism allows us to make sense of this. The qualifying phrase ‘from the perspective of ...’ is dropped: these facts are not taken to obtain ‘relative to different perspectives’, rather, they are taken to obtain without co-obtaining. The facts that obtain across time, do not all co-obtain; the phenomenal facts that appear to obtain across subjects do not all co-obtain, and so on.

The cases of perspectival variance support the relevant fragmentalist views if we think that the way the world is given to us in experience can be evidence for the way that world is. It will thus be an epistemological assumption of this dissertation that observations - understood in a suitably loose sense - provide evidence of the nature of properties that are instantiated in the world. Many of the arguments that I offer in defence of various fragmentalist views will presume that experience is always at least a *prima facie* guide to reality.

The overall aim will be to show that the possibility of worldly fragmentation bolsters certain non-standard views of various issues, that long-standing metaphysical problems can be given new solutions, and that we might gain a firmer theoretical hold on some of the most nebulous phenomena, such as the passing of time and the conscious aspect of experience. Besides the proposals of various fragmentation-based views, the overarching claim is simply that worldly fragmentation allows us to make sense of things that did not make much sense before.

The structure of the dissertation was dictated by what I took to be the best line of defence for fragmentalism: namely to show how it can help in

a wide range of areas. The possibility that the world is fragmented and that incompatible facts obtain, requires quite a change in mindset. As long as fragmentalism is applied to one area only - say, only to time, or only to conscious experience - it will always seem somewhat of an overkill, too radical for the motivations in question, a metaphysician's play-thing. It seemed to me that only an accumulation of cases could make for real plausibility and show that accepting worldly fragmentation is not some ad-hoc fix for some philosophical issue.

This relates to an important disclaimer. Whole dissertations can and have been written about the topics that are discussed in the chapters of this dissertation, about time, consciousness, colours, and so on. Needless to say, I cannot reach the detail that such specialized treatments reach. The main objective in each individual chapter is not a full treatment of the topic in question but to propose a fragmentalist understanding of it. This means that I will be relatively quick when I discuss worries or apparent objections to the more standard theories of these phenomena in the understanding that the worries are not meant to refute these theories but only to help motivate and shape a new contender.

Below I offer a brief overview of the individual chapters. The chapters are not stand-alone essays and typically build on what came before them. The sections that are labelled as appendices can be skipped without loss of continuity but, otherwise, the chapters are meant to be read in the order in which they occur. Whenever, in citation, two dates occur separated by a slash, as in Berkeley (1713/1979), the first is the original year of publication and the second is the year of the publication whose page numbers are cited and whose details are found in the bibliography. The contents of the seven chapters are as follows.

Chapter 1 serves as a general introduction to fragmentalism - laying the conceptual foundations of the view that is explored in the rest of the dissertation. The concept of co-obtainment is introduced and elucidated. Simple model theory is used to fix the inferential role of co-obtainment, and to show

that fragmentalism is not incoherent. The main aim is to put the fragmentalist framework on a solid conceptual footing.

Chapter 2 discusses the characterization of fragmentalism that is offered by Kit Fine in his ‘Tense and Reality’ (2005) and which serves as a major source of inspiration of the view defended in the thesis. A critical discussion raises open questions for the conceptual machinery of Fine’s fragmentalism, and raises issues for one of Fine’s arguments in favour of his fragmentalist A-theory: the argument from truth. I then offer a formulation of a fragmentalist A-theory within the framework of *Chapter 1*, and show how it solves some but not all of the issues raised in the critical discussion.

Chapter 3 considers what is required for a proper understanding of a real and objective passing of time. I argue that both the standard A-theory and the fragmentalist A-theory of *Chapter 2* fail to capture the passage of time. In light of this, I explore a theory of time that accepts the passage of time as the most basic temporal notion, instead of the usual A-theoretic and B-theoretic notions. This passage theory of time, as I will call it, enables us to offer a new solution to the problem of change.

Chapter 4 discusses a no-subject view of experience. First I discuss three phenomenological considerations in favour of a no-subject view and propose a way of precisifying it. The proposal is that a phenomenal fact consists in something *appearing* to be the case. Then I discuss a problem for the no-subject view that shows how - without resort to fragmentalism - the view leads to solipsism. In the final section, I discuss how the resulting view offers an independently plausible understanding of the unity of consciousness as well as of the disunity of consciousness that seems to arise in cases of subjects whose hemispheres have been disconnected.

Chapter 5 considers what is required for things to be genuinely coloured in the world. A well known obstacle to this is the widespread occurrence of cases where the same thing exhibits incompatible colours to different observers, even though none of these observers seems privileged over the other. I will discuss this well-known problem, and offer a fragmentalist solution to it.

Chapter 6 argues that, just as the paradox of conflicting apparent colours threatens the view that the world is genuinely coloured, the special theory

of relativity predicts similar sort of cases that threaten the view that objects have lengths, shapes and masses, and that events have durations and occur in a certain kind of temporal order. I offer a fragmentalist interpretation of the relevant facts, which doesn't have the antirealist conclusions concerning the frame-dependent properties. I also revisit the discussion of the passage of time, and show how the offered account fits with the passage theory of time offered in *Chapter 3*.

Chapter 7, finally, offers a general discussion of perspectival phenomena and the fragmentalist interpretations of them. The various cases discussed clearly share a common structure, which I will draw out. I will discuss the question of how widely fragmentalism applies, and whether it undermines the importance of objectivity in our inquiries concerning the world. I also offer some very general considerations in favour of fragmentalism: relating to the possibility of a uniform approach to perspectival phenomena, and the inductive evidence in favour of it.

Chapter 1

Introducing Fragmentalism

There is a metaphysical structure to the world, and we implicitly assume that structure to be one way instead of another. We assume that the world is unified in the sense that any fact co-obtains with every other fact. Indeed, this assumption may well be part of the reason why we think that something cannot bear incompatible properties, be red all over and blue all over for example.

We do not realize that we assume that the world is unified in the relevant sense because we do not ordinarily conceptualize the relevant metaphysical structure. And yet, through the use of models and analogies, we can see that there must be such a metaphysical structure and that it could be different from the way we presume it to be. In this section we use these models and analogies to engineer a concept that latches onto this metaphysical structure, so that we can render our implicit assumption explicit and formulate an alternative view of the world, on which the world is not a unified place.

1.1 The possibility of worldly fragmentation

Imagine that we face each other, and that Sophie and Nora stand between us. We observe things differently: you observe Sophie as being on the left of Nora, I observe Sophie as being on the right of Nora. Sophie and Nora are observed to have incompatible two-place relations from our different perspectives. We

think that Sophie cannot be both on the left and on the right of Nora, and, given that no orientation is privileged as the one way of seeing the world as it truly is, nor is Sophie just on the left of Nora, or just on the right of her. Since there is nothing special about the case, we draw the general conclusion that there are no two-place relations of *being on the left of* or *being on the right of* instantiated by pairs of things out there in the world.

The standard response to such cases of perspectival variance is this: the ways things appear to observers in a case of perspectival variance are *merely* ways things appear to observers in the relevant circumstances. They are not ways things genuinely are, that is, they do not feature in a conception of the world as it is in itself (B. Williams 1979/2005:229). This response typifies a whole family of views. Views differ in particular on the status they assign to the ways things appear to the observers: are they perhaps not even properties that things appear to have (but then how to characterize the incompatibility between the observations?), or are they properties that we simply misattribute to things (but then whence this systematic error?), or are they properties that are somehow secondary or less-than-fully real yet such that we are somehow correct in attributing them to things (but then how to make sense of this metaphysical status hovering between the real and unreal?).¹ The finer-grained differences aren't relevant here. All that matters for our purposes is that any version of the standard response will at least deny, somehow, that things genuinely are the ways they appear to the observers.

The standard response cannot be avoided by saying that the apparent properties or relations merely turn out to have a higher adicity - that these cases simply reveal a hidden argument place. There is no such thing as 'increasing the adicity' of a property whilst leaving it in place. We really just deny that the apparent n -adic property or relation is instantiated and replace it with a different $n+1$ -adic relation (see Boghossian 2006). The latter is distinct from the former, not just because it has a different adicity, but also because the $n+1$ -adic relation doesn't characterize the way things appear to us. Sophie and Nora appear in incompatible ways. If it simply

¹See Rosen (1994) and Stroud (2000).

appeared to you that *Sophie is on the left of Nora relative to you*, and it appeared to me that *Sophie is on the right of Nora relative to me*, then there would be no incompatibility in the ways things appeared to us. But there *is* such an incompatibility, and hence those three place relations aren't the relevant ways things appear to be to us. Sophie and Nora appear to have the two-place relations, and the standard response is to deny that these are instantiated.

Why do we think that something cannot have incompatible properties, such as being round and being square? It's not a simple question of logical form. These facts have the form of '*a* is F' and '*a* is G' and there's nothing in these syntactic forms from which we can read their incompatibility. And yet we somehow see that something could not be both round and square.

Perhaps we think that something cannot be both round and square because we cannot imagine what such an object would be like. If we have an image of *a* as round and an image of *a* as square, we cannot combine these two images in a single more detailed image of *a*. Note that we here assume that the world is a unified place: the images need to be deemed squeezable into *one* before their collective contents can be taken to reflect what things are possibly like.

But now ask yourself this: why couldn't the two images just collectively depict the way the world is, why is the inconceivability of a conjoined image evidence that at most one of the images can reflect what *a* is like, why couldn't the plurality of images *as such* be representative of the world? Why does the world necessarily have to be a metaphysically unified place?

I believe that we are able to adopt an alternative conception according to which the world can be disunified. A fragmented world is only adequately represented by a multiplicity of images and, within such a plurality, two images may have incompatible contents: there may be an image in which *a* is round and an image where *a* is square, or there may be an image where Sophie is on the left of Nora, and an image where Sophie is on the right of Nora. A fragmented world is a world that corresponds to multiple perspectives on the world, indeed, we can say that a fragmented world is a perspectivally constituted world. How this is so will, hopefully, become clearer throughout

this chapter and indeed, throughout the dissertation.

In order to make sense of a fragmented world, we must first learn to distinguish (1) the case in which A and B merely each obtain, and (2) the case in which they not only each obtain but also co-obtain. Again, we do not normally distinguish these cases. We assume that when A and B each obtain, they *thereby* co-obtain. This leaves no room for the notion of co-obtainment that we require to make sense of a fragmented world. This notion of co-obtainment must be such that it may be the case that two facts both obtain without co-obtaining in the relevant sense. From now on, when I talk of the co-obtainment of facts, I will be referring to co-obtainment in its newly introduced sense.

The talk of co-obtaining *facts* should ultimately be dispensable. It would be odd to think that we can make sense of a fragmented world only if we are committed to an ontology of facts. Instead of attributing a co-obtainment relation to facts, our official view will express the co-obtaining of facts directly, using a sentential connective ‘ \circ ’. I propose we read, for instance, ‘Sophie is sitting \circ Nora is sitting’ as ‘Sophie is sitting *insofar as* Nora is sitting’. This reading is admittedly not ideal. In ordinary language, the phrase ‘insofar as’ has a bare conjunctive reading but also an asymmetric reading on which the left side is in some sense explained by the right side - and it’s only the conjunctive reading that fits (though certainly doesn’t exhaust) the theoretical sense the notion is intended to have. From now on ‘insofar as’ is used only in a conjunctive sense.² I will also occasionally resort back to talk about facts, for ease of expression, in the understanding that this can be translated into talk about the involved objects.

What is it for facts to co-obtain; and more importantly, what is it for facts to both obtain without co-obtaining? I do not believe that there is any definition of co-obtainment; if it latches onto something in the world, it latches onto a *sui generis* metaphysical structure. This means that the fragmentalist incurs the important burden of elucidating the notion in other

²Why not simply read ‘Sophie sits \circ Nora sits’ as ‘Sophie’s sitting *co-obtains with* Nora’s sitting’? One may do this of course but it requires we transform the embedded sentences into imperfect gerundial nominals, such as *the tree’s being leafless*. This makes for awkward sentences and leads away from straightforward talk about the involved object.

ways. As mentioned in the introduction, I believe such elucidation proceeds on the basis of paraphrases, metaphors, analogous cases, and models.

We can start close to home and think, not of a fragmented world, but of a fragmented body of information. As Lewis pointed out, it's plausible to think that our own beliefs are like this:

I speak from experience as the repository of a mildly inconsistent corpus. I used to think that Nassau Street ran roughly east-west; that the railroad nearby ran roughly north-south; and that the two were roughly parallel. (By 'roughly' I mean 'to within 20°'). So each sentence in an inconsistent triple was true according to my beliefs, but not everything was true according to my beliefs. Now, what about the blatantly inconsistent conjunction of the three sentences? I say that it was not true according to my beliefs. My system of beliefs was broken into (overlapping) fragments. Different fragments came into action in different situations, and the whole system of beliefs never manifested itself all at once. The first and second sentences in the inconsistent triple belonged to - were true according to - different fragments; the third belonged to both. [...] I think the same goes for other corpora in which inconsistencies are successfully quarantined. The corpus is fragmented. Something about the way it is stored, or something about the way it is used, keeps it from appearing all at once. It appears now as one consistent corpus, now as another. The disagreements between the fragments that appear are the inconsistencies of the corpus taken as a whole. (Lewis 1982: 436).

I think that Lewis's description is plausible: we may believe each of an inconsistent pair of sentences without, in some sense, believing them together. And, more generally, it's not hard to see how bodies of information can be fragmented in a certain sense. The collection of pictures was just another instance of it.

Now think of the world itself as being just like such a fragmented body of information. Think of the world as a fragmented corpus of facts instead of a

fragmented body of information about the facts. Something about the nature of the world, about the way the facts obtain, keeps the facts from obtaining all at once or all together. The world appears now as one consistent corpus of facts, now as another. The disagreements between the fragments are the inconsistencies of the world taken as a whole.

In addition to the analogy between fragmented systems of information and a fragmented world, we can resort to direct paraphrases. We may think of the co-obtainment of facts as a type of metaphysical ‘glue’ holding between the facts. When facts co-obtain, we could say that they form a unified qualitative manifestation of the objects involved, a single situation within which things are a certain way. Indeed, in the case when two facts involve the same object and the instantiations of two distinct properties, it’s the co-obtaining of these facts that is naturally understood as the *co-instantiation* of the relevant properties. For example, ‘*a* is round \circ *a* is blue’ describes *a* as co-instantiating *being round* and *being blue*. Whenever, on the other hand, *a* is round and blue but not round *insofar as* it is blue, this is aptly imagined by imagining *a* as round but not blue and imagining *a* as blue but not round; by imagining these facts one after the other (see figure 1.1).

The loose talk of co-obtainment as the composition of a unified qualitative scene can be taken further. If one is comfortable with the talk of facts, one may also find it helpful to initially think of a fragmented world in mereological terms. Let $[A]$ be a name for *the fact that A*. One can liken the co-obtaining of *A* and *B*, i.e. it being the case that $A \circ B$, to the existence of a binary sum $[A]+[B]$. A fact is then, as it were, part of a larger fact and the unity assumption corresponds to a principle of unrestricted composition applied to facts: any collection of facts whatsoever compose a fact. In contrast, the idea that a world can be fragmented is analogous to the rejection of the unrestricted composition of facts. The world is fragmented when there is a collection of facts that may fail to compose a single fact; there are pairs of facts which are such that there is not a single whole that they are both parts of. They obtain, but within the context of distinct facts. Co-obtaining facts compose a unified chunk of world, as it were, whereas facts that obtain but fail to co-obtain fail to compose a unified chunk of world.

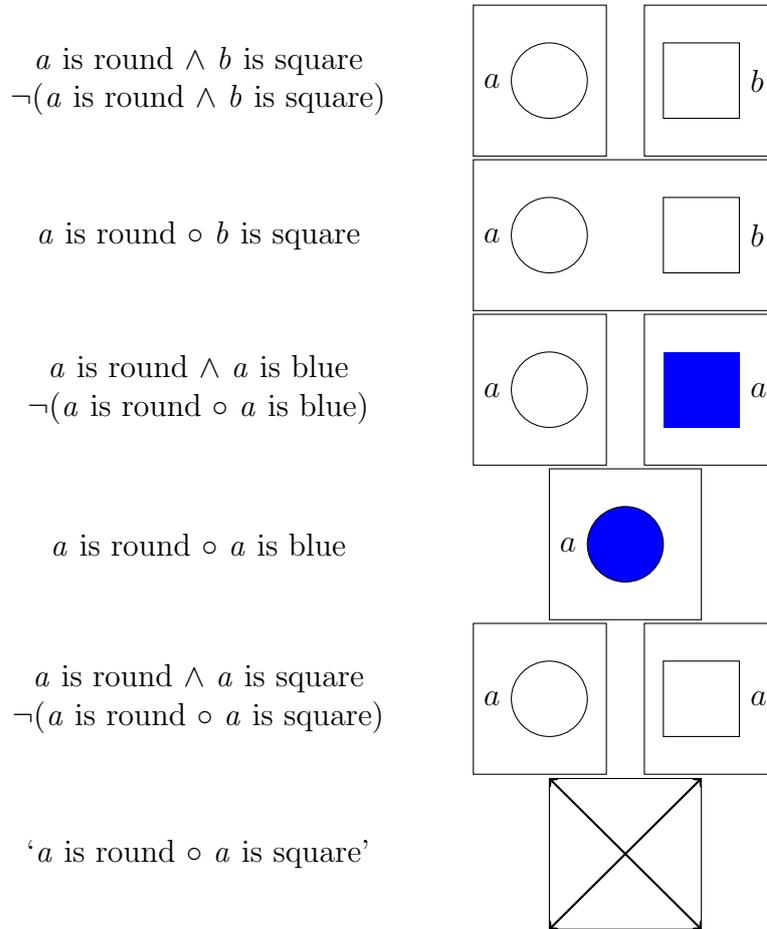


Figure 1.1: Fragmentation and co-obtainment

Finally, and this is probably the most helpful way of coming to grips with co-obtainment, we can think of fragmentation in terms of ‘perspectival facts’. To return to the case of Sophie being on the left and right of Nora, we can describe this situation as:

Sophie is on the left of Nora \wedge Sophie is on the right of Nora \wedge \neg (Sophie is on the left of Nora \circ Sophie is on the right of Nora).

Sophie is on the left of Nora and Sophie on the right of Nora, but Sophie is *not* on the left of Nora *insofar as* Sophie is on the right of the Nora. We are only ever able to observe one of the facts. When I observe Sophie as being on the right of Nora, and I walk around them to see them from your

side, it seems as if the earlier observed fact is no longer there. The two facts are like two sides of a coin, one side can be facing you only insofar as the other doesn't. Facts that fail to co-obtain obtain in a perspectival way: one obtains only insofar as the other doesn't.

We noted that the incompatibility of facts isn't always a matter of logical form, as in the case of '*a* is round' and '*a* is square'. But sometimes it is, as in the case of '*a* is round' and '*a* is *not* round'. Call pairs of sentences one of which is the negation of the other *contradicting* sentences. If we not only allow that *a* is round and that *a* is square, but also that *a* is round and *not* round, then we accept a metaphysics that requires revisions in our logic, given that - according to classical logic - a true contradiction implies that everything is the case. I don't think that we need to take this route. We should distinguish between two kinds of conflicting facts:

'*A*' and '*B*' state *contrary* facts iff they cannot both obtain, i.e. iff necessarily $\neg(A \wedge B)$.

'*A*' and '*B*' state *incompatible* facts iff they cannot co-obtain, i.e. iff necessarily $\neg(A \circ B)$.

It seems clear that contradictory sentences always state contrary facts - even for the fragmentalist: nothing can possibly be round and *not* round. Contradictory sentences stand in contrast to sentences such as '*a* round' and '*a* is square', neither of which is the negation of the other, and yet which clearly express facts that conflict in some way. We adopt a fragmentalist view when we think that such sentences state incompatible facts, and not contrary facts, and hence that they can both obtain (though not co-obtain).

It's highly implausible that all conflicting facts are only expressed by sentences which are really of a contradictory logical form - which would leave no room for a fragmentalist view about any matters. Wittgenstein held this radical reductive claim in the *Tractatus* (1921/1961): he believed that when the truth of one seemingly atomic sentence seems to constrain the truth of another seemingly atomic sentence, philosophical analysis will reveal a definition of the terms involved according to which the sentences do after all

express contradictory propositions. So, in the case of ‘the rose is red’ and ‘the rose is blue’, the colour terms would have to admit of definitions that show that they do not - contrary to what their surface structure suggests - express atomic propositions but logically complex propositions that contradict one another (6.3751). But it’s implausible to believe that such definitions will always be forthcoming, indeed, that this is implausible is the well-known ‘exclusion problem’ and played a crucial role in Wittgenstein’s own abandonment of his Tractarian view, see in particular his ‘Some Remarks on Logical Form’ (1929). Without the implausible reductive claim about philosophical analysis, we all already admit a distinction between incompatible and contradictory matters. The fragmentalist does not introduce the contrast between incompatible sentences and contradictory ones, it’s already there in any plausible philosophy of logic. The fragmentalist only takes the distinction further, by revising our understanding of metaphysical incompatibility whilst leaving our understanding of contradictions in place.

We noted already that all contradictory sentences state contrary facts, but it’s also natural to hold that - vice versa - all contrary facts are stated by contradictory sentences. Note however that this isn’t obligatory. If a fragmentalist believes that ‘the box is full’ and ‘the box is empty’ state contrary facts, this doesn’t force the fragmentalist to also believe that the one sentence expresses the negation of the other - she may simply be convinced that it’s impossible that a box is both empty *and* full whereas, in contrast, she may be convinced that it’s possible that the box is red all over *and* blue all over. It’s not an in-built commitment of the fragmentalist framework that all contrariety is reflected in the logical form of the sentences stating those facts.

Having said this, though, when we accept that there are merely incompatible facts, it *does* become plausible to think that contrariness *always* results from the involvement of negation. A box cannot be both empty *and* full, the fragmentalist may think, and so ‘is empty’ means ‘is not full’ (or ‘is full’ means ‘is not empty’). One might be worried that this commitment, if we accept it, makes fragmentalism overly sensitive to the logical form of sentences, given that sentences may only state merely incompatible facts when

their logical form is not contradictory. But this worry gets things the wrong way around. Whether we think that the box can be both full and empty does not depend on what we think the correct logical form is of the sentences ‘the box is full’ and ‘the box is empty’, rather, - *if* the fragmentalist commits herself to the claim that all contrary facts are stated by contradictory sentences - then what logical form should be attributed to the sentences depends on whether she thinks the box can be both full and empty. How do we decide whether ‘the box is full’ and ‘the box empty’ express atomic predications or, respectively, the proposition *that the box is not empty* or the proposition *that the box is not full*? It seems to me that a major role is played by our metaphysical view of the world. To give a well-worn analogy: we think that the truth of the sentence ‘Nobody arrives’ is not a matter of some object, Nobody, having a certain property, arriving. We think of the truth of the sentence as consisting in the fact that there is no person that arrives. The logical form we attribute to the content expressed by ‘Nobody arrives’ is in this way informed by our metaphysical views, in particular by the view that the world can underwrite the truth of the sentences without containing an object that is Nobody. Similarly, how ‘the box is full’ and ‘the box empty’ are eventually to be understood will depend on whether we think their truth is plausibly understood in terms of the box instantiating two properties, or whether one sentence is rather the denial of the instantiation of some property. So, if two facts are contrary whenever they are stated by contradictory sentences, this does not make the fragmentalist’s metaphysical view depend on the logical form of sentences, it rather constrains what logical form she will attribute to sentences.

What about compatible facts, such as *a*’s being round and *a*’s being orange, do they necessarily co-obtain? I see no reason why they should. Think back to the plurality of images: one image may show *a* as round and green, and another may show *a* as square and orange, so that *a*’s being round and *a*’s being orange co-obtain as little as *a*’s being round and *a*’s being square. We can distinguish the possibility of compatible facts co-obtaining from the possibility in which they each obtain but fail to co-obtain, and this means that their co-obtaining must be a contingent and substantive matter.

This means that there are two ways in which modal space turns out to be richer than we normally assume it to be. First: for any two compatible facts, such as a 's being round and a 's being orange, it's possible that these co-obtain and it's possible that they fail to do so, i.e. we have it that $\diamond(a \text{ is round} \circ a \text{ is orange})$ and $\diamond\neg(a \text{ is round} \circ a \text{ is orange})$.³ And, second: for any two incompatible facts, such as a 's being round and a 's being square, it's possible that they both obtain but impossible that they co-obtain, i.e. we have it that $\diamond(a \text{ is round} \wedge a \text{ is square})$ though $\neg\diamond(a \text{ is round} \circ a \text{ is square})$. When we try to imagine a round square, and fail, what we try to imagine is something that is round *insofar as* it is square. The fact that this is inconceivable, and indeed impossible, doesn't tell against it being possible that something is round *and* square without also being the one *insofar as* it is the other.

This diverges radically from our common conception of the world. At the very least we do not think our world to be like that; we do not think that there are cases where incompatible facts obtain. We can now see that this is a substantive assumption about the actual world, which I will call the unity assumption:

The unity assumption: necessarily, if two facts obtain they co-obtain, i.e necessarily, if A and B , then $A \circ B$.

The world is assumed to be a unified place, and this leaves no room for incompatible facts: if two incompatible facts both obtain then, given the unity assumption, they co-obtain, but incompatible facts cannot co-obtain and hence, there can be no incompatible facts.

Our standard conception of the world is based on the unity assumption: the unity of the world is presumed and the denial of incompatible facts is deemed to be rationally demanded of us. The unity assumption is moreover not recognized for what it is, a substantive assumption about the metaphysical structure of reality. That it is indeed a substantive assumption should

³The diamond ' \diamond ' stands for metaphysical possibility, which behaves as it ordinarily does. Modal space simply includes fragmented possible worlds, i.e. worlds at which various co-obtainment claims are true.

be granted by anyone who can make sense of the world being fragmented in the sense proposed here.

Someone might want to counter this claim about the possible fragmentation of the world with a tollensing move: if it's indeed possible that the world is fragmented, then it's possible that two incompatible facts obtain (that $\Diamond(p \wedge q)$, where p and q are incompatible), but it isn't possible that two incompatible facts obtain and hence it isn't possible that the world is fragmented. This objection is misguided however. If there is a way of making sense of a view, it is dialectically dubious to refuse to make sense of it, and then deny the view because it does not make sense. The proposal is to understand the incompatibility in terms of the impossibility of the co-obtainment of them to obtain (and not, as usual, as the impossibility of the conjunction of the them). Allowing that it's metaphysically possible that facts obtain without co-obtaining thus explains how incompatible facts can obtain. Given this explanation, the proposal cannot be rejected on the grounds that it's unintelligible that two incompatible facts obtain. We don't have a stalemate here, the fragmentalist has the upper hand: the fact that the orthodoxy sees no difference between conjunction and co-obtainment as introduced here explains both why the tollensing move can sound reasonable as well as why it really isn't.

One may wonder whether a framework like this should not simply accept the truth of contradictions and, if not, why not. After all when you observe Sophie and Nora, one might say that it doesn't *just* appear to you that Sophie is on the left of Nora but also that it appears to you that Sophie is *not* on the right of Nora. To answer this and related worries, we need to have a closer look at the inferential role of co-obtainment.

1.2 A basic semantics for co-obtainment

We can further solidify our understanding of co-obtainment by specifying the inferential role that the notion of co-obtainment plays. I know of at least two ways to go about this. We can directly lay down the formal features that

seem plausible for the notion, and simply stipulate that co-obtainment is to be commutative, non-transitive, etc. We then arrive at a conception of a fragmented world through the co-obtainment facts thus stipulated. I find this first approach risky. There is no guarantee that the intuited formal features indeed come together to form the exact notion that captures the presumed fragmentation of the world.⁴ Besides, given that this is quite an unfamiliar conception of the world, on what basis are we to intuit that co-obtainment is to have this or that formal feature? We are not already accustomed to reason in terms of co-obtainment as it is intended here. It is thus less risky to start with models of fragmented worlds and specify the truth conditions of co-obtain in such a way that it encodes the fragmentation that we are after. Using these models, we can then read off how the co-obtainment interacts with the logical connectives. To keep things as simple as possible, I will only set out the semantics of a very basic propositional logic that allows for the fragmentation of the world. I discuss the interaction between co-obtainment and metaphysical possibility in the first appendix to this chapter (§1.4. I discuss a first-order logic in the second appendix to this chapter (§1.5), for those who want to see a more detailed picture. The semantics goes in many ways back to the ‘discussive logic’ of Jaśkowski (1948/1969) - with the important difference that Jaśkowski’s discussive logic is paraconsistent, whereas the logic below isn’t.⁵

We already saw that a useful toy model for a fragmented world is a collection of images that only depict all of reality collectively: something is the case when it’s depicted in one of the images in the collection, and isn’t the case when not depicted by any of them. These images allow us to make

⁴What is worse, we may jumble together bits of inferential role that constitute an empty notion. Prior’s ‘tonk’ (1960) illustrates this nicely: if we stipulate that we can infer from A , that A -tonk- B and that we can infer from A -tonk- B that B , then this introduced notion, call it tonk, allows us to prove anything from anything. As not everything is the case; we know that no ‘tonk facts’ obtain. Note that, when we deem the notion to be empty it seems we appeal directly to our conception of what the world is like, namely a conception on which not everything whatsoever is true, and we do not appeal to combination of the bits of inferential role.

⁵For other closely related logics, see Rescher and Brandom (1980), Priest (2008) and, in particular, Restall (1997). For an accessible introduction to non-adjunctive logic, see Priest (2007: §4.2), and Varzi (1997).

sense of incompatible facts given that one might depict me as sitting, for example, whereas another depicts me as standing. Where I'm depicted as sitting, I'm not also depicted as standing: so for two facts to be depicted *in the same picture* is a good model for the co-obtaining of the depicted facts. We will now abstract from the inessential aspects of this toy model (such as its consisting of multiple images), to arrive at the following models.

Let the set of sentences S consist of atomic sentences p, q, r, \dots and be such that, if A and B are sentences, so are $\neg A$, $A \wedge B$, $A \vee B$ and $A \circ B$ (besides these, nothing else is in S). A model M is a pair $\langle W, v \rangle$, where W is a set of points and v is a function that assign either 1 or 0 to each of the atomic sentences relative to each point w in W .

The valuation v for the atomic sentences relative to points is extended to a valuation for all the sentences via the following recursive clauses (where w ranges over points in W):

$$v_w(A \circ B) = 1 \text{ iff } v_w(A) = 1 \text{ and } v_w(B) = 1$$

$$v_w(A \wedge B) = 1 \text{ iff } v_w(A) = 1 \text{ and } v_w(B) = 1$$

$$v_w(A \vee B) = 1 \text{ iff } v_w(A) = 1 \text{ or } v_w(B) = 1$$

$$v_w(\neg A) = 1 \text{ iff } v_w(A) \neq 1$$

Truth in a model, written $M \Vdash A$, is defined via the following recursive clauses (where p is an arbitrary atomic sentence):

$$M \Vdash p \text{ iff } \exists w (v_w(p) = 1)$$

$$M \Vdash A \circ B \text{ iff } \exists w (v_w(A \circ B) = 1)$$

$$M \Vdash A \wedge B \text{ iff } M \Vdash A \text{ and } M \Vdash B$$

$$M \Vdash A \vee B \text{ iff } M \Vdash A \text{ or } M \Vdash B$$

$$M \Vdash \neg A \text{ iff } M \not\Vdash A$$

We define validity and logical truth as follows (where Σ is a set of sentences):

An argument from Σ to A is *valid*, written $\Sigma \vDash A$, iff, for every model M , if $M \Vdash \Sigma$ then $M \Vdash A$.⁶

A formula A is *logically true*, written $\vDash A$, iff, for every model M , $M \Vdash A$.

Note that the points correspond to the images of the toy models.⁷ The v -clauses say what is true and false of each them. There are v -clauses for conjunction, disjunction, negation and co-obtainment because the logic needs to handle the embedding of logically complex sentences in co-obtainment sentences. For example, $A \circ (B \circ C)$ is true in a model if A and $B \circ C$ are true *at a point*, and this requires that $B \circ C$ has a truth-value *at points*. The same applies to $A \circ (\neg B)$, $A \circ (B \vee C)$ or $A \circ (B \wedge C)$.

Before discussing things in more detail it may help to first run through the machinery stepwise. Start with the evaluations of atomic sentences relative to points and run the standard Boolean evaluations for complex sentences relative to points, treating co-obtainment as conjunction. The atomic sentences and co-obtainment sentences that are true at some point in W are all true in the model. And these truths are the basis for standard Boolean evaluations of negations, conjunctions and disjunctions in the model: taking negations of things that are false in the model to be true in the model, conjunctions of two truths to be true in the model, etc.

We can use the machinery to see more clearly how the fragmentalist doesn't require a paraconsistent logic, and to see more clearly what formal features co-obtainment has. Let me start with the former.

A sentence is true *at a point* if and only if its negation isn't true there. This means that we never have a point where both a sentence and its negation are true. That is:

$$\vDash \neg(A \circ \neg A)$$

⁶By $M \Vdash \Sigma$ I mean that $M \Vdash B$ for all $B \in \Sigma$.

⁷Alternatively, the points can be interpreted as the possible worlds known from standard modal logics, so that a single fragmented world (here represented by a single model) corresponds to a set of possible worlds in a frame of modal logic; cf. Restall (1997).

It cannot be the case that something obtains *insofar as* it doesn't obtain.

Similarly, any sentence in our language is *true in a model* if and only if its negation isn't true *in the model*. This means that the law of excluded middle and the law of non-contradiction hold:⁸

$$\models A \vee \neg A$$

$$\models \neg(A \wedge \neg A)$$

To illustrate an important consequence of this, consider a model where we have w_1 at which p is true but q isn't, and w_2 at which q is true but p isn't. As there are points at which p and q are true, they are true in the model. This means that $\neg q$ isn't true in the model. But given that $\neg q$ is true *at* w_1 , $p \circ \neg q$ is true in the model. So here, $\neg q$ is false, yet true *insofar as* p is true (i.e. $\neg q$ is false but $p \circ \neg q$ is true). The fragmentation gives rise to negative sentences being true *insofar* certain other things are true, even though they are false. You observe that Sophie is *not* on the right of Nora, but this is the case only *insofar as* Sophie is on the left of Nora.

It's important to note that this captures a natural understanding of negation in this context, reflecting how we think of local absences versus global absences. We can compare it with the way existence-at-a-location and existence simpliciter interact: an object exists when there is a location at which it exists but it doesn't fail to exist when there is a location at which it doesn't exist. An object doesn't exist only when there is *no* location at which it exists. Whereas local existence suffices for global existence, local non-existence doesn't suffice for global non-existence. There is a natural asymmetry here. The same sort of asymmetry is found, in a fragmented world, between atomic facts and their negations. When an atomic fact obtains *insofar as* other facts obtain (or 'within a fragment'), this suffices for it to obtain simpliciter (or 'within the world as such'), but when the fact is absent *insofar as* other facts

⁸We also have:

$$B \models A \vee \neg A$$

$$A \wedge \neg A \models B$$

obtain (or absent ‘within a fragment’), this does not suffice for the fact to be absent simpliciter (or absent from the world at large). Put more directly: if a is F *insofar as* b is G, this suffices for a to be F; but if a is not F *insofar as* b is G this does not suffice for a not to be F because it may still be F *insofar as*, say, c is H.⁹ There is no reason why the fragmentalist should have a non-standard understanding of negation; and hence no reason why the fragmentalist requires a paraconsistent logic or admit true contradictions.¹⁰

It can easily be checked that the semantics is sound with regard to the deductive reasoning of classical propositional logic. To prove that the semantics is sound with regard to each natural deduction rule we only need to appeal to the truth in a model clauses for the relevant connectives, and the clauses above are just those of any standard semantics for classical sentential logic. They are the standard Boolean truth conditions. And so any soundness proof for a standard semantics for sentential logic will work here.¹¹

Let me now turn to elucidating the formal properties and inferential role of co-obtainment. It can easily be seen from the semantics that co-obtainment is commutative and associative:

$$A \circ B \models B \circ A$$

⁹Compare this to Restall (1997), who takes an atomic sentence to be false in the model when false at a point, and consequently arrives at a paraconsistent logic (which, to be sure, is precisely what he aims for).

¹⁰To be sure, this is not to say that one could not develop the notion within a paraconsistent framework; it’s just to say that there is nothing in the intuitive picture that requires or of itself leads to such a treatment.

¹¹To illustrate, consider conjunction introduction (\wedge I) and conjunction elimination (\wedge E). (\wedge I): If $X \vdash A$ and $Y \vdash B$, then $X, Y \vdash A \wedge B$. Proof: If the last rule applied in a proof is \wedge I, it ends with $X, Y \vdash (A \wedge B)$ arrived at from $X \vdash A$ and $Y \vdash B$. Assume that $X \models A$ and that $Y \models B$ and let M be a model that satisfies $X \cup Y$. This is a model that satisfies X and thus satisfies A . And it is a model that satisfies Y and thus satisfies B . By the right-to-left direction of the clause for \wedge , M satisfies $A \wedge B$. So, if $X \models A$ and $Y \models B$, then $X, Y \models A \wedge B$.

(\wedge E): If $X \vdash A \wedge B$, then $X \vdash A$. If the last rule applied in a proof is \wedge E, it ends with $X \vdash A$ arrived at from $X \vdash A \wedge B$. Assume that $X \models A \wedge B$ and let M be a model that satisfies X . This is a model that satisfies $A \wedge B$ and thus, by the left-to-right clause for \wedge satisfies A . So, if $X \models A \wedge B$, then $X \models A$.

What drives the proofs are the clause for the truth of conjunctions in a model, and the proposed semantics features the standard Boolean clause. It can be easily checked that the same goes for the other logical connectives, and hence that the semantics are sound with regard to classical propositional logic.

$$A \circ (B \circ C) \not\equiv (A \circ B) \circ C$$

The commutativity of co-obtainment is underwritten by the fact that if A and B are true at some point, they aren't true in some particular order. To see why co-obtainment is associative, note that if we have $A \circ (B \circ C)$ this means that there is a point w where both A and $B \circ C$ are true (see the v -clause for \circ), and $B \circ C$ is only true at w if both B and C are true at w . This means that all three, A , B and C are true at w , which means that $A \circ B$ must be true at w together with C , and hence that $(A \circ B) \circ C$ is true in the model.

Co-obtainment fails to be idempotent:

$$A \circ A \not\equiv A$$

$$A \not\equiv A \circ A$$

To see why we have $A \circ A \not\equiv A$, consider a model where p is false at w_1 but true at w_2 . In this model, $\neg p \circ \neg p$ is true given that there is a point where each is true (viz. w_1). And yet $\neg p$ is not true simpliciter, given that p is true at w_2 . To see why we have $A \not\equiv A \circ A$, consider a model where p is true but q is false at w_1 and where p is false but q is true at w_2 . Here $p \wedge q$ is true in the model, but $(p \wedge q) \circ (p \wedge q)$ isn't true, as there is no single point at which $p \wedge q$ is true.

The fact that it's not a logical truth that a fact co-obtains with itself can seem worrying, after all, how on earth can a fact fail to co-obtain with itself? It seems however that the inference only fails for conjunctions. If $A \wedge B$ did co-obtain with itself, one expects it to characterize a *unified* bit of world as being such that $A \wedge B$ even though it may not be characterized by $A \circ B$ and so not not be a unified bit of world. This seems incoherence. The model theory can be complicated to accommodate the inference, but I'm not convinced that this is something that should be accommodated.

Relatedly, we have the failure of simplifying and adjunctive rules for co-obtainment:

$$A, B \not\equiv A \circ B$$

$$A \circ B \not\equiv A$$

For the failure of adjunction, consider a model where we have w_1 at which p is true and w_2 at which q is true. Here p is true and q is true because they are atomic sentences and there are points at which they are true. But $p \circ q$ isn't true, given that there is no point at which p and q are both true. The failure of adjunction is of course precisely what is required, and is in some sense the core feature of co-obtainment. For the failure of simplification, consider a model where p is true at w_1 and where q and $\neg p$ are true at w_2 . Here $q \circ \neg p$ is true in the model, but $\neg p$ is not true in the model, given that p is true at w_1 .

Co-obtainment is also non-transitive:

$$A \circ B, B \circ C \not\equiv A \circ C$$

Consider a model where we have a point w_1 at which p and q are true but r isn't, and a point w_2 at which q and r are true but p isn't. In such a model, $p \circ q$ and $q \circ r$ are true, but $p \circ r$ isn't. Put informally, the failure of transitivity allows fragments to overlap.

Some of the failures of these types of inferences are solely due to the fact that we allow any kind of sentence to be embedded in co-obtainment sentences. For example, co-obtainment sentences that only embed atomic sentences are idempotent, and there is a restricted form of simplification for atomic sentences embedded in co-obtainment claims:

$$p \circ p \not\equiv p, \text{ where } p \text{ is an atomic sentence.}$$

$$A \circ p \equiv p, \text{ where } p \text{ is an atomic sentence.}$$

For atomic idempotence, note that if one has a point at which p is true, one thereby has a point at which $p \circ p$ is true, and vice versa. For atomic simplification note that if $A \circ p$ is true, then there is a point w at which A and p are true, but that suffices for p to be true in the model.

I take the above to include the most important principles for the purpose of solidifying our inferential grasp of the notion of co-obtainment. There are of course many more principles governing the co-obtainment notion; one can, for example, easily check that the following hold:

$$A \circ B \vDash A \circ A$$

$$A \circ (B \circ C) \vDash B \circ C$$

$$A \circ (B \wedge C) \not\vDash B \wedge C$$

$$A \circ (B \wedge \neg B) \vDash C$$

$$A \circ B \vDash A \circ (B \vee C)$$

$$A \circ (B \vee C) \vDash (A \circ B) \vee (A \circ C)$$

$$(A \circ B) \vee (A \circ C) \vDash A \circ (B \vee C)$$

$$A \circ (B \wedge C) \vDash (A \circ B) \wedge (A \circ C)$$

$$(A \circ B) \wedge (A \circ C) \not\vDash A \circ (B \wedge C)$$

Note that we have complete distributivity for co-obtainment and disjunction but not for co-obtainment and conjunction. Next to these inferences, we may also note that, for any two sentences A and B , at least one of these four options obtains: $A \circ B$ or $A \circ \neg B$ or $\neg A \circ B$ or $\neg A \circ \neg B$.

It cannot be stressed enough that the model draws out the inferential role of co-obtainment in ways we happen to find elucidating; the set-theoretic machinery is merely a heuristic tool to elucidate the logical structure of the co-obtainment notion. Sentences aren't true 'relative to' or 'at points' in our object language, the language which reflects how the fragmentalist understands the world. The points do not correspond to anything in the fragmentalist's ontology. Certain facts obtain *insofar as* other facts do, the rest is but metaphor and model.

1.3 Summary and concluding remarks

To sum up, we have established the features of co-obtainment:

$$A \circ B \not\vDash B \circ A \quad \text{[Commutative]}$$

$A \circ (B \circ C) \not\equiv (A \circ B) \circ C$	[Associative]
$A \circ A \not\equiv A$	[Non-idempotent]
$A \not\equiv A \circ A$	[Non-idempotent]
$A \circ B, B \circ C \not\equiv A \circ C$	[Non-transitive]
$A, B \not\equiv A \circ B$	[No adjunction]
$A \circ B \not\equiv A$	[No simplification]
$\models \neg(A \circ \neg A)$	[LNCC]

Each of these typify characteristic features of the notion of co-obtainment as it will be understood here.

By no means is the inferential role of some notion all there is to meaning of that notion. But, at the very least, these principles constrain the application of the notion when certain truths are assumed to hold and thus helps us judge when the notion can be applied and when it can't. The point of the informal paraphrases we saw in the previous chapter (e.g. 'co-obtaining facts constitutes a single qualitative scene or manifestation of the objects involved') is precisely to offer some sense of the content of the notion that goes beyond its inferential role. And, if all is well, we hopefully continue to build our understanding of what it is for facts to co-obtain when we turn to concrete applications of the framework.

1.4 Appendix: incorporating modality

We already made certain modal claims concerning the co-obtainment of facts, and more will be made. There is thus some interest in the way modality interacts with the notion of co-obtainment. This is what I take to be the most straightforward understanding of this interaction: amongst the possible worlds, we now find fragmented possible worlds or, better, amongst the possibilities we find the possibilities of the world being fragmented in various

ways. One choice point that arises is whether metaphysical possibility admits of perspectival variance: whether something may be possible *insofar as* one thing is the case, but impossible *insofar as* something else is the case. I do not find such variance plausible, at least not for metaphysical possibility. After all, metaphysical possibility is possibility relative to the metaphysical laws, and it seems that something only counts as a *metaphysical* law precisely if it applies no matter what. I'm not aware of any phenomenon that suggests or requires metaphysical possibility to be a perspectival matter. And so, I believe that metaphysical possibility does not admit of perspectival variance: all fragments agree on what is metaphysically possible and metaphysically impossible. Below I offer a semantics that reflects this decision. As I will aim to capture modality as characterized by S5, I will not use any accessibility relation in the semantics.

Let the set of sentences S consist of atomic sentences p, q, r, \dots and be such that, if A and B are sentences, so are $\neg A$, $A \wedge B$, $A \circ B$ and $\diamond A$ and $\Box A$ (besides these, nothing else is in S). A model M is a pair $\langle \mathcal{F}, \mathcal{W}, v \rangle$, where:

\mathcal{F} is a set of points,

\mathcal{W} is a set consisting of non-empty subsets of \mathcal{F} and is such that $\bigcup \mathcal{W} = \mathcal{F}$,

v is a function that assigns either 1 or 0 to each of the atomic sentences relative to the points ϕ in \mathcal{F}

One can think of \mathcal{F} as the set of all possible fragments, and of \mathcal{W} as all the possible worlds that can be built from these possible fragments.¹²

The valuation v for the atomic sentences relative to points ϕ in \mathcal{F} is extended to a valuation for *all* sentences relative to each point ϕ in \mathcal{F} as well as to each set ω in \mathcal{W} via the following recursive clauses (where ω ranges over points in \mathcal{W} and ϕ ranges over points in \mathcal{F}):

¹²Do not confuse the set \mathcal{W} in the current semantics with W in the basic semantics of the previous section; whereas the former represents possible worlds, the latter represented the fragments. Though it may have been helpful to think of the fragments as being like worlds in the basic semantics, that becomes confusing in the current setting.

$$v_\phi(A \circ B) = 1 \text{ iff } v_\phi(A) = 1 \text{ and } v_\phi(B) = 1$$

$$v_\phi(A \wedge B) = 1 \text{ iff } v_\phi(A) = 1 \text{ and } v_\phi(B) = 1$$

$$v_\phi(\neg A) = 1 \text{ iff } v_\phi(A) \neq 1$$

$$v_\phi(\diamond A) = 1 \text{ iff } \exists \omega (v_\omega(A) = 1)$$

$$v_\phi(\Box A) = 1 \text{ iff } \forall \omega (v_\omega(A) = 1)$$

$$v_\omega(p) = 1 \text{ iff } \exists \phi \in \omega (v_\phi(p) = 1)$$

$$v_\omega(A \wedge B) = 1 \text{ iff } v_\omega(A) = 1 \text{ and } v_\omega(B) = 1$$

$$v_\omega(\neg A) = 1 \text{ iff } v_\omega(A) \neq 1$$

$$v_\omega(\diamond A) = 1 \text{ iff } \exists \omega_1 (v_{\omega_1}(A) = 1)$$

$$v_\omega(\Box A) = 1 \text{ iff } \forall \omega_1 (v_{\omega_1}(A) = 1)$$

$$v_\omega(A \circ B) = 1 \text{ iff } \exists \phi \in \omega (v_\phi(A \circ B) = 1)$$

Note that the truth at a world clauses v_ω for the non-modal sentences correspond to the truth in a model clauses of the basic semantics, and that the truth at a fragment clauses v_ϕ for the non-modal sentences correspond to the truth at a fragment clauses in the basic semantics (cf. §1.2).

We define validity and logical truth as follows (where Σ is a set of sentences):

An argument from Σ to A is *valid*, written $\Sigma \vDash A$, iff, for every model M and every ω in \mathcal{W} : if $v_\omega(\Sigma) = 1$, then $v_\omega(A) = 1$.¹³

A formula A is *logically true*, written $\vDash A$, iff, for every model M and every ω in \mathcal{W} : $v_\omega(A) = 1$.

¹³By ' $v_\omega(\Sigma) = 1$ ' I mean that $v_\omega(B) = 1$ for all $B \in \Sigma$.

This completes the model theoretic specification of the logic.

Note that the clause for the truth of $\Diamond A$ at a world ω is the same as the clause for the truth of $\Diamond A$ at a fragment ϕ . This reflects the choice that metaphysical possibility is not a perspectival matter and ensures that no fragment disagrees with any other on what is possible (more on this below). The reason why the valuation v is extended simultaneously to a valuation of sentences relative to points in \mathcal{F} and worlds in \mathcal{W} is that we can have co-obtainment claims concerning possibilities (which are determined by what modal space is like) and modal claims concerning co-obtainment (which are determined by the fragments of the possible worlds).

The clauses may seem to have an air of circularity about them, after all, sometimes truth relative to worlds (v_ω relies on truth relative to fragments (v_ϕ and sometimes truth relative to fragments relies on truth relative to worlds. But there is no real circularity here. It's helpful to see the assignment of truth values in stages. At the first stage, we have the given truth values of atomic sentences relative to all possible fragments. This provides us with truth values of logically complex sentences that do not involve or embed diamonds or boxes anywhere (but embed only co-obtainments, conjunctions, etc.), relative to fragments and relative to the worlds that include these fragments. This provides us in turn with truth values of embeddings of these non-modal sentences under diamonds and boxes, relative to fragments and relative to the worlds that include these fragments. This provides us in turn with truth values of co-obtainments, conjunctions and negations of these diamond and box-embedded sentences, relative to fragments and relative to the worlds that include these fragments. This in turn provides us with truth values of embeddings of such complexes under diamonds and boxes again, relative to fragments and relative to the worlds that include these fragments. And so on. For any sentence, there is a definite stage at which it's assigned a truth value and whenever it's assigned a truth value at some stage we never return to it to assign it a different truth value. There is no circularity.

To get a feel for how the machinery works consider how the model determines the truth value relative to a world of the following four kinds of sentences:

$\diamond(p \circ q)$. This is true at a world ω , i.e. $v_\omega(\diamond(p \circ q)) = 1$ whenever there is a world ω_1 such that $\exists \phi \in \omega_1 (v_\phi(p) = 1 \text{ and } v_\phi(q) = 1)$. Put more informally: $\diamond(p \circ q)$ is true at a world if there is a world that includes a fragment at which both p and q are true.

$\diamond p \circ \diamond q$. This is true at a world ω , i.e. $v_\omega(\diamond p \circ \diamond q) = 1$, when there is a point ϕ in ω such that $v_\phi(\diamond p) = 1$ and $v_\phi(\diamond q) = 1$. Now, $v_\phi(\diamond p) = 1$ if there is world ω_1 such that $\exists \phi_1 \in \omega_1 (v_{\phi_1}(p) = 1)$ and $v_\phi(\diamond q) = 1$ if there is a world ω_2 such that $\exists \phi_2 \in \omega_2 (v_{\phi_2}(q) = 1)$. Put informally: $\diamond p \circ \diamond q$ is true at a world when there is a world that includes a fragment at which p is true and there is a world that includes a fragment at which q is true.

$\Box(p \circ \diamond q)$. This is true at a world ω , i.e. $v_\omega(\Box(p \circ \diamond q)) = 1$, when every world ω is such that $v_\omega(p \circ \diamond q) = 1$, which is the case when every world ω is such that there exists a ϕ in ω such that $v_\phi(p) = 1$ and $v_\phi(\diamond q) = 1$, and the latter is the case when there is a world ω_1 such that $v_{\omega_1}(q) = 1$. Put informally: $\Box(p \circ \diamond q)$ is true at a world when every world includes a fragment at which p is true and there is a world that includes a fragment where q is true.

$\neg \diamond \neg A$. This is true at a world ω when $\diamond \neg A$ is false at ω , which is the case when there is no world at which $\neg A$ is true. There is no world at which $\neg A$ is true whenever there is no world where A is false. That is to say, there is no world at which $\neg A$ is true whenever, for every world ω_1 , $v_{\omega_1}(A) = 1$. Informally: $\neg \diamond \neg A$ is true at a world when A is true at all worlds. (This shows that we can define ‘necessarily A ’ (‘ $\Box A$ ’) as $\neg \diamond \neg A$).

Note that, depending on the main connective of the sentence we encounter, we either look at worlds or the fragments included in some world.

Note also that if it’s possible that A *insofar as* something else, B , obtains, then it’s thereby possible that A :

$$(\diamond A) \circ B \models \diamond A$$

The reason is that $(\Diamond A) \circ B$ is true at a world if it includes a fragment where both $\Diamond A$ and B are true, but there is only a fragment where $\Diamond A$ is true if there is a world ω where A is true, and that also suffices for $\Diamond A$ to be true at the world the fragment is part of.

Vice versa, if it's possible that A , and some B obtains insofar as C obtains, then both B (and C) and the co-obtainment of B and C co-obtain with the possibility that A . That is:

$$\Diamond A, B \circ C \models (\Diamond A) \circ B \text{ (and } \Diamond A, B \circ C \models (\Diamond A) \circ C)$$

$$\Diamond A, B \circ C \models \Diamond A \circ (B \circ C)$$

If there is a world ω at which A is true, then it's true at each possible fragment that there is a world ω at which A is true, and hence $\Diamond A$ is true at each fragment. If $B \circ C$ is true, then there is also a fragment where B and C is true and, given that $\Diamond A$ is true at each fragment, $\Diamond A$ will be true at the fragment where B and C are true, and hence co-obtain with each of them, as well as with their co-obtainment $B \circ C$.

For much the same reason, if two things are each possible, then the one is possible insofar as the other is possible:

$$\Diamond A \wedge \Diamond B \models \Diamond A \circ \Diamond B$$

If $\Diamond A \wedge \Diamond B$ then A is true at some world and hence $\Diamond A$ is true at all fragments, and similarly B is true at some world and hence $\Diamond B$ is true at all fragments as well, and hence all fragments in all worlds are such that both $\Diamond A$ and $\Diamond B$ are true at them, which suffices for $\Diamond A \circ \Diamond B$ to be true in any world. We thus have a restricted form of the unity assumption, holding only for modal facts: there is no fragmentation in what is metaphysically possible or not.

Much more could be said about modality in a fragmentalist setting, and matters will become more complicated (and more interesting) when we consider different kinds of modal notions. I hope the above suffices to clarify one simple and straightforward way of incorporating metaphysical possibility. In short: modal space now includes fragmented possible worlds in addition to the fully unified worlds, representing possible ways the world could be

fragmented. What is possible (and necessary) is assumed to be an invariant matter, all fragments look across the same modal space in the very same way.

1.5 Appendix: incorporating quantification

When we consider the way co-obtainment interacts with quantification, this time the choice arises whether existence can be a perspectively variant matter or not, that is, whether something may exist *insofar as* certain things are the case, yet fail to exist *insofar as* certain other things are the case. This time I think we should decide in favour of variance and allow fragments to disagree about what exists. After all, on what grounds would we assume an asymmetry between existing and having a property? Moreover, we will have some occasion to resort to the perspectival variance of existence in applications of fragmentalism. As things are complicated enough, I will only consider the first-order extension of the basic semantics instead of the modal extension discussed in the previous appendix.

The vocabulary of the language now comprises the following elements:

variables: v_0, v_1, v_2, \dots

constants: k_0, k_1, k_2, \dots

for every natural number $n \geq 0$, n -place predicates $P_n^0, P_n^1, P_n^2, \dots$

connectives: $\circ, \wedge, \vee, \neg$

quantifiers: \forall, \exists

brackets: $(,)$

I will use x, y, z for arbitrary variables, and a, b, c for arbitrary constants. I will use F_n, G_n, H_n for arbitrary n -place predicates. I will omit the subscript in cases where it is clear from context. Any constant or variable is called a *term*.

The syntax of the language is as one would expect:

If t_1, \dots, t_n are any terms and F is any n -place predicate, $Ft_1\dots t_n$ is an atomic formula.

If A and B are formulas, so are the following: $(A \circ B)$, $(A \wedge B)$, $(A \vee B)$, $(\neg A)$.

If A is any formula, and x is a variable, then $\forall xA$, $\exists xA$ are formulas.

Nothing else is a formula.

I will omit outer brackets in formulas from now on. Any occurrence of a variable in an atomic formula is said to be *free*. The free occurrences of variables in $\neg A$, $A \wedge B$, $A \vee B$, $A \circ B$ are just the free occurrences of variables in A and B . And, the free occurrences of variables in a quantified formula, $\forall xA$ or $\exists xA$ are the variables other than x that are free in A . If an occurrence of a variable is not free in a formula A , then it is said to be *bound*. A formula with no free variables is said to be *closed*. $A_x(a)$ is the formula obtained by substituting a for each free occurrence of x in A .

A model M for the language is a triple $\langle D, W, v \rangle$, where:

D is a domain of quantification,

W is a set of points,

v assigns:

- (1) to each constant a of the language a member $v(a) = d$ in D ,
- (2) to each point w in W a subset of D , so $v(w) \subseteq D$,
- (3) to each n -place predicate F , relative to each point w in W , a subset of $v(w)^n$.

Intuitively, $v(a)$ is the object named by a , $v_w(F)$ the set of n -tuples that satisfy F at point w (and which may vary from point to point), and $v(w)$ is the domain at point w . I will write $v(w)$ as D_w , the domain of point w .

Given a model, truth values are assigned relative to all closed formulas, which I will call sentences. Now it may be that some object d in the domain

is not the denotation of any constant in the language even though it's in the extension of predicates in that language. To nevertheless capture the right truth conditions for the quantifiers, we will specify the truth conditions in terms of possible extensions of the language. If M is some model for a language, then we can consider extensions of that model in which a name c , which is not already in the language, is added to the language and given a denotation by v . I will write M_d^c for an extended model in which v assigns c the denotation d , and which is otherwise exactly like M . I will say that an object d satisfies a formula A_x in a model M , written $M \Vdash A_x[d]$, to mean that, if we were to add a name and assign it d as denotation, this name would satisfy A_x . That is:

$$M \Vdash A_x[d] \text{ iff } M_d^c \Vdash A_x(c)$$

Similarly, I will write $v_w(A_x[d]) = 1$ in M to mean that $v_w(A_x(c)) = 1$ in M_d^c . That is to say, the valuation v determines that an object d satisfies a formula A_x in a model M if and only if an extension of M in which v assigns d to c is such that $v_w(A_x(c)) = 1$.

Given these stipulations, v is recursively extended to assign 1 or 0 to each formula relative to each point w in W :

$$v_w(Fa_1\dots a_n) = 1 \text{ iff } \langle v(a_1), \dots, v(a_n) \rangle \in v_w(F)$$

$$v_w(A \circ B) = 1 \text{ iff } v_w(A) = 1 \text{ and } v_w(B) = 1$$

$$v_w(A \wedge B) = 1 \text{ iff } v_w(A) = 1 \text{ and } v_w(B) = 1$$

$$v_w(A \vee B) = 1 \text{ iff } v_w(A) = 1 \text{ or } v_w(B) = 1$$

$$v_w(\neg A) = 1 \text{ iff } v_w(A) \neq 1$$

$$v_w(\exists x A) = 1 \text{ iff for some } d \in D_w, v_w(A_x[d]) = 1$$

$$v_w(\forall x A) = 1 \text{ iff for all } d \in D_w, v_w(A_x[d]) = 1$$

I will call these the v -clauses for the connectives.

Truth in a model, written $M \Vdash A$, is defined via the following recursive clauses

$M \Vdash Fa_1 \dots a_n$ iff $\exists w(v_w(Fa_1 \dots a_n) = 1)$

$M \Vdash A \circ B$ iff $\exists w(v_w(A \circ B) = 1)$

$M \Vdash \neg A$ iff $M \not\Vdash A$

$M \Vdash A \wedge B$ iff $M \Vdash A$ and $M \Vdash B$

$M \Vdash A \vee B$ iff $M \Vdash A$ or $M \Vdash B$

$M \Vdash \exists x A$ iff for some $d \in D$, $M \Vdash A_x[d]$

$M \Vdash \forall x A$ iff for all $d \in D$, $M \Vdash A_x[d]$

Finally, we understand validity and logical truth in the usual way (where Σ is a set of sentences):

An argument from Σ to A is *valid*, written $\Sigma \vDash A$, iff, for every model M , if $M \Vdash \Sigma$ then $M \Vdash A$.

A formula A is *logically true*, written $\vDash A$, iff, for every model M , $M \Vdash A$.

Thus far the semantics.

To see the semantic machinery at work, consider how the model determines various kinds of formulas of a language to be true in the model:

Fa . This is an atomic sentence, and hence true if there is some point w at which it is true, i.e. if there is a point where a is in the extension of F .¹⁴

$\neg Fa$. This is true if Fa is false. Fa is false if there is *no* point at which Fa is true, which is the case when a is not in the extension of F anywhere.

$Fa \circ Ga$. This is true if there is a point at which both Fa and Ga are true, i.e. if there is a single point where a is in the extension of both F and G .

¹⁴When I say that a is in the extension of F , I really mean that the object d , denoted by a , is in the extension of F .

$Fa \wedge Ga$. This is true if both Fa and Ga are true. Fa and Ga are both true if there is a point for each at which it is true (and this does not have to be the same point).

$Fa \circ \neg Ga$. This is true if there is a point at which both Fa and $\neg Ga$ are true. And $\neg Ga$ is true at a point if Ga is false *at that point*.

$\exists xFx$. This is true if there is some object d such that $F[d]$ is true. The truth of the latter is determined in virtue of the corresponding atomic formula in extended models: it is true if in any extension that assigns d the name c , Fc is true, and since this is an atomic sentence, this is the case if there is a point in the extended model at which Fc is true.

$\neg \exists x \neg Fx$. This is true if $\exists x \neg Fx$ is false. $\exists x \neg Fx$ is false if there is no object d which is such that $\neg F[d]$ is true. This in turn means that, for every object $d \in D$, $F[d]$ is true, i.e. every object is in the extension of F at some point. (Of course, this means that $\forall xFx$ is true. Indeed it can be checked that we can, more generally, define $\forall xA$ as $\neg \exists x \neg A$).

$A \circ \forall xFx$. This is true if there is a point w where A is true and where $\forall xFx$ is true (i.e. if $\exists w(v_w(A) = 1$ and $v_w(\forall xFx) = 1$). Now $v_w(\forall xFx) = 1$ if and only if, for all $d \in D_w$, d is in the extension of F relative to w . (Note that we are only concerned with D_w here and not D).

The last of these implies the failure of an important inference. Everything may be such that B *insofar as* certain things are the case, and yet not everything may be such that B , that is:

$$A \circ \forall xB \not\equiv \forall xB$$

For example, consider the special case of $A \circ \forall xFx$. The fact that v assigns a subset of the domain to each point reflects that everything can be F *insofar as* certain facts obtain, even though not everything is F . This reflects that the satisfaction of predicates by objects is a perspectival matter.

In particular, let $=$ be any predicate whose extension relative to a point in a model is always and only pairs with the same object as first and second member (i.e. whose extension are pairs such as $\langle d, d \rangle$). Then a model M can be such that $M \Vdash \forall x(x = a \vee x = b) \circ A$ even though $M \not\Vdash \forall x(x = a \vee x = b)$. That is, a and b may be all there is *insofar as* A and yet not be all there is, given that c exists *insofar as* B . This reflects again the decision to treat existence as a possibly perspectival matter.

Interestingly, analogous inferences *do* go through for existence. Note first of all that we have:

$$Fa \circ A \vDash Fa$$

If there is a point where A holds and where a is in the extension of F , this suffices for Fa to be true in the model. And the same holds for the existential generalization of the atomic sentence:

$$\exists xFx \circ A \vDash \exists x(Fx)$$

If $\exists xFx \circ A$ is true in a model, then there is a point where both $\exists xFx$ and A hold. The former holds only when there is some object in the extension of F at that point, and this suffices for $\exists xFx$ to be true in the model.

The same reasoning applies in particular for existence claims. If there is a model M such that $M \Vdash \exists x(x = a) \circ A$, then $M \Vdash \exists x(x = a)$. Just as Fa is true if there is a point where the object denoted by a is in the extension of F , so $b = a$ is true if there is a point where the objects denoted by a and b are in the extension of $=$ (which only picked out predicates in models that behave like numerical identity as we know it). If an object exists *insofar as* A , then the object exists.

With regards to the other direction, it's not the case that if $\forall xA$ or $\exists x(A)$ is true simpliciter, it holds at all fragments:

$$\forall xA, B \circ C \not\equiv \forall xA \circ B$$

$$\exists xA, B \circ C \not\equiv \exists xA \circ B$$

Imagine we have a model where every object in the domain is in the extension of a predicate F at some point. Then $\forall xFx$ is true in the model. But that every object in the domain is in the extension of F at some point allows that there are points where not all objects are in the extension of F . That every object is such that it's F somewhere does not mean that everywhere all objects are F , in particular, not everything may be F where B holds. In the case of existence, it may be that some object d is in the extension of F at some point, but that does not guarantee that it's in the extension of F at the very point where B holds.

There is much more to be said about quantification in fragmented worlds; but I hope that this minimal treatment suffices to illustrate how we can incorporate quantification in the fragmentalist framework and take the existence of an object to be no less perspectival than its instantiation of properties. In short: the boundaries of what exists may include such and such *insofar as* these facts obtain, but include so and so *insofar as* those other facts obtain.

Chapter 2

Fine on Fragmentalism

Fine introduces fragmentalism in his ‘Tense and Reality’ (2005). In this chapter, I will discuss fragmentalism as it is characterized by Fine.

The issues that I will raise are more or less a record of what kept me from adopting the fragmentalist position as Fine characterized it, and which informed the development of the fragmentalist framework discussed in *Chapter 1*. The critical discussion raises open questions for the conceptual machinery of Fine’s fragmentalism, and discusses one of Fine’s arguments in favour of fragmentalism: the argument from truth. The final section offers a formulation of tense realism within the framework of *Chapter 1*, and shows how it solves some of the issues raised in the critical discussion. It does not solve all issues, however. I think the problems that emerge from the discussion of Fine’s argument from truth remain, and motivate a view of time that is not based in tense. Such a view is explored in *Chapter 3*.

2.1 Fine’s conception of fragmentalism

Fine’s defence of fragmentalism is tentative; he makes it clear from the start that he is out to defend two conditional claims, namely that ‘if one is going to be a realist about tense, then one should be a non-standard realist and that, if one is going to be a non-standard realist, then one should be a fragmentalist’ (Fine 2005b:262). He supports these conditional claims by

arguing that, compared to standard realism about tense, fragmentalism is in a better position to account for the passage of time, that it can make sense of the truth of tensed utterances, and that it's compatible with the special theory of relativity. He also argues that it's conceptually more stable than another non-standard view of tensed facts (called external relativism) and he shows how fragmentalism might help us to better understand subjectivity.

Fine summarizes his conception of fragmentalism as follows:

Under such a view, reality will be fragmentary. Certain of the facts constituting reality will 'cohere' and some will not. Any fact is plausibly taken to belong to a 'fragment' or maximally coherent collection of facts; and so reality will divide up into a number of different but possibly overlapping fragments. (Fine 2005b: 281).

To properly understand this description, various bits of terminology need to be explained.

The above description features the notion of facts constituting or belonging to reality. This reflects an important methodological starting point of Fine's approach, namely that we sharply distinguish between stating what is the case and stating what is the case *in reality* (see Fine 2005b: §2; and Fine 2001: §8-10). There are at least two ways in which one can introduce the distinction. One may draw the distinction in terms of differences in the assertability conditions of certain claims: when engaged in metaphysics, we are not typically willing to assert everything we are willing to assert in ordinary contexts. A nominalist will in ordinary contexts assert that there are numbers between 1 and 5, but not when engaged in metaphysics. There is all the things the metaphysician ordinarily says, and then there are the things she is willing to say *and* takes to describe how things are *fundamentally speaking*, or *in reality*. The standards on what we are willing to assert are higher when engaged in metaphysics.

One may also introduce the distinction in terms of a much older distinction between reality and appearance. Fine:

The idea of a realist or anti-realist metaphysics involves a fundamental distinction between what is real and what is mere ap-

pearance. Intuitively, there is nothing more to the metaphysician's world than what he takes to be real. So, for example, on Democritus's view there will be nothing more to the world than atoms in the void. But there might appear to be something more to the world than what he takes to be real – there might appear to be minds or morality, say, in addition to the atoms – and hence the distinction between appearance and reality. Once we accept this distinction between appearance and reality, the metaphysician must explain why what appears to be something more is not genuinely something more. (Fine 2007: 23)

We come to the world employing all sorts of representations of it; but we arguably also come with a concept of reality, or of the world as it is in itself, and we do not typically take all aspects of all our representations to reflect what we take the world to be like in it itself.

Fine's conception of reality is tied in with his view on metaphysical grounding.¹ On this view, some facts are taken to obtain *in virtue of* or *because* certain other facts obtain.² To take a stock example, the fact that $A \wedge B$ obtains *because* the fact that A obtains and the fact that B obtains. The conjunctive fact obtains because the relevant conjuncts obtain. Similarly, if the fact that $A \vee B$ obtains and the fact that A obtains, then the fact that $A \vee B$ obtains *in virtue of* the fact that A . Such grounding relations are, in turn, assumed to reflect the structure of reality: if some fact is grounded in other facts, that is, if the obtaining of the former is *metaphysically explained* in terms of the obtaining of others, this gives us reason to think that it is thereby not one of the facts that obtain *in reality*.

In order to make the distinction more explicit between *what is the case* and *what is the case in reality*, Fine proposes that we introduce a sentential operator ' \mathfrak{R} ' ('it is the case in reality that...') to form sentences of the form ' $\mathfrak{R}A$ ' which say that *it is the case in reality that A* . For ease of expression, Fine

¹See Fine (2001) and Fine (2012).

²Fine takes an operator approach to grounding, where grounding claims are made in terms of sentences of the form ' A because B '. This stands in contrast to a view on which grounding is understood as a relation amongst entities (objects, events and states of affairs), as in Schaffer (2009).

allows the more loose talk of facts, which belong to, compose or constitute reality. This is a double reification of facts and of reality, turning them into things we refer to (Fine 2005b: 268) instead of things we state to be the case. It's important to keep in mind that this just loose talk, and that on the official view, we do not refer to facts and to reality as entities. They are not included in the official ontology of the view.

Fine discusses how various views of time employ different concepts of reality. He formulates four general metaphysical principles whose conflict lie at the heart of McTaggart's argument for the unreality of time, and the resolution of which results in the different theories of time. Fine's exact formulations of these principles are as follows (see 2005: 271):³

Realism: Reality is composed of tensed facts.

Neutrality: No time is privileged, the tensed facts that constitute reality are not oriented towards one time as opposed to another.

Absolutism: The constitution of reality is an absolute matter, i.e. not relative to a time or other form of temporal standpoint.

Coherence: Reality is not contradictory, it is not constituted by facts with incompatible content.

The above principles are formulated using the loose talk of reality composed of facts. They can be translated back into principles featuring the reality statements, in the following way:

Realism: there are true tensed sentences A such that we take it that $\mathfrak{R}A$. For example: it is the case that $\mathfrak{R}(\text{KF is sitting})$.

Neutrality: where A and B truthfully describe how reality is at different times, we take it that $\mathfrak{R}A$ and $\mathfrak{R}B$. For example: we take it that $\mathfrak{R}(\text{Aristotle is sitting})$ and $\mathfrak{R}(\text{KF is sitting})$.

³Fine also discusses a more sophisticated version of McTaggart's argument (2005: §4), with slightly more complicated principles. The simpler argument, and the simpler formulation of the relevant principles suffice for our purposes.

Absolutism: it's not the case relative to times that $\mathfrak{R}A$, it's the case simpliciter.

Coherence: where A and B truthfully describe incompatible states, it's not the case that $\mathfrak{R}A$ and $\mathfrak{R}B$. For example: it's not the case that $\mathfrak{R}(\text{KF is sitting})$ and $\mathfrak{R}(\text{KF is standing})$.

These principles conflict as follows. By Realism there will be some tensed fact that is the case in reality. Let it be the fact that KF is sitting. This fact constitutes the way that reality is at some moment in time. By Neutrality, there will equally be facts that constitute the way that reality is at different moments in time. As it is plausible to think that Reality varies across time, there will be such facts that are incompatible with the first, such as the fact that KF is standing. By Absolutism, we deny that the fact that KF is sitting belongs to reality *relative to one time* and that the fact that KF is standing belongs to reality *relative to a different time*; the facts simply belong to reality. As they are incompatible facts, this contradicts Coherence.

Given the conflict, at least one of the principles needs to be rejected. Different views reject different principles. Fine recognizes four views:

Anti-realism: there are no true tensed sentences A such that $\mathfrak{R}A$.

Standard realism: if $\mathfrak{R}A$ and $\mathfrak{R}B$, then A and B describe how reality is at a single time, namely the current time.

External relativism: what is the case in reality is relative to times, e.g. it will be the case that $\mathfrak{R}(\text{KF is sitting})$ relative to one time t and $\mathfrak{R}(\text{KF is standing})$ relative to a different time t^* . There are thus multiple realities, each relative to a different moment in time.⁴

⁴This fits Fine's remarks that, on external relativism, the constitution of reality by facts is relative to times (2005: 280). In a later discussion, Fine discusses how fragmentalism is superior to external relativism because it allows us to deny that times are 'fundamental constituents of reality' (2005: 308). This might be taken to suggest that times feature inside the scope of reality statements, and not outside of it as I formulated external relativism here. But Fine remarks: 'surely any reasons for thinking that times are not basic should apply across the board, not only to their role in the specification of the facts that are real but also in the formulation of the reality claims themselves' (2005: 310). Reasons to deny that times are fundamental constituents are equally reasons to deny that reality statements are relativized to them, as on external relativism.

Fragmentalism: incompatible facts are the case in reality, e.g. we allow that $\mathfrak{R}(\text{KF is sitting})$ and $\mathfrak{R}(\text{KF is standing})$.

As a result of rejecting different principles, each view has a different concept of reality. Fragmentalism emerges here as one way of replying to McTaggart's argument, accepting its conclusion that reality is in some sense incoherent.⁵ On the fragmentalist view of time, reality is constituted by all the tensed facts that obtain at some time or other, even though some of these will be incompatible such as the fact that KF is sitting and the fact that KF is standing.

Fragmentalism doesn't simply reject the *Coherence* principle. Many of the facts that belong to reality are taken to stand in a relation that Fine calls 'coherence'. He proposes that we take this relation of coherence between facts as a primitive notion (2005: 281). Using this relation, we can then identify maximal coherent collections of facts, the 'fragments' of reality. If any two facts cohere, then they are compatible. So all the facts that belong to the same fragment are at least compatible facts.

Each moment of time is identified with a fragment. This means that the fact that KF is sitting coheres also at least with all the facts that obtain simultaneously with it, so that there is a maximal coherent collection of facts (i.e. a fragment) whose contents corresponds exactly to the contents of a moment in time. Times are therefore not *sui generis* entities over and above that which is the case when it is that time.

The division of the world into internally coherent fragments plays an important role in the semantics that accompanies the metaphysical picture, see (2005: 282, 295-298). The semantic machinery is devised in such a way that the metaphysical picture should not be taken to legitimate contradictory utterances, even though it admits contradictory facts to reality. The starting assumption is that an utterance comes with an in-built specification of the

⁵A note on terminology. Where Fine uses 'fragmentalism' for the specific fragmentalist view of tensed facts, I will also use 'fragmentalism' for the general view that the world is fragmented in some way, just as we talk of relativism about this or that and relativism in general. Sometimes, when I refer to fragmentalism, I trust that it will be clear from context whether I refer to a fragmentalist view of a particular phenomenon, such as time, or to fragmentalism as understood as a general metaphysical picture.

fragment within which the verifying facts are to be found, a ‘focal content’, and that the utterance will be false when the required facts are not found within the specified fragment, even though they are found in a different fragment. Fine understands this in-built specification of a time, the focal content, to be a new type of meaning. An utterance is naturally taken to be focussed on the time of utterance, and this is captured in a semantic principle that Fine lays down for both external relativism and fragmentalism (2005: 295):

Relative Link: an utterance is true if and only if what it states is verified by the facts that obtain at the time of utterance.⁶

Because all the facts that belong to a single fragment are compatible with each other, and because the utterance is focused on a single fragment only, and not on the fragmented reality in its totality, no utterance of a contradiction comes out as true.

In a nutshell then, Fine’s fragmentalism pictures reality at large as a series of successive collections of tensed facts that are internally coherent but not coherent with each other. This is not taken to vindicate contradictory utterances however because it’s assumed that our situated language-use only makes claims about the collection of tensed facts that it’s focused on and not about reality at large.

2.2 Critical discussion I: on tense and truth

As mentioned, Fine offers a range of arguments for favouring fragmentalism over standard realism. One of these arguments is that standard realism cannot maintain a very plausible understanding of the relation between the truth of utterances and reality as it is conceived on such a view. Fine offers in fact two closely related arguments. I will only discuss the second of these.

⁶The formulation is Fine’s. In the case of fragmentalism, the principle is not to imply that facts ‘obtain at a time’ (contrary to Absolutism). The fragmentalist will hold that an utterance is true if and only if what it states is verified by facts that cohere with the utterance’s being made.

The fragmentalist can avoid the first argument by insisting on the *relative link* principle mentioned in the previous section. The fragmentalist's response to the first argument seem right to me. It's only the response to the second argument that I find unsatisfactory, and so this is what I will focus on.

The main intuitions behind the argument are simple.⁷ Say that I utter 'it rains' when it does indeed rain. The utterance is clearly true. One might also think that the token utterance doesn't change what it says when the sun starts shining and, given tense realism, one will surely think that it has tensed truth-conditions, claiming that it rains. One might also think that my token utterance back then does not change to being false when the sun starts shining a moment later. I will have no inclination to say that the utterance made back then is now false. The truth of that very token utterance is only affected by the way things were back then, and not by the way they are now. So the token utterance continues to have the truth-value it has. This means that we want to say all of the following things: that the token utterance has tensed truth conditions, that it was and remains true, and that it said and continues to say that it rains. The tension in this is not hard to spot: if my token utterance of 'it rains', which I made yesterday, is still true today, and still says that it rains, then this seems to imply that it rains today - but this may not be the case, it may have rained yesterday but not today.

Fine makes explicit a host of assumptions that drive the argument. He formulates these using (1) the notion of *truth* as applied to token utterances, (2) a relation of *stating*, which holds between a token utterances and a proposition (a token utterance states a proposition), and (3) a relation of *verification*, which holds between a stated proposition and a fact (a proposition is verified by a fact) (2005: 288-289). Also, still following Fine's methodology, we distinguish sharply between the obtaining of a fact and its belonging to reality: for the fact that P to obtain is for it to be the case that P , for the same fact to belong to reality is for it to be the case that $\mathfrak{R}P$. The assumptions that drive Fine's take on the argument are the following:

*Link**: an utterance is true only if the proposition stated by the utter-

⁷The argument goes back to Evans (1985) and Mellor (1986). Fine draws upon a discussion of the argument by Percival (2002).

ance is verified by facts.⁸

Truth-value stability: if an utterance is true, then it is always true.

Content stability: if an utterance states the proposition that P , then it always states the proposition that P .

Factuality: if some facts in reality verify the proposition that P then those facts obtain.

Conditionality: if some facts in reality verify the proposition that P and those facts obtain, then P .

Now say that I make the utterance ‘KF is sitting’ and say that this token utterance is true (because I make the utterance when KF is indeed sitting). The tense realist takes this utterance to state a tensed proposition, in this case the proposition that KF is sitting. By Content stability, the utterance always states the proposition that KF is sitting. As we assume that the utterance is true, then by Truth-value stability, the utterance is always true. By Link*, this must mean that it is always the case that some facts that belong to reality – f_1, f_2, \dots – verify the proposition that KF is sitting. By Factuality, this means that those facts, f_1, f_2, \dots always obtain. And given that these facts both verify the proposition that KF is sitting and always obtain, it follows from Conditionality that KF is always sitting – which is false; it’s not always the case that KF is sitting.

Fine argues that the fragmentalist – contrary to the standard realist – has room to block the argument in a satisfactory manner. Fine takes it that the Factuality assumption can be broken down into two more specific theses:⁹

⁸This is the left to right direction of the stronger Link principle: an utterance is true only if and only if what it states is verified by the facts. This principle drives the first version of the argument but is not accepted by fragmentalism, which only accepts the ‘Relative Link’ principle mentioned above.

⁹Whereas on page 292, Fine claims that these two principles - Reality 1 and Reality 2 - support Factuality, on page 297 Fine claims that they support Conditionality instead of Factuality. Reality 1 and Reality 2 straightforwardly imply Factuality however, so I take his latter claim to be a slip.

Reality 1: if some facts verify a proposition then these facts belong to reality.

Reality 2: if a fact belongs to reality, then it obtains.

In order to rebut the argument, Fine argues that the fragmentalist can reject *Reality 2* and thus deny that if a fact belongs to reality, it obtains. To make fully explicit how this blocks the argument, we may replace factuality with *Reality 1* and *Reality 2*:

1. I make the utterance u , which states the proposition that KF is sitting, and is true when uttered.
2. *Content stability*: if an utterance states the proposition that P , it always states the proposition that P .
3. u always states the proposition that KF is sitting.
4. *Truth stability*: if an utterance is true, it is always true.
5. u is always true.
6. *Link**: an utterance is true only if the proposition stated by the utterance is verified by facts $f_1, f_2 \dots$
7. the proposition that KF is sitting is always verified by facts f_1, f_2, \dots
8. *Reality 1*: if a proposition is verified facts f_1, f_2, \dots , then those facts belong to reality.
9. f_1, f_2, \dots always belong to reality.
10. *Reality 2*: if a fact belongs to reality, then it obtains.
11. f_1, f_2, \dots always obtain.
12. *Conditionality*: if some facts in reality verify the proposition that P and those facts obtain, then P .

13. it's always the case that KF is sitting

Fine proposes that the fragmentalist rejects *Reality 2* and that we accept that the facts that verify the tensed proposition that KF is sitting *belong to reality*, yet deny that they obtain. Thereby the antecedent of *Conditionality* is not satisfied, and hence we are not forced to say that KF is always sitting.

The rejection of *Reality 2* is the rejection of factivity for reality statements. Without the reification of facts, it reads:

Reality 2: if $\mathfrak{R}A$, then A .

To motivate the rejection of factivity, Fine writes:

We might naturally take a fact to belong to reality if it belongs to *a* reality. It is then the second Reality assumption, that any fact belonging to reality obtains, which should be given up; for the fact may relate to one reality and the obtaining to another. In stating that a fact belongs to reality, we adopt a general perspective but, in stating that a fact obtains, we adopt the current perspective; and it is because of this shift in perspective that we cannot generally assert that the facts belonging to reality will obtain. (Fine 2005b: 297-298).

Take the claim that ML is currently dead. As fragmentalist, I might be committed to the claim that $\mathfrak{R}(\text{ML is dead})$ if there is a time where ML is dead. If the claim that $\mathfrak{R}(\text{ML is dead})$ implies that ML is dead, then I'm committed to the claim that ML is dead, which seems mistaken if we read as a present tensed claim about how things are right now. The present tensed character of the fact forces me to adopt the point of view of the current time when stating that the fact obtains; but given that all times are on par, this introduces the tensed contents of past and future times into the present. And this results in many falsehoods (and, indeed, contradictions), for example, I'm not currently dead.

It seems to me that Fine's solution for the fragmentalist is problematic. First of all the rejection of factivity conflicts with a natural concept of reality,

of what it is for something to be the case in reality. It's natural to think that there is everything that is the case and that there is a privileged subset of these matters that are the case *and* reflect what reality is like in itself. When Fine first glosses the notion of reality he remarks that 'whatever is really the case (belongs to metaphysical reality) may, with some plausibility, be taken to be the case (belong to mere reality)' (2005: 267; noted in Merlo 2013: 5). The denial of factivity goes against this initial gloss of the notion. To be sure, when providing the initial gloss, Fine is only introducing a generic notion of reality, which becomes refined in different ways by the different views that he discusses. The initial gloss is by no means presented as a core feature shared by all the different views. But it highlights that factivity is part of the intuitive understanding of reality, and belongs to a core feature if anything does. If we use the fact talk and deny the factivity of reality, then we have to say that some facts belong to reality without being the case at all, or alternatively, that some of the things that belong to reality are not facts. This, it seems to me, is not a plausible notion of reality, and it's quite a cost of fragmentalism if it requires this.

Relatedly, it's indeed worrying if the fragmentalist blandly states that I'm dead and reads this as having a present tensed character but, I'm not sure whether it's any less worrying if the fragmentalist states that it's the case in reality that ML is dead (i.e. that $\mathfrak{R}(\text{ML is dead})$). That is to say, it's not clear to me that when we state that a tensed fact belongs to reality, we 'adopt a general perspective' instead of a perspective on what things are currently like, not if the fact that is said to belong to reality is a tensed fact which concerns the way things are now. The claim that it is now the case that I'm dead *in reality* does not sound much better than the claim that it's now the case that I'm dead. It seems we need a stronger motivation to deny the factivity of reality.

There is another worry concerning Fine's solution. I think that, even with the denial of the factivity of reality (*Reality 2*), there are other plausible principles concerning the connection between reality and truth, which still allow the problematic derivation to go through. Consider again *Conditionality*:

Conditionality: if some facts verify the proposition that P and those facts obtain, then P .

This principle mediates the connection between reality and truth via the mere obtaining of the facts. It is this that allows the fragmentalist to block the problematic argument, given that she denies that the facts that belong to reality obtain. But it's unclear why the relation between reality and truth needs to be mediated through the obtaining of the facts that belong to reality. The following is a more direct principle:

*Conditionality**: if some facts verify the proposition that P and *those facts belong to reality*, then P .

If some facts verify a proposition, then their *belonging to reality* suffices for the proposition to be verified. This seems a natural thought concerning the relation between reality and truth: if some facts are such that they verify a proposition if they obtain, then surely they also verify a proposition if they obtain *in reality*. But in that case the argument goes through again:

1. I make the utterance u , which states the proposition that KF is sitting, and is true when uttered.
2. *Content stability*: if an utterance states the proposition that P , it always states the proposition that P .
3. u always states the proposition that KF is sitting.
4. *Truth stability*: if an utterance is true, it is always true.
5. u is always true.
6. *Link**: an utterance is true only if the proposition stated by the utterance is verified by facts $f_1, f_2 \dots$
7. the proposition that KF is sitting is always verified by facts f_1, f_2, \dots

8. *Reality 1*: if a proposition is verified facts f_1, f_2, \dots , then those facts belong to reality.
9. f_1, f_2, \dots always belong to reality.
10. *Conditionality**: if some facts in reality verify the proposition that P and those facts belong to reality, then P .
11. it's always the case that KF is sitting

Notice that *Reality 2* doesn't feature in this line of reasoning, and hence it's beside the point that the fragmentalist can reject it. So, to avoid the argument, the fragmentalist must now either reject *Conditionality** or one of the other principles as well.

But there is an even simpler argument. This time starting from the other end, the truth of the utterance. If an utterance is currently true, and it currently states the proposition that P , then, plausibly, we can infer from this that it is the case that P . That is:

*Conditionality***: if an utterance currently states the proposition that P and the utterance is currently true, then P .

This is a natural disquotational principle for utterances: if I tell you that an utterance of mine is currently true, and that it currently asserts that it is raining, then surely this should be enough information for you to conclude that it is raining. This principle makes the argument much quicker, and thus harder to avoid:

1. I make the utterance u , which states the proposition that KF is sitting, and is true when uttered.
2. *Content stability*: if an utterance states the proposition that P , it always states the proposition that P .
3. u always states the proposition that KF is sitting.
4. *Truth stability*: if an utterance is true, it is always true.

5. u is always true.
6. *Conditionality***: if an utterance currently states the proposition that P and the utterance is currently true, then P .
7. KF is sitting.

The rejection of *Conditionality*** seems awkward: it means that I should make room for utterances which are now true, and which now state that matters are the case, without those things being indeed the case. This strongly suggests that something is wrong with *Content stability* or *Truth stability*.

To sum up, the fragmentalist doesn't only require costly revisions in the concept of reality, even with the non-factive conception of reality, she also needs to reject further plausible principles concerning the relation between the truth of utterances and the obtaining of the relevant proposition stated by the relevant utterance.

2.3 Critical discussion II: on (in)coherence

Fine's paper is very rich, covering many different topics and issues. It's not a paper solely about fragmentalism. As a result Fine's characterization of the view is brief and leaves open certain important questions concerning the conceptual foundations of the view. The critical discussion revolves around two questions: first the question of the exact sense in which reality is incoherent according to Fine's fragmentalism, and second the question of exactly how we are to understand the coherence relation that structures the fragments.

Fine remarks that 'the fragmentalist takes reality to be contradictory' and that she denies that 'all contradictions can be ironed out' (2005: 282, 280). This sounds worrying of course, and will be enough to deter many from the view. But it's not clear exactly in what sense reality is incoherent according to Fine's fragmentalism.

We typically think of contradictions as conjunctions of contradictory conjuncts. So one might wonder first of all whether conjunctions of contradictory

conjuncts are ever the case in reality, that is, whether the fragmentalist takes it that $\mathfrak{R}(A \wedge \neg A)$. Fine will reject this, given that he thinks we should ‘explain the obtaining of a conjunctive fact in terms of the obtaining of its conjuncts’ (2005: 281). The relevant sense of explaining here is presumably that of grounding (as understood in his 2001 and 2012): the obtaining of a conjunction is grounded in the obtaining of each of the conjuncts, and hence conjunctions are not the case in reality. If conjunctions are not the case in reality than, ipso facto, neither are conjunctions of contradictory conjuncts the case in reality. So reality cannot be contradictory in this sense, according to Fine.

But then, in what sense is it incoherent exactly? There is still a range of open options that differ on what is allowed to contradict with what. Consider the following two distinct ways in we may allow reality to be contradictory:

- I. Allow that $\mathfrak{R}(A)$ and $\neg\mathfrak{R}(A)$, i.e. allow that it both is and is not the case in reality that A .
- II. Allow that $\mathfrak{R}(A)$ and $\mathfrak{R}(\neg A)$, i.e. allow that it is the case in reality that A and that it is the case in reality that not A .

These different options come with different commitments on the part of fragmentalism. Let me discuss them separately.

According to option I, reality is contradictory in the sense that $\mathfrak{R}(A)$ and $\neg\mathfrak{R}(A)$. In other words, we say that the statement ‘ $\mathfrak{R}(A)$ ’ is both true and not true, or that the fact that A both belongs and doesn’t belong to reality. If the fragmentalist allows such contradictions, the metaphysical view could not plausibly assume classical logic is the right logic, given that classical logic is explosive: a contradiction implies everything. So either fragmentalism must commit itself to a rejection of classical logic, or this is not the way reality is incoherent according to fragmentalism.

According to option II, reality is contradictory in the sense that $\mathfrak{R}(A)$ and $\mathfrak{R}(\neg A)$. In other words, the fact that A and the fact that not A are both the case in reality. Note first of all that, if we allow that $\mathfrak{R}(A)$ and that $\mathfrak{R}(\neg A)$ and we want to avoid contradictions, we must again deny the

factivity of reality, that is, we must deny that if $\mathfrak{R}A$, then A . As Merlo notes in a discussion of Fine's view, if the notion of reality is factive, then 'incoherence is fated to spread from reality (the totality of what is really the case) to the world (the totality of what is the case)' (2013: 6). This denial of factivity of reality seems implausible to me, as we already discussed in the previous section.

Also, if this conception is not to imply option I, we must furthermore reject that if $\mathfrak{R}\neg A$, then $\neg\mathfrak{R}A$. So we must assume that, if it is the case in reality that *not A*, this does not mean that it's not the case in reality that A . The reality of a negated fact does not suffice for that fact to be absent from reality. This also seems in tension with a plausible understanding of reality. Indeed, when Fine introduces the loose talk of facts and reality, the idea is precisely that we tend to think of reality as a kind of 'container' for the facts. To say with regards to the container, *not A*, this is straightforwardly read as saying precisely that A is absent from the container. If we accept that the fact that *not A* can belong to reality, even when the fact that A does as well, this seems to require a non-standard understanding of negation and undermines our thinking of reality as a kind of realm or container.

Fine states explicitly that fragmentalism takes reality to be 'irredeemably incoherent' (see 2005: 281) and unsurprisingly, fragmentalism has been resisted in the literature on the grounds that it offers an incoherent conception of reality (see, e.g., Deng 2013: 10-11). But before we can critically assess fragmentalism on this basis, we need to be clear on the precise sense in which it is incoherent. It appears that the less problematic the contradictions that are accepted, the more problematic the conception of reality becomes.

The question of the exact sense in which reality is incoherent is connected with the question of how we are to understand the coherence relation that structures the facts into fragments. The notion of 'coherence' plays a central role. There are however some open questions concerning the notion.

Note first of all that—though the relation is called 'coherence' by Fine—it must be quite different from the ordinary relation of coherence that holds amongst propositions. This is shown clearly by Correia and Rosenkranz (2012: 312). Dubbing Fine's notion of coherence 'coherence*', to distinguish

it from the ordinary understanding, they write the following:

[S]upposing, say, that there is no time, past, present or future, at which Socrates is furious and Plato is anxious, the fragmentalist should say that there is no fragment which comprises both the fact that Socrates is furious and the fact that Plato is anxious, and accordingly that these two facts do not cohere*; and yet, these two facts cohere, in so far as it could have been the case that, at some time, Socrates is furious and Plato is anxious. Coherence* entails coherence, but the converse does not hold. (Correia and Rosenkranz 2012: 312).

If Fine's notion of coherence* were like the ordinary notion of coherence, then the reduction of times to fragments would be inadequate due to an over-generation of fragments. We would get more times than there actually are.

So the fragmentalist's notion of coherence is not charitably taken to be the ordinary notion of coherence. But if we cannot identify the notion of coherence with a primitive that we are already familiar with, then the fragmentalist incurs the burden of elucidating it. On this, Fine writes:

One might want to take the notion of coherence as fundamental in addition to the notion of reality. One would then expect there to be various substantive 'rules of coherence' concerning the conditions under which a set of facts would be coherent and the way in which the coherence of one set of facts might constrain the coherence of another set. For example, in the classic tense-logical case, one would want that if *f* were to cohere (i.e. to be simultaneous) with *g* and *g* were to cohere with *h* then *f* would cohere with *h* (though, within a relativistic setting, this rule would have to be dropped). (Fine 2005b: 281-282).

So coherence is a fundamental notion that is to be elucidated in terms

of rules that govern it.¹⁰ The example rule is that the relevant notion of coherence is transitive.

The transitivity of coherence requires that fragments do not overlap. However, intuitively, they do overlap. Imagine that KF and ML are both sitting, and then ML stands up while KF continues to sit. If the fact that KF is sitting coheres with the fact that ML is sitting (given the earlier situation), but also coheres with the fact that ML is standing (given the later situation), then the fragments would overlap and the transitivity of coherence would dictate that the fact that ML is sitting coheres with the fact that ML is standing, which certainly do not cohere. So if we indeed assume transitivity, the fragments should not overlap in this way.

We can avoid the overlap if we assume a suitable view of the identity of facts and allow distinct token facts of the same type, so that each fragment consists only of its own collection of momentary facts. Now Fine seems indeed to accept such token momentary facts in the guise of token events, writing that ‘it might be supposed, for example, that there are [...] token events that can occur only at a given time’ and that ‘each current token momentary event would give rise to a tensed fact, the event’s occurring, which could obtain only at the current time’ (Fine 2005b: 318). If we apply this suggestion in the above example, there is the first token event e_1 of KF’s sitting, and there is the second token event e_2 of KF’s sitting, and only the first coheres with the (token) event of ML’s sitting, and only the second coheres with the (token) event of ML’s standing. In terms of the official idiom, we now say that $\mathcal{R}(\text{event } e_1, \text{ of KF's sitting, occurs})$ and that $\mathcal{R}(\text{event } e_2, \text{ of KF's sitting, occurs})$, where e_1 and e_2 are distinct events.

On this view, fragmentalism comes with a commitment to token events. But such time-specific entities do not seem to be in the spirit of tense realism, after all, we distinguish the different token events first and foremost by virtue

¹⁰One might have thought that fragmentalism is simply the view that denies that if $\mathcal{R}A$ and $\mathcal{R}B$, then $\mathcal{R}(A, B)$, so that our overall conception of reality would take this form: $\mathcal{R}(A, B)$, $\mathcal{R}(A, \neg B)$, ... We do not then need a separate notion of coherence. But this is not Fine’s view, in fact, in a footnote, he suggests the opposite direction of reduction, that ‘given a primitive relation of coherence, one might take a fact to belong to reality when it is self-coherent’ (Fine 2005b: 281 fn.13).

of the fact that one occurs at one time and the other at a different time. In fact, depending on one's view of events, one might want to say that the event e_1 , of KF's sitting, occurs in virtue of or because KF sits at t , where t is the time that the e_1 is specific to.¹¹ Instead of saying that $\mathcal{R}(\text{event } e_1, \text{ of KF's sitting, occurs})$ we would then say that $\mathcal{R}(\text{KF is sitting at } t)$, but the latter is arguably not a tensed fact. So the fragmentalist not only incurs a commitment to token events, she also requires a particular view of events that is somehow congenial to her realism about tense.

Independent of the commitment to events, I find it more natural to think that it is not all flux, with new token entities popping into existence at every moment in time, but that there is also stability in the continued obtaining of certain facts. If there is indeed such stability across time, then — contra Fine — we should rather deny that coherence is transitive, even in non-relativistic cases.

Another open question is the way in which the notion of coherence bears on our understanding of what it is for facts to be incompatible. The central thesis of fragmentalism is that reality can be constituted by incompatible facts. But what is it for facts to be incompatible? If we think that facts are incompatible if and only if they cannot both obtain in reality, then one will be unable to make sense of a view that says that incompatible facts can both obtain in reality. That is, if we plug this definition of incompatibility into the formulation of the rejected *Coherence* principle, fragmentalism becomes the view that pairs of facts that cannot belong to reality belong to reality. So if we cling to the standard understanding of incompatibility, then fragmentalism risks being outright unintelligible.

The fragmentalist must offer some understanding of incompatible facts that allows us to see how two incompatible facts can both belong to reality. The most promising proposal is that facts are incompatible if and only if

¹¹Some examples: Quine (1960: 171) holds that events are spatiotemporal regions; Lewis (1986a) holds that events are properties of spatiotemporal regions; Kim (1993) holds that events are compounds $\langle x, F, t \rangle$ of an object x , property F , and time t ; or Chisholm (1976: 126) holds that events are states of affairs that are concretized at a place and time. None of these seem to be in the spirit of tense realism. For a tense realist critique of irreducible reference to events, see Prior (1962/2003).

it is metaphysically impossible that they cohere. We must understand coherence in such a way that the compatibility of facts can be understood as the possibility of such coherence. But to the extent we do not know what exactly can cohere with what, we do not know whether this is an adequate understanding of the compatibility and incompatibility of facts.

To be sure, it would be unfair to demand an exhaustive definition of a notion that Fine wants to introduce as a primitive. The complaint is just that we need some elucidation, but both the name that Fine picks for the notion (namely ‘coherence’) as well as the formal feature attributed to it (namely transitivity) appear to point us in the wrong direction.

2.4 A fragmentalist A-theory

Fine aims to argue that fragmentalism is the best way to believe in the reality of tensed facts. Here I want to illustrate what tense realism looks like in the framework of *Chapter 1*, and clarify certain differences with Fine’s view. This version of tense-based fragmentalism answers the open questions that have been raised in the previous section, §2.3. The outlined view fails to answer the deeper worries raised in §2.2 however, and in part for that reason, I do not endorse the theory outlined here. I will discuss the issues that remain at the end of this section; in the next chapter I will propose a fragmentalist conception of time that, I will argue, is more satisfactory.

Within the framework of *Chapter 1*, Fine’s notion of coherence is understood as the co-obtainment of facts. A first point to note here is that co-obtainment is such that, if the proposed framework is combined with a Finean view of the concept of reality, it can be a much more natural concept. For example, we can allow each of the following inferences:

$$\mathfrak{R}(A) \models A$$

$$\mathfrak{R}(A \circ B) \models \mathfrak{R}B \circ \mathfrak{R}A$$

$$\mathfrak{R}\neg A \models \neg\mathfrak{R}A$$

The most important of these concerns the factivity of reality. This isn't simply plausible; it's also a gain in neutrality. The fragmentalist can adopt Fine's meta-metaphysical commitments and distinguish between what is the case and what is the case in reality as Fine has proposed, but she can choose not to adopt these meta-metaphysical commitments. She does not have to be concerned with what is the case in reality but may simply hold a fragmentalist view with regard to what is the case. So let me proceed in that way, and ignore reality statements from here on, in the understanding that anything that will be claimed for any sentence whatsoever will hold *ipso facto* for a sentence concerning reality along the lines proposed by Fine.

Resorting to the framework proposed in *Chapter 1*, our description of a fragmented world that is constituted by tensed facts will take the following form:

$$\begin{aligned}
& \dots \wedge [\neg P(\text{Aristotle is sitting}) \circ \neg(\text{Aristotle is sitting}) \circ F(\text{Aristotle is sitting}) \circ F\neg(\text{Aristotle is sitting}) \circ \dots] \wedge \dots \\
& \dots \wedge [\neg P(\text{Aristotle is sitting}) \circ \text{Aristotle is sitting} \circ F(\text{Aristotle is sitting}) \circ F\neg(\text{Aristotle is sitting}) \circ \dots] \wedge \dots \\
& \dots \wedge [P(\text{Aristotle is sitting}) \circ \text{Aristotle is sitting} \circ \neg F(\text{Aristotle is sitting}) \circ F\neg(\text{Aristotle is sitting}) \circ \dots] \wedge \dots \\
& \dots \wedge [P(\text{Aristotle is sitting}) \circ \neg(\text{Aristotle is sitting}) \circ \neg F(\text{Aristotle is sitting}) \circ F\neg(\text{Aristotle is sitting}) \circ \dots] \wedge \dots^{12}
\end{aligned}$$

Note the use of the standard past tense operator 'P' ('it was the case that...') and standard future tense operator 'F' (it will be the case that...). So, 'P(Aristotle is sitting)' says that it was the case that Aristotle is sitting and 'F(Aristotle is sitting)' says that it will be the case that Aristotle is sitting. Furthermore, any unembedded sentence 'A' is read as present tensed, stating how things are now. All the tensed facts that obtain at a single moment in time are understood as co-obtaining together.

Note that we will say, for example, that (1) ML is sitting \circ KF is sitting, as well as that (2) ML is sitting \circ KF is standing. We can allow for such

¹²As we have been understanding the notion as a binary connective, it should be clear that I have left out some unnecessary bracketing. It might well turn out that co-obtainment is more aptly conceived of as a multigrade or even infinitary connective.

overlap since co-obtainment was stipulated to be non-transitive: the joint truth of these two co-obtainment claims do not imply that (3) ML is sitting \circ ML is standing. This thus answers the transitivity issue that I raised for Fine's notion of coherence.

We also denied that any two compatible facts necessarily co-obtain. This answers the worry raised by Correia and Rosenkranz that, if we understand Fine's notion of coherence as ordinary coherence, there is an overgeneration of times. There is no temptation to understand co-obtainment as ordinary coherence or any reason why two facts that cohere in the ordinary sense also co-obtain.

We have also discussed in *Chapter 1* how the incompatibility of facts can be understood in terms of the impossibility of their co-obtainment, which we distinguished from contrariety, understood as the impossibility of the obtaining of their conjunction. We thus distinguish between two coherence principles:

Incoherence 1: there are sentences A and B , which state incompatible facts, and which are such that it's the case that A and B . For example: it's the case that KF is sitting and the case that KF is standing.

Incoherence 2: there are sentences A and B , which state contrary facts, and which are such that it's the case that A and B . For example: it's the case that KF is sitting and not the case that KF is sitting.

The proposed understanding of fragmentalism accepts *Incoherence 1* and denies *Incoherence 2*. Given that the notion of incompatibility that features in *Incoherence 1* is understood in terms of co-obtainment, there is no problem in making sense of the principle.

So we have answers to all the issues raised in §2.3: we allow the factivity of reality, we can allow the descriptions of the facts at different times to overlap, and we have made clear in what sense reality harbours conflicting facts and in what sense it doesn't.

Unfortunately, I don't think the outlined theory helps in any way with the issues raised in §2.2. The argument from truth suggests to me that there

is something deeply suspect about tense realism, *any* kind of tense of realism. Just to refresh our memory, consider again two of the principles that drove the argument from truth, *Truth stability* and *Content stability*:

Truth-value stability: if an utterance is true, then it is always true.

Content stability: if an utterance states the proposition that P , then it always states the proposition that P .

The plausibility of Content stability and Truth stability seem to me to be rooted in an atemporal understanding of predication, the sort of predication for which there is no room in a tense-based framework, and yet which seems to be part of our ordinary ascriptions of truth to token utterances. As Fine notes:

Our ordinary notion of truth, as applied to utterances, appears to be stable. Suppose I utter the words ‘I am sitting’ while sitting; and suppose that a few minutes later I stand up. Someone may then ask ‘is that utterance KF made five minutes ago true?’. The correct answer is surely ‘Yes’, despite the fact that I am now standing. (Fine 2005b: 293).

Fine seems absolutely right to me, the correct ordinary answer is surely ‘Yes’. But I think the intuitive sense in which we say that an utterance of five minutes ago is true is not a temporal sense. When someone asks ‘is that utterance KF made five minutes ago *now* true?’, I would deny that it is; and similarly when someone asks ‘is that utterance KF made five minutes ago *always* true?’, I wouldn’t answer Yes. The utterance isn’t *now* true, nor is it *always* true, it is *just* true. There seem to be no temporal aspect to the predication of truth to the utterance. The problem with a tensed framework is that it cannot but assimilate such an atemporal (or sempiternal) understanding of predication to an eternalist understanding of it, and the reason for this lies in the convention that - in a tense-logical framework - any sentence states how things are currently: that A is equivalent to $\text{NOW}(A)$. We will discuss this convention in more detail in §3.1. If there is to be an

atemporal predication, this convention must be denied: it must be possible that *A* even though it is not now the case that *A*.

Atemporal predication is plausibly involved in our descriptions of abstract objects.¹³ If I say that 3 is more than 2, I do not make a claim about how numbers are right now, nor about they always are; it's simply a claim attributing a relation to numbers, without situating this in time in any way. Compare this with spatial location. Say we introduce sentences of the form, 'THERE(*A*)', saying that it's over there the case that *A* and furthermore introduce the convention that any sentence '*A*' is read saying that it's right here the case that *A*.¹⁴ One problem of such a framework is that we can state that things are the case without implying that it's the case *here*, *there* or *anywhere* in space. If I say that '3 is larger than 2' I would be making a claim about the way things are here; and I'm not doing that. The problem is not that it is parochial; I could prefix the statement with 'Everywhere(3 is larger than 2)' to widen the spatial region I'm concerned with, but the real problem remains: I just do not mean to imply anything about the way things are in a spatial region, no matter how wide that region is. I simply mean to attribute a relation to two things, the numbers 3 and 2, in a spatially innocent sense. Similarly, at any time at which you ask me whether 3 is more than 2, I will answer Yes, but this does not mean that I make a claim *about* the way things are at that time. My claim is temporally innocent. Similarly, at any time at which you will ask me whether some particular token utterance is true, I will answer Yes, but this does not mean that I make a claim about the way things are at that time.

The *Content stability* and *Truth stability* principles, are intended by Fine to capture 'the way in which we are willing to adopt an eternal perspective of what the truth of a tensed utterance might require of reality' (2005: 294). The eternalist perspective is one within which the token is always true. But given that there is an intuitive discrepancy between the claim that I will always answer Yes when asked whether a past utterance is true (as such), which

¹³Fine (2005a: §) notes that the predications of certain properties, such as being self-identical or being a man, are also not plausibly understood as tensed predications.

¹⁴For a discussion of such a location-modal framework, see Simons (2006).

seems very plausible, and the claim that I will answer Yes when asked whether a past utterance is always true, which sounds a whole lot less plausible to me, it seems that we are not so much willing to adopt an eternal as an atemporal (or sempiternal) perspective.

The tensed framework leaves little room for atemporal predication. If I can indeed make tensed claims without adopting the the point of view of the current moment, this undermines tense realism. The whole point of resorting to tensed notions in one's metaphysics is that facts constitute the way things are right now, and that this is the only way the world is. So I do not see any plausible way out of the argument, not without abandoning tense realism.

2.5 Summary and concluding remarks

I do not want to claim that the proposed understanding of fragmentalism is the only way of filling out the view proposed by Fine, or that it is the view that offers the best fit with Fine's characterizations. I do want to claim however that the form of fragmentalism defended in this dissertation is a sensible version of the view. If one insists on tense realism, and is attracted to the fragmentalist take on it, then the outlined view might be attractive. Let me make a slightly stronger claim: if one is a tense realist fragmentalist, then this should be the view that one adheres to. But, as discussed, I do not think that we should not ultimately cling to tense realism. We have already seen one reason for this: the argument from truth. In the next chapter I will discuss another reason, and then propose a version of fragmentalism that is not based in tense but in a primitive notion of passage.

Chapter 3

Passage in a Fragmented World

Time passes. The fact that you are reading these words passes into the fact that you are reading these words instead, which passes into the fact that you now reading these words, and so it goes. When we describe the contents of one time as passing into the contents of a next time, we describe the world from an atemporal point of view. There are the contents of each time and, *falling outside of these*, there is the passing of each into another. When we adopt such an atemporal view on which all times are on par, in order to admit the passage between them, it seems that we adopt an incoherent view, given that it's only the passage of a fact into a contrary fact that can make for a genuine change of the objects involved in those facts. Many conclude that passage and change are therefore incoherent notions that we should dispense with. Fragmentalism enables us to explore a view of time that accepts this passing of one fact into another as the most basic temporal notion, instead of the usual A-theoretic and B-theoretic notions, even though it involves contrary facts.

The chapter consists of the following four sections. Section 1 argues that both the standard A-theory and fragmentalist A-theory fail to capture the passage of time. Section 2 spells out a passage-based view of time, which I will call the passage theory. Section 3, finally, offers a new solution to the problem of change on the basis of the proposed passage theory of time. Also included is an appendix that discusses the truth conditions of temporal language in the

context of the passage theory, and a second appendix that discusses whether there are translations from the passage theory into a standard A-theory. The main objective of this chapter is simple: to describe an understanding of time that, I believe, should be of intrinsic interest. I will be fairly quick when I discuss worries concerning the more standard theories in the understanding that the worries are not meant to refute these theories but only to offer reasons for exploring a different view.

3.1 Passage and A-theories

There is a widespread view that for time to pass is for certain tensed facts to obtain. We have met the tensed framework in *Chapter 2*. Let an A-theory be any theory that resorts to tensed descriptions of the world.¹ Just as in *Chapter 2*, the tensed facts are stated with the help of a past tense operator ‘P’ (‘it was the case that...’), a future tense operator ‘F’ (‘it will be the case that...’) and there is a further convention that any unembedded sentence ‘A’ is read as stating how things are *right now*.² So, ‘P(A)’ says that it was the case that A or, equivalently, that it’s right now the case that it was the case that A.³

One might think that time passes when certain things *were* the case and *will* be the case that aren’t the case right now. Thus, for example, Prior:

‘It was the case that *p*, but is not now the case that *p*’ - this formula continues to express what is common to the flow of a

¹This is not how McTaggart originally introduced the A-theory, but I think it’s nevertheless the defining characteristic of the relevant family of views. Though a tensed conception is often combined with a presentist view (according to which only current objects exist), it naturally features in other conceptions of time as well. Even on a growing block view, one will want to say that the block *will* include more. And even on a moving spotlight view, one will want to say that a different time *will* be qualitatively privileged (see Sider 2001:22).

²The convention doesn’t hold for sentences embedded under tense operators. Kamp (1971) showed that prefixing embedded sentences under a now operator can affect the truth-value of the sentence it’s embedded in.

³The most detailed examination of this framework is still found in Prior (1967).

literal river on the one hand (where it was the case that such-and-such drops were at a certain place, and this is the case no longer) and the flow of time on the other. (Prior 1962/2003: 19).

Though this formula - it was the case that p and it is now not the case that p - may be true only if there is passage in the world, I do not think that it captures or describes this passage.

Given that any sentence A states what is the case right now, any fact whatsoever is a current fact; in particular, that something *will* be the case and that something *has been* the case are all themselves current facts. At the heart of the standard A-theory lies the convention that any sentence states what is the case right now, that any sentence is merely descriptive of the current state of the world. This means that any sentence only ever specifies the contents of a single moment in time. To state that something obtains now is just to describe more of the current stage in history, and not the passing from that stage of history to the next. When we understand the world in terms of the tense-logical framework, we understand it as consisting of a single time-slice that includes forward-looking and backward-looking facts. Price made the objection vivid: ‘[W]hat did God need to create, in order to create the whole of reality, as our exclusive presentist describes it? Not a long series of worldstages, but just a single moment, complete with its internal representation of a past and future’ (Price 2009: 5).

The closest that the standard A-theory comes to capturing the passage of time is in the constant rewriting of its description of the world. It states that the world is (now) this way. And then we wait. And then it states that the world is now *this* way. But the crucial bit is in the waiting, that is where time passes, and that isn’t captured in any of the descriptions that the theory offers us. To emulate passage is not to capture it, just as we do not capture the nature of redness by writing up our theory in red ink. The tensed descriptions do not offer an animated picture of the world simply because it includes bits that fix what is to come, and what came before.⁴ What it really offers us, across time, are the still snapshots of that which passes away *if and*

⁴Cf. Park (1972) and Savitt (2002).

when time passes and not a picture of that very passing itself. A passing picture isn't a picture of passage.⁵

The culprit seems to be the convention that any statement is a statement about the way things are right now, which privileges the perspective of a single time. It's a natural thought, then, that we should free ourselves from this, and describe the world from an atemporal point of view on which all times are on par. This brings us of course to the fragmentalist A-theory I proposed in §2.4. We step back from our temporally embedded perspective, and take the tensed contents of any time to be all equally part of reality. We admit the tensed facts of past times as they were back then (so we admit the fact that Aristotle is sitting), and we also admit the tensed facts of future times as they will be in due time (so we admit the fact that a human is walking on Mars). All the tensed contents of all times are all deemed equally real.

Does the fragmentalist A-theory succeed in capturing the passage of time? It seems not. We have at most the contents of which we might want to say that one passes over into another, and still not the passing itself. Fine notes this:

[C]learly, something more than the equitable distribution of presentness is required to account for the passage of time. But at

⁵Fine offers a slightly different objection to the standard A-theory. Consider an A-theoretic view according to which there is a range of times and one time *is* present, some times *were* present, and some times *will be* present. And compare this to a tenseless view on which there is a range of times, ordered by an earlier than relation. Fine argues that the only substantive difference between such tense-based and tenseless views is that the tense-based view singles out one time as the present. He argues as follows. Imagine we also single out one time *t* within the tenseless view, and call it the present. Does this single change in the tenseless view suffice for it to capture the passage of time? Clearly not. All we have is the same static time series, only now with one of its elements designated. And yet it's entirely equivalent to the tense-based view: when the tense-based theory says that a certain time is present the tenseless theory also says it is present. When the tense-based theory say some time *will be* present, the tenseless view says that the time is later than *t* and *t* is present; and when the tense-based view says that some time was present, the tenseless view says that it's earlier than *t* and *t* is present. '[The tense-based] conception of temporal reality,' Fine writes, 'once it is seen for what it is, is as static or block-like as the anti-realist's, the only difference lying in the fact that his block has a privileged centre' (Fine 2005b: 287).

least, on the current view, there is no obvious impediment to accounting for the passage of time in terms of a successive now. We have assembled all of the relevant NOWs, so to speak, even if there remains some question as to why the relationship between them should be taken to constitute a genuine form of succession. (Fine 2005b:288).

Fragmentalism is only a necessary part of a theory that captures the passage of time. It's not the full story.

So, then, what is the full story? Tallant, in a discussion of Fine's fragmentalism, sees various obstacles to the addition of a relation that accounts for passage. First, he notes:

[S]uch a relation would have to be *in* neither of the fragments that it relates – it must bridge the gap between them. The first (obvious) problem is that it is entirely unclear what sort of relation is suited to relating distinct fragments of reality. (Tallant 2013: 12-13)

Tallant is right: passage cannot be a further way things are 'at' a moment in time. The relevant relational fact could not simply be *just-more* momentary content, after all, we are not interested in further additions to the momentary states of the world, we are interested in the passage of each state into another momentary state. Tallant continues:

The second problem is that, even if we can locate a relation to relate the distinct fragments, it remains unclear how this relation is to suffice for passage. [...] Only particular relations can generate temporal order. If that is right, and the relation between distinct fragments of reality is temporal, then presumably said relation will have to be the tenseless 'earlier than' and 'later than' relation, that is the fundament of the B-theory. (Tallant 2013: 13).

We need a relation that doesn't just hold 'at' a time and which is yet a *temporal* relation. But Tallant is too quick in thinking that it must then be the earlier-than relation that is added to the fragmentalist's view.

In fact, we can be sure that it's of no help to add the earlier-than relation, as the temporal order that this introduces *is* already captured by the tensed contents of Fine's fragments. Let $frag_1, frag_2, \dots$ refer to the fragments, understood as certain sets of tensed facts, and let $[A]$ refer to *the fact that A*.⁶ We can then define an earlier-than notion, symbolized with ' \leq ', as follows (following Meyer 2013: 61):

$$frag_1 \leq frag_2 \text{ iff, if } [A] \in frag_1, \text{ then } [PA] \in frag_2.$$

The earlier-than relation can therefore be defined in terms of the tensed contents of the fragments. So it's not temporal order that is lacking from the fragmentalist's view of the world. We have no problem in ordering the fragments in such a way that they follow the trajectory of actual history, and *yet* we feel that the very passing of time is absent. Mere temporal order just doesn't make for the passage of time.

This of course aggravates the first question, of what the fragmentalist could add to the 'equitable distribution of presentness' in order to capture the passage of time. There are at least two possibilities. Either we find something, distinct from passage itself and distinct from the earlier-than relation and tensed notions, which - perhaps together with the tensed facts - constitutes the passing of one fragment into the next. I don't see what this further notion could be. A second possibility is that we started in the wrong place by assuming that we could analyze the passage of time in terms of tensed notions or a combination of tensed notions and something else. Perhaps there is something to our concept of passage that is *sui generis*, and cannot be captured in any other terms. The very fact that many of us can recognize that passage seems lacking from the standard and fragmentalist A-theories, suggests that at least many of us possess a concept of passage that is not exhausted by any of the theoretical primitives currently at play in these theories; and if we possess such a concept, we are free to add it to our theories wholesale and precisify this concept of passage without trying

⁶Of course, as we discussed, this talk of facts and fragments as things is mere loose talk, and not the idiom that reflects the fragmentalist's conception of the world (Fine 2005b: 268. This does not matter for the point at issue here.). The difference between the strict and loose talk doesn't matter for the current discussion.

to define it.⁷ We forego an analysis of the passage of time and say that the contents of one time pass over into the contents of another time *in the very sense in which this seems absent from the more standard A theories*.

Could we simply replace the earlier-than notion in B-theoretic views with the notion of passage? It's not a matter of introducing some two-place notion and labelling it 'passage'; it needs to live up to its label. We will see that Fine was right to assume that a fragmentalist framework is required for this. We will discuss this in more detail below but, roughly, the reason is that passage should make for genuine change (cf. McTaggart 1908:459), and there is genuine change only if a fact passes into a contrary fact. But we also saw in our brief discussion of the standard A-theory that we require an atemporal perspective on which all times are on a par, and so we need to make room for contrary and incompatible facts in our conception of the world. This, I believe, requires a conception on which the world is fragmented.

3.2 A Passage theory of time

The fragmentalist framework affords us with an atemporal view of the world without being forced to take the facts that constitute the world at various times all to co-obtain, or be compatible. We take the way that reality is at a single time to consist in the co-obtainment of facts that capture the instantiation by objects of the various properties they have at that time. Each such fact co-obtains with every other in the relevant collection, and not all of which co-obtain with the facts that constitute the world at a different time. So the overall conception of the world, thus far, is reflected in a description of the following form:

$$\dots \wedge [\neg(\text{Plato is alive}) \circ \neg(\text{Aristotle is alive}) \circ \dots] \wedge \dots$$

$$\dots \wedge [\text{Plato is alive} \circ \neg(\text{Aristotle is alive}) \circ \dots] \wedge \dots$$

⁷The notion of passage is of course the subject of much debate. Influential criticisms of passage are found in Smart (1949), D. C. Williams (1951) and Price (2011). Considerations in defence of passage are found in Earman (1974), Norton (2010) and Maudlin (2007: Ch. 4). The assumption of this chapter is of course that there is a meaningful notion of passage.

... \wedge [Plato is alive \circ Aristotle is alive \circ ...] \wedge ...

... \wedge [\neg (Plato is alive) \circ Aristotle is alive \circ ...] \wedge ...⁸

The long co-obtainments that feature in these descriptions state what is the case at various times as facts that all mutually co-obtain. Roughly, then, whenever two facts obtain simultaneously, the fragmentalist will say that they co-obtain.

But of course, nothing thus far represents the passing of time. We have already seen that the addition of tensed facts to the fragments do not help us any further in describing the passing of time (§3.1). Indeed, I don't believe that the above facts are sensibly understood as present tensed facts at all but should rather be understood as tenseless descriptions. If the claim that 'Aristotle is alive' is understood as saying 'Aristotle is *now* alive', then, in treating all times on a par, the fragmentalist is claiming that Aristotle is now alive. But Aristotle isn't now alive. The fragmentalist's predications are more naturally understood as the tenseless predications used by the B-theorist.⁹ To say that Aristotle is alive is just to say that Aristotle instantiates a certain property. It's not to say that he instantiates the property *now*, nor that he *always* or *eternally* instantiates the property; he simply has the property. The predication of properties and relations to objects is stripped from any temporal or (or indexical) meaning.

Instead of resorting to tense operators, let us appeal directly to the very notion of passage that seems to be lacking in other theories. We use many metaphors when we describe the passage of time: we compare time for example to 'a river' that 'flows' and 'carries' us into the future. The resort to such metaphors is taken by many to show that passage is a confused notion.¹⁰ Another (obvious but somewhat neglected) possibility is that passage is a sui-generis phenomenon, any description of which in different terms is bound

⁸As we have been understanding the notion as a binary connective, it should be clear that I have left out some unnecessary bracketing. It might well turn out that co-obtainment is more aptly conceived of as a multigrade or even infinitary connective.

⁹Cf. Tallant: 'If, as Fine claims, the light of presence is spread equitably throughout all time, then it is simply unclear that this is *substantively* different from anti-realism about tense.' (2013: 13).

¹⁰See, e.g., Smart (1949) and D. C. Williams (1951).

to be metaphorical precisely because it is a basic temporal phenomenon.¹¹ The metaphors are helpful to elucidate the notion of passage in certain ways, and to fix its sense, but they are also risky. When we describe the passage of time as a ‘flow of time’ or a ‘moving now’, this suggests that passage is a type of change or movement. But, I believe, the passage of time has to be carefully distinguished from change and movement. As Maudlin notes:

Except in a metaphorical sense, time does not move or flow. Rivers flow and locomotives move. But rivers only flow and locomotives only move because time passes. The flow of the Mississippi and the motion of a train consist in more than just the collections of instantaneous states that have different relative positions of the waters of the Mississippi to the banks, or different relative positions of the train to the tracks it runs on. The Mississippi flows from north to south, and the locomotive goes from, say, New York to Chicago. The direction of the flow or motion is dependent on the direction of the passage of time. Given the essential role of the passage of time in understanding the notion of flow or motion or change, it is easy to see why one might be tempted to the metaphor that time itself flows. (Maudlin 2007: 110).

Notions such as ‘flow’, ‘change’ and ‘movement’ are parasitic on the passage of time; the latter is a precondition for any of the former to occur and should not be confused with it.¹² Indeed, I will assume that the passing of time is constitutive of stability as much as it is constitutive of change across time. The fact that you sit passes into the fact that you sit, which passes into the fact that you sit, and so on for as long as you continue to sit. The passing of a fact into itself is what constitutes the stability of that fact across the stretch of time that is constituted, in part, by that passage.

¹¹The suggestion seems to be made by Broad: ‘I do not suppose that so simple and fundamental a notion as that of absolute becoming can be analysed...’ (1938/1976: 281).

¹²Cf. Prior: ‘A natural first move towards extricating ourselves from [certain] perplexities is to admit that talk of the flow or passage of time is just a metaphor. Time may be, as Isaac Watts says, *like* an ever-rolling stream but it isn’t really and literally an ever-rolling stream. But *how* is it like an ever-rolling stream?’ (Prior 1962/2003: 7).

To introduce the notion of passage into our metaphysics, we can introduce a sentential passage-operator ‘ \leftrightarrow ’ (‘... passes into ...’) and we can again turn to models that precisify the notion of passage. We can use a simple adaptation of the basic semantics for co-obtainment that we saw in §1.2. We add sentences of the form $A \leftrightarrow B$ to the language, and take a model M to be a triple $\langle T, O, v \rangle$, where T is a set of points, v is a function that assigns 1 or 0 to the atomic sentences relative to each point in T , and O is a set of pairs of points in T representing an order relation on T , being irreflexive, anti-symmetric and transitive.

The valuation v for the atomic sentences relative to points in T is first extended to a valuation for all the sentences via the following recursive clauses (where t ranges over points in T):

$$\begin{aligned} v_t(A \circ B) &= 1 \text{ iff } v_t(A) = 1 \text{ and } v_t(B) = 1, \\ v_t(A \wedge B) &= 1 \text{ iff } v_t(A) = 1 \text{ and } v_t(B) = 1, \\ v_t(\neg A) &= 1 \text{ iff } v_t(A) \neq 1, \\ v_t(A \leftrightarrow B) &= 0. \end{aligned}$$

Note that any passage sentence (i.e. any sentence of the form ‘ $A \leftrightarrow B$ ’) is false at each of the points in T . This captures the intuition that passage is not itself part of that which passes. We will discuss this below. The valuation v is further extended to an evaluation of the sentences relative to each pair $\langle t_1, t_2 \rangle \in O$:

$$\begin{aligned} v_{\langle t_1, t_2 \rangle}(A \leftrightarrow B) &= 1 \text{ iff } (A = B \text{ or } A = \neg B \text{ or } B = \neg A) \text{ and} \\ &v_{t_1}(A) = 1 \text{ and } v_{t_2}(B) = 1, \\ v_{\langle t_1, t_2 \rangle}(\neg A) &= 1 \text{ iff } v_{\langle t_1, t_2 \rangle}(A) \neq 1, \\ v_{\langle t_1, t_2 \rangle}(A \circ B) &= 1 \text{ iff } v_{\langle t_1, t_2 \rangle}(A) = 1 \text{ and } v_{\langle t_1, t_2 \rangle}(B) = 1, \\ v_{\langle t_1, t_2 \rangle}(A \wedge B) &= 1 \text{ iff } v_{\langle t_1, t_2 \rangle}(A) = 1 \text{ and } v_{\langle t_1, t_2 \rangle}(B) = 1, \\ v_{\langle t_1, t_2 \rangle}(p) &= 0. \end{aligned}$$

Note first of all that the only true passage sentences are those that feature either the same sentence on both sides, or a sentence and its negation. Note secondly that atomic sentences are false at the points in O . We will discuss

the motivation for these clauses below. The clauses for the points in T and O together fix the truth of each sentence in a given model, which is defined via the following recursive clauses (where x ranges over $T \cup O$):

$$\begin{aligned}
M \Vdash p &\text{ iff } \exists x(v_x(p) = 1), \\
M \Vdash A \circ B &\text{ iff } \exists x(v_x(A \circ B) = 1), \\
M \Vdash A \leftrightarrow B &\text{ iff } \exists x(v_x(A \leftrightarrow B) = 1), \\
M \Vdash A \wedge B &\text{ iff } M \Vdash A \text{ and } M \Vdash B, \\
M \Vdash \neg A &\text{ iff } M \not\Vdash A.
\end{aligned}$$

Validity and logical truth are defined as before.

This model-theoretic machinery fixes the formal properties of the passage notion, and thus precisifies our ordinary notion without defining it. To see the machinery at work consider the following time-series (with columns representing points in T , and the rows stating sentences that are true relative to those points):

$t_1,$	$t_2,$	$t_3,$	t_4
$A,$	$\neg A,$	$A,$	$\neg A$
$B,$	$B,$	$\neg B,$	$\neg B$
$\neg C,$	$C,$	$C,$	$\neg C$

A model like this determines the following sort of truths. There is the passage of the individual facts, for example, in the case of A we have: $A \leftrightarrow \neg A$ (t_1 to t_2), $\neg A \leftrightarrow A$ (t_2 to t_3), $A \leftrightarrow A$ (t_1 to t_3) and $\neg A \leftrightarrow \neg A$ (t_2 to t_4). These passage facts do not co-obtain, for example, we have that $\neg[(A \leftrightarrow \neg A) \circ (\neg A \leftrightarrow A)]$. Other passage facts do co-obtain however, forming bundles of passings as it were. We have for example: $(A \leftrightarrow \neg A) \circ (B \leftrightarrow B) \circ (\neg C \leftrightarrow C)$ (t_1 to t_2): A 's passing into $\neg A$ co-obtains with B 's passing into B which both co-obtain with $\neg C$'s passing into C . Co-obtaining with these is also the passage of logically complex facts, in particular we have: $(A \circ B \circ \neg C) \leftrightarrow \neg(A \circ B \circ \neg C)$ (t_1 to t_2): the fact that $A, B, \neg C$ all co-obtain passes into the fact that they no longer all co-obtain. The passage of these large co-obtainment facts can be fruitfully compared to the images that 'pass by' when running a film: where the particular images represent the co-obtaining

of the depicted facts (as in the toy model of §1.1), which flow into an image that represent a slightly different set of facts as co-obtaining, and so on.

Let us have a closer look at some of the proposed features of the passage of time, and the rationale behind them. As noted, the only true passage sentences are those that feature either the same sentence on both sides, or a sentence and its negation. The fact that I sit doesn't pass into the fact that it rains, even if it rains at some later time. The fact that I sit either passes into the fact that I sit, or into the fact that I do *not* sit. This reflects the idea that, as Maudlin notes, 'the passage of time underwrites claims about one state "coming out of" or "being produced from" another' (2007:110). The fact that it rains doesn't 'come out of' the fact that I sit, at least not in the sense of 'coming out of' that is sensibly said to be constitutive of time.

One might want to object that the clause is too restrictive in focusing on contrary facts, leaving out the passage from one fact into a fact that is merely incompatible with the first, such as from my sitting into my standing. Introducing such passage-facts raises various difficulties however. There are various technical questions (how do you make the models 'see' the incompatibility of the facts?), but also metaphysical difficulties. In particular, if we introduce the passage from the rose's being red into its being brown (i.e. allow that $R \leftrightarrow B$ can be true), we introduce this passage as something *over and above* the fact that the loss of redness co-obtains with the gain of brownness (i.e. the fact that $(R \leftrightarrow \neg R) \circ (\neg B \leftrightarrow B)$); and it's not at all clear that the passage from one into an incompatible fact really is something over and above the latter bundle of passage facts. We would make things much more complicated with the sole purpose of introducing facts of which it's not at all clear they add anything to what is already captured in the simpler picture given here.

Given that passage may involve contrary facts, it should be clear that passage cannot be a factive notion:

$$A \leftrightarrow B \not\equiv A \text{ and } A \leftrightarrow B \not\equiv B$$

Say that the rose's being red passes into its not being red. If passage were factive, this would imply that the rose is red and not red; and that cannot

be the case. Across time we only accept, say, that the rose is red, and that it's not red *insofar as* it's blue. But though it's not the case that the rose is not red, the passage *from* the rose's being red *to* its not being red *does* genuinely obtain, and so passage is not in general factive. We will see below that this is a central feature of the notion of passage.

The model theory sharply distinguishes between truth at a point in T and truth at an ordered pair in O . At points in T , atomic sentences are true and passage sentences aren't; at points in O , passage sentences are true and atomic sentences aren't. This reflects the idea that, as we saw Tallant put it, '[The required relation] would have to be *in* neither of the fragments that it relates [but] bridge the gap between them' (2013: 12-13). The points in T represent the contents of times, or what happens *in* time. The points in O are representative of the way in which the contents of times pass into the contents of other times, and the very passing of such contents is not itself part of the contents that pass. In more metaphysical terms: the passing of time itself is not *in* time, but constitutes it.

That the passing of time does not lie *in* time has the consequence that it cannot be properly said to pass 'at a certain rate', and this defuses the worry that an objective passage of time gives rise to a regress of time series. Black:

If the claim that time always flows had a literal sense, we ought to be entitled to ask how fast time is flowing. And if so, there would have to be a supertime for measuring the rate of flow of ordinary time. With regard to that super-time, a similar question would immediately rise, viz., how fast it was flowing; and this would imply still another time for its answer. And so on without end. (Black: 56-57).

The simple reply to this is that time doesn't *literally* flow (which is not to say that describing the passage of time as a flow can, metaphorically, bring out the transitory aspect of time). As we already noted: 'flows' are changes that are brought about *as and because time takes things forwards*. It does not 'take' time for time to pass. Changes of objects take time, and for a change to occur *at a certain rate* is *just* for the change to occur together with certain

other fixed kinds of change. A car's movement from M to N happens at a rate of one 100 kilometre per hour when, say, the distance between M and N is 100 kilometre and the dial on a clock shows 1 when the car is at M and changes as the car moves to show 2 when the car arrives at N. The passage of time itself is not however a change that can be measured against other changes; indeed it's not change at all.

Note, thirdly, that since the relation O is transitive, O includes pairs of points in T that are not adjacent in the order induced by order. The contents of one time do not just pass into the contents of the next time (if there is one), they also pass into the contents of *any* time that comes after it. Is this right? Does the way things are now pass into the way things will be in two moments from now? If time passes to the next moment, and the next moment passes into two moments from now, we can say that the current moment *thereby* passes into the way things are two moments from now. The passage from current facts to the next imply the passage from this moment into two moments from now; there seems little room to deny the overarching passage.¹³

These are obviously subtle intuitions, the sort of intuitions that carry some but not too much weight, and which can conflict: there is also something to the thought that the contents of one time become those of the next *only*. There is however also a more theoretical pressure to think of passage in the proposed way. Time is plausibly thought to be continuous, that is, passage is naturally thought to be a passage through a continuous series and not naturally thought to consist of staccato jumps between discrete units of time.

¹³This consideration might be taken to suggest that the passage from the current facts to those of two moments hence is somewhat like a determinable, realized in the passing between the facts of intervening times. Say the following are successively true: A , $\neg A$, A . It then seems tempting to invoke some suitable notion of metaphysical grounding or realization and say that, for example, $A \leftrightarrow A$ because or in virtue of the fact that $A \leftrightarrow \neg A$ and $\neg A \leftrightarrow A$. This may however imply infinite chains of grounding facts in the case where time is continuous, and it may be problematic when we have a time series in which we have: A , A , A , are we going to say that $A \leftrightarrow A$ because or in virtue of the fact that $A \leftrightarrow A$ (and $A \leftrightarrow A$)? This would be an objectionable case of self-grounding. This may in turn be avoided if we hold that the fundamental kind of passage is that of maximal co-obtainment facts into another; but this may run contrary to the intuition that the passing of maximal facts is build up from the passing of facts that are 'part' of those co-obtainment facts. There is more to be said here.

If we were to deny the overarching passage, and replace the order relation in the model theory with its transitive reduction (i.e. a non-transitive order relation), so that one moment only passes into the very next moment, we would then be at a loss to account for the passage of time if time is continuous, i.e. if time is such that between any two moments of time there is another, so that there is no ‘next’ time. Put differently: if time is continuous, which seems plausible, and if the way things are *now* seem to us to pass into the way things are *now*, we cannot deny that overarching passage has thereby occurred.

Fourthly, it should be uncontroversial that the passing of time is an inherently directional affair; that we pass *from* certain facts *to* possibly distinct facts, and hence that passage is not symmetric:

$$A \hookrightarrow B \not\Leftarrow B \hookrightarrow A$$

The lack of symmetry is what provides the direction of time. Much has been written about the direction of time. There are various types of physical asymmetries running along the direction of time, such as the increase of entropy, the expansion of the universe, and the causation of events. Many have explored reductive views of the direction of time, such as the view according to which a state of a system is earlier than another state when the entropy of the second state is larger than that of the first.¹⁴ The passage theory stands in contrast with such reductive views. The direction of time does not emerge from anything that is the case *at the various times*, but part and parcel of the phenomenon of passage. The reductive views make it difficult to see what the content is of various important physical laws. If the direction of time is reduced to the increase in entropy, for example, this seems to suck the content out of the Second Law of Thermodynamics, which surely doesn’t just state that entropy increases as, well, as entropy increases (Maudlin 2007: 129). The non reductive view of passage enables a straightforward reading of physical laws that concern the mentioned asymmetries. On the one hand there is, for example, the increase in entropy, and on the

¹⁴See Mellor (1998: Ch. 12), Zeh (2007) and the papers in Savitt (1995).

other hand there is the direction of passage. So the relevant physical law relates one thing to another: the objective passage of time is what the various asymmetries are relative to.

Much more could be said about the inferential role of passage, in particular about its interaction with the notion of co-obtainment. But these are more of logical than metaphysical interest. Let me simply note that the following claims can all easily be checked in light of the model theory:

$$A \leftrightarrow p \vDash p \text{ and } p \leftrightarrow A \vDash p$$

$$(A \leftrightarrow C) \circ (B \leftrightarrow D) \vDash (A \circ B) \wedge (C \circ D)$$

$$\not\vDash (A \leftrightarrow B) \circ (\neg A \leftrightarrow B) \text{ and } \not\vDash (A \leftrightarrow B) \circ (A \leftrightarrow \neg B)$$

To sum up the picture so far: passage is either that of one fact into itself or that of one fact into a contrary fact, it is neither factive nor symmetric, but it is transitive. There are clearly various choice points, and the choices made above are by now means set in stone. Just as there are many systems of tense logic, so we can naturally expect various systems of passage logic. We could for example explore different logics by changing the order relation O in the model theory.¹⁵ The above seems to me to be one sensible and relatively simple way of precisifying our ordinary notion of passage to a point where it can be put to metaphysical work.

3.3 The problem of change

The passage theory may seem to come dangerously close to a B-theory, given that it too advocates that we describe the world from an atemporal perspective, and it too takes atemporal predications to be more basic than tensed predications. Does the passage theory not just describe a B-theoretic world without a genuine passage of time? It doesn't. When we adopt a B-theoretic conception of the world, we are unable to offer a satisfactory account of the

¹⁵This is how we obtain various systems of tense logic; see e.g. Burgess (2002).

change of objects across time. The passage theory offers a new solution to the problem of change, one which doesn't face the no-change objections faced by accounts of persistence that are based in a B-theory of time.

Imagine a straight young tree growing in the dunes. Because of the relentless sea-wind, the young tree eventually grows into a crooked tree. The young tree is straight, the older tree bent. But the young tree is one and the same thing as the older tree and we subscribe to *the principle of the indiscernibility of identicals*: that, if $x = y$, then for any property F, x is F if and only if y is F. If the young tree and the older tree are one and the same object, and we describe it from an atemporal perspective, then *it* is both straight and bent. This is the so-called problem of change (also known as the problem of temporary intrinsics, or the problem of identity across time).

In the context of a B-theory of time, the task is often taken to be the following. Avoiding contradiction requires that we say, in some sense or other, that there are different times t and t^* and that the tree is straight at t and bent at t^* . Our task is to account for what the truth of these claims consist in; what is it for example for the tree to be 'straight at t '?¹⁶ The following three are the most well-known B-theoretic accounts of what the truth of such time-relativized statements consists in (Lowe 1988: 73):

(A) *Perdurantism*: ' a -at- t is F', where ' a -at- t ' is the name of a temporal part.

¹⁶For example, Johnston: 'This then is the problem of persistence cast in the formal mode: explain the role of temporal qualification in our attributions of change, where explaining does not just mean opting for a style of appending ' t 's and ' t^* 's but defending the views about properties, the nature of time and the nature of persisting individuals which justify this style of appending.' (Johnston 1987: 115). See also, amongst others, Sider (2001: 93), Ehring (1997), and Lowe (1988).

Note however that this way of setting up the problem of change is only appropriate on the assumption of a B-theory of time, and even then it's somewhat objectionable, as it skews our intuitions. The very sensible sounding claim that the same object continues to exist *through time* or *as time passes* becomes, on this way of setting up the problem of change, the odd-sounding claim that the same object is 'wholly present' at different times in the way in which universals are 'wholly present' at different locations, a multiply located object (Lewis 1986b: 202). Our default view is surely endurantism, but not *that* kind of endurantism. The incompatibility with a plausible form of endurantism is bound up with the issue that I will raise below.

(B) *Relationalism*: ‘ a is F-at t ’, where ‘_is F-at_’ is a predicate expressing a relation holding between a and t .

(C) *Instantiationism*: ‘ a is-at- t F’, where ‘_is-at- t _’ is a copula that expresses a time-relative type of instantiation of a property.

Each of these accounts solves the problem of change. But each account also faces the objection that it fails to capture real change. That perdurantism faces a no-change objection is well-known; that the other accounts face the same objection isn’t discussed so much. I will briefly discuss the no-change objection to each approach, before considering a generalized no-change objection to accounts of persistence formulated in the context of a B-theory of time.

(A) *Perdurantism*. The perdurantist account holds that the tree has changed when it has a straight temporal part located at one time and a bent temporal part located at a later time (see e.g. Lewis 1986b: 202-204). There are two distinct objects that are, respectively, straight and bent, instead of one object that is first straight and then bent. This is the way in which it avoids saying that there is one object that has two incompatible properties. The no-change objection to this view is well-known. Neither of the temporal parts changes, the one is simply straight and the other bent. But the spatiotemporal sum that they are part of doesn’t change either, given that it instantiates neither the straightness nor the bentness property and so certainly cannot be said to change *from* being straight to being bent. So neither the temporal parts, nor their sum can be plausibly described as changing (Mellor 1998: 89).

One could plead that having a straight temporal part can be considered ‘a way’ for the sum to be straight, so that there is a sense in which the sum has the properties of being straight and bent, and can be said to vary qualitatively across time *in this sense*. But if having straight and bent temporal parts were indeed genuine ways for the sum to be straight and bent, then the sum is straight and bent after all, bringing us back to the incompatible facts the perdurantist is out to avoid. Having a straight temporal part cannot really be ‘a way’ for the sum to be straight, and hence the sum cannot be said to vary

or change in properties across time.¹⁷ Perdurantism avoids the incompatible facts because the persisting object doesn't have the incompatible properties, but precisely to that extent does it fail as a satisfactory account of change.

(B) *Relationalism*. The relationalist account holds that the tree changes when it stands in a *straight-at* relation to one time and a *bent-at* relation to a later time (as in Mellor 1981: Ch. 7).¹⁸ Here we revise our conception of the involved properties: we are now given an object that stands in two relations to two distinct entities (namely to distinct times), instead of two incompatible monadic properties had by the object across time. The account faces a no-change objection much like perdurantism does. Note first of all that standing in incompatible relations to distinct objects doesn't make for the qualitative change of that object. That I'm taller than one person but smaller than another in no way means that I change (Rodríguez-Pereyra 2003: 191-192).

One might plead that there is something special about relations *to times*; maybe incompatible relations can be said to make for change when they are had *to times*. But even standing in incompatible relations to two distinct times doesn't in general make for change: a single moment of time stands itself in the earlier-than relation to one time and the later-than relation to another time, but surely such a moment of time isn't itself thereby changing. We might again try to say that for the tree to be *straight-at- t_1* is a way for the tree to be straight, and for the tree to be *bent-at- t_2* is a way for the tree to be bent, so that the holding of the relations does make for the tree to be first straight and then bent *in this sense*. But now we no longer avoid the incompatible facts, for the tree to instantiate to instantiate the two relations is for it to be straight and bent, and so it is both, which the relationalist wants to avoid. Just as with the perdurantist account, the relationalist account avoids the incompatible facts but fails to offer a satisfactory account of change.

(C) *Instantiationism*. There are various accounts that might be offered of the time-relative predication invoked by - what I call - instantiationist

¹⁷Cf. Haslinger (1989: 119f.). For a (to my mind unconvincing) reply, see Lewis (2002: 5).

¹⁸Mellor changed his mind and opted for a version of instantiationism in his 1998).

solutions to the problem of change. One might for example be an adverbialist and take the time-relative instantiation to be a *way of having* a certain property (as e.g. in Johnston 1987:128). Or, alternatively, one might think that there are type-facts or type-events which we describe as *the tree's being straight* and *the tree's being bent*, which only have the tree and the respective properties as constituents, but which are 'tokened' or 'located' at different times (as e.g. in Haslanger 2003: §9.3 and Mellor 1998: Ch. 8).¹⁹ There is a genuine question whether this account is really any different from the relational account - but set this aside (see Lewis 2002). Even on its own terms, views like these face the same type of no-change objection we have seen above. Consider the type-fact approach. If the mere existence of the type fact of *the tree's being straight* is a genuine way for the tree to be straight, then the existence of the incompatible fact-types should suffice for the tree to be both straight and bent, in which case the view fails to avoid incompatible facts. If, on the other hand, the existence of the two fact-types doesn't suffice for the tree to be both straight and bent, then the tree simply isn't straight *simpliciter* or bent *simpliciter*, and since it's neither of them, it certainly doesn't change in being first the one and then the other. Yet again, we either fail to avoid incompatible facts or fail to account for change.

I do not intend for any of this to be meant as a conclusive refutation of these accounts of persistence. For the present purposes it is not the individual no-change objections that are the point, it is the pattern running through them. Let *A* describe what the truth of 'the tree is straight at *t*' consists in and let *B* describe what the truth of 'the tree is bent at *t**' consists in. The B-theoretic accounts face a dilemma. First horn of the dilemma: if for it to be the case that *A* does not suffice for the tree to be straight, and for it to be the case that *B* does not suffice for the tree to be bent, then *A* and *B* do not allow us to capture the tree's change from being straight to being bent. Second horn of the dilemma: if for it to be the case that *A* does suffice for the tree to be straight, and for it to be the case that *B* does suffice for the tree to be bent, then *A* and *B* do not allow us to avoid the claim that the tree is both straight and bent. There is either no change, or there are

¹⁹For more non-standard views within this camp, see MacBride (2001).

incompatible facts. Given that the same type of objection applies to very different accounts, this is evidence that it's an essential feature of change, real change, that it involves contrary facts across time.

The generalized no-change objection leaves little room for a coherent account of change. A natural response is thus to say that the generalization in fact shows why we should not be worried about the no-change objection and evaluate the accounts of persistence on other merits. The demand for *genuine* change is an implicit demand for incompatible facts, and hence is not a demand that anyone should be worried about. The objection overgeneralizes.

But another diagnosis of the generalized no-change objection is this: the problem lies with the B-theoretic background view presumed by these accounts of persistence. The B-theory treats the contents of different times on a par and attempts to construct the change of objects from the contents of different times. But change, as we ordinarily understand it, could only emerge from genuinely conflicting facts across time, which the B-theory cannot coherently admit on pain of incoherence. The B-theorist doesn't simply take an atemporal perspective on time, it also pictures the world as one unified spatiotemporal block, no part of which can conflict with any other part. From such a world, change cannot be made to emerge.

The no-change objection doesn't overgeneralize if there is an account of change to which the objection doesn't seem to apply. I believe such an account of persistence falls out of the passage theory of time, which allows us to offer a new solution to the problem of change. Change is plausibly identified with certain kinds of temporal passage. More precisely, different kinds of change are straightforwardly identified with different kinds of passage:

Ceasing to be a certain way: a ceases to be F iff $Fa \leftrightarrow \neg Fa$. For example: the tree ceases to be straight iff $\text{Straight}(a) \leftrightarrow \neg \text{Straight}(a)$.

Coming to be a certain way: a becomes F iff $\neg Fa \leftrightarrow Fa$. For example: the tree becomes straight iff $\neg \text{Straight}(a) \leftrightarrow \text{Straight}(a)$.

Qualitative change: a changes from F to G iff $(Fa \leftrightarrow \neg Fa) \circ (\neg Ga \leftrightarrow Ga)$ and necessarily $\neg(Fa \circ Ga)$. Put informally, an object changes

qualitatively when it loses a property *insofar as* it gains a property that is incompatible with the first. For example: the tree changes from being straight to being bent iff $(\text{Straight}(a) \leftrightarrow \neg\text{Straight}(a)) \circ (\neg\text{Bent}(a) \leftrightarrow \text{Bent}(a))$.²⁰

The passage theory describes a world in which things genuinely change, or so I want to claim. The crucial difference with B-theoretic accounts is that the passage theory can appeal directly to a cross-temporal phenomenon, the passage of time, which involves genuinely conflicting matters. Remember that passage is not factive, for example, we cannot infer from the fact that $\text{Straight}(a) \leftrightarrow \neg\text{Straight}(a)$, that $\neg\text{Straight}(a)$. But, within the passage theory, we do not need the conflicting facts to obtain in order to somehow construct change from them, as on B-theoretic accounts, we only need the *passage from a's being F to a's not being F*. That suffices on itself for the genuine change that we aim to capture. Put differently, there is no need to admit that a is F and a is not F in order to say that the a changes from being F to not F if we can identify the change from a 's being F to a 's not being F with the passage from a 's being F to a 's not being F.

The result is an endurance account in the sense that one and the same object is involved in facts that obtain across time. The result is however not an endurance account in the sense that the same object is 'wholly present' or 'located' at multiple times. When the fragmentalist takes an atemporal view of the world, she doesn't picture it as a unified block, or a series of temporal locations. We have made no reference to times at all, and we do not need to in order to avoid incoherence. The account is thus neutral on the question whether we should admit times to our ontology; if there is such a need, it does not stem from having to solve the problem of change.

Note also that the offered account requires no revisions in our conception of the objects that are involved in change, nor revisions in the intrinsic properties they have across time, nor revisions in the way objects have these properties. The account has no reason to reject any of the following plausi-

²⁰Note that $(\text{Straight}(a) \leftrightarrow \neg\text{Straight}(a)) \circ (\neg\text{Bent}(a) \leftrightarrow \text{Bent}(a))$ is equivalent to $[\text{Straight}(a) \circ \neg\text{Bent}(a)] \leftrightarrow [\neg\text{Straight}(a) \circ \text{Bent}(a)]$.

ble conditions, each of which is rejected by other accounts of persistence (cf. Haslanger 2003):

Identity condition: If an object persists through change, the object before the change is one and the same as the one existing after the change.

Contrariness condition: The way an object is before the change conflicts with the way the object comes to be through the change.

Proper subject condition: The object undergoing the change is the very thing that is the conflicting ways before and after the change.

The tree before the change is identical to the tree after the change; it changes from being straight to not being straight, which are genuinely conflicting (indeed, contrary) ways for a thing to be; and it is the tree itself whose being straight passes into its not being straight. The properties of being straight and being bent remain intrinsic properties. The offered account does not fiddle anywhere with the ontology of objects or properties.

To sum up, when we try to account for change within the context of a B-theory, we seem to always arrive at accounts that fail to capture real change. This does not happen within the passage theory of time. Change can be identified with certain kinds of passage, or the co-obtainment of certain kinds of passage. These kinds of passage involve genuinely contrary facts, and hence do not face a no-change objection. The passage theory is therefore not just another B-theory.

3.4 Summary and concluding remarks

One important gap in the discussion so far is the relation of the passage theory to the special theory of relativity. We have assumed a single foliation of facts constituting one definite order in which things come to pass, and we know that things cannot be that simple. We will discuss this in *Chapter 6*.

One may have noticed how the dialectic has retraced the steps of McTaggart's argument for the unreality of time. McTaggart's argument is well-known, targeting both the A-theoretic and B-theoretic conceptions of time. The A-theoretic determinations of events give rise to an inconsistency when we take an atemporal standpoint and collect together the ways events are *throughout* history, so that the same event seems both past, present and future.²¹ An A-theorist can reply that the mistake is to resort to such an atemporal perspective: if one assumes such an atemporal perspective one will land in inconsistency all right but this just shows that we should never assume such an atemporal perspective and say that an event is both past, present and future (Prior 1967: 5-6). The objection raised in §1 plugs this hole in McTaggart's argument: if we merely describe the way things are now, so that any event is either only past, or only present, or only future, then our description is a snapshot of a single moment, within which no passage or change is to be found either. Passage occurs *from* one moment *to* the next; we *must* stand back from a single moment of time and assume an atemporal perspective if we are to make room for a real passage of time.

McTaggart argued against the B-theoretic conception of time by insisting on the essential connection between passage and change. The same assumption was made in the discussion above. McTaggart:

It would, I suppose, be universally admitted that time involves change. A particular thing, indeed, may exist unchanged through any amount of time. But when we ask what we mean by saying that there were different moments of time, or a certain duration of time, through which the thing was the same, we find that we mean that it remained the same while other things were changing. A universe in which nothing whatever changed (including the thoughts of the conscious beings in it) would be a timeless universe. (McTaggart 1908: 459).

McTaggart's assumption seems plausible to me. Not only is the passage of

²¹On the crucial role of the atemporal perspective in McTaggart's argument, see Dummett (1960: 503).

time a necessary condition for change, its sufficing for genuine change under certain conditions is, vice versa, a necessary condition for the passage of time. If passage did not make for change it would not be passage. The conditions that McTaggart used to show that time is unreal have been used in the passage theory as the necessary conditions that a primitive notion of passage has to meet. The resulting theory is not incoherent, and yet it offers everything that McTaggart demanded of a temporal world. According to the passage theory, McTaggart's fault was to introduce a false choice between two families of primitives, the A and B notions.

When we form our theories about phenomena, our choice of primitives is all-important. If a certain choice of primitives leaves us unable to make sense of an apparently real phenomenon, this is evidence that our choice of primitives isn't right. And this appears to apply to the A and B-theoretic primitives: they fail to capture the passage of time. If we accept the passage of time as a fundamental temporal phenomenon, we can aim to understand it better in other ways, in particular by identifying the various principles and inferential rules that govern it. As we elucidate and precisify passage in this way, we should expect that related phenomena become better understood as well. This, I argued, happens with the phenomenon of qualitative change: once we understand the world in terms of the passage of time, and see what is required of a world in which time passes, a simple account of qualitative change falls into place. What is needed now is a more detailed evaluation of the passage theory in comparison with the A-theories and B-theories. But I believe that things look promising. Explanatory power is everything and the passage theory seems to answer long-standing questions that the standard theories fail to answer (how can there be a genuine passage of time? how can there be genuine change?). At the very least however, I hope that it might open up new avenues for us to explore.

3.5 Appendix: truth-conditions

In this appendix I want to discuss the truth conditions that are plausibly assigned to different bits of temporal language if the passage theory is correct.

One may wonder why this is of any interest. Not so long ago (roughly before the 1980s), semantic considerations were thought to be of central importance to the debate between A-theorists and B-theorists. The crucial question seemed to be whether what tensed sentences say (i.e., whether the proposition that it expresses) is ‘tensed’ in the sense that its truth-value can change with time, or whether the sentence expresses different, eternally true propositions at different times. Because it’s very implausible to think that tensed sentences can be synonymous with tenseless sentences, this seemed an important objection to the B-theory, after all, if there are tensed propositions expressed by sentences, then if they are true, there are thereby tensed facts. The dialectic has however changed. The so-called ‘new B-theory of time’ is happy to admit that tensed sentences have tensed meanings - in some sense - and are therefore not synonymous with tenseless sentences. What is denied now is that this is relevant to our metaphysics of time. It’s argued, roughly, that meaning must consist in propositions (or truth conditions) + some element *x*. This ‘element *x*’ is given different accounts by different philosophers of language. To name two prominent accounts: Kaplan (1977) takes meaning to consist roughly in propositions + ‘character’ and sees the character of a sentence as a body of linguistic conventions or rules associated with it. Perry (1979) takes meaning to consist in what is believed + the way it is believed, where the latter is a kind of cognitive guise or cognitive significance, characterizing not the object of belief but the state of believing that object.²² So we individuate ‘meanings’ more finely without individuating the truth conditions of the truth bearers more finely: an utterance at noon 25/03/2012 of the tensed sentence ‘I am a father’ and the tenseless sentence ‘I am a father at noon 25/03/2012’ can have different meanings, yet the same truth conditions, namely, they are both true if and only if I am a

²²Both Kaplan and Perry resist treating this ‘element *x*’ as a Fregean sense. One reason for this is that they think that the semantic value of indexicals must be the worldly object picked out in context: when I use ‘I’ it must pick out me, and not some concept of me.

father at noon 25/3/2012. This allows the B-theorist to argue that semantic considerations bear only on whether two sentences are associated with the same ‘element x’ instead of the truth conditions, and this neutralizes the metaphysical relevance of those semantic considerations.²³ They concern the linguistic rules associated with the sentences, or the way in which we believe what we believe when we utter the sentence.²⁴

Regardless of all this, it can still be revealing to consider the truth conditions we can assign to bits of temporal language. In doing this, we can check first of all, whether the theory provides enough resources for the various bits of language to come out true. For example, we have had no need thus far to accept times into our ontology; but one could think that we must countenance times in order to account for sentences that explicitly use dates to state when something occurred. The question of expressive power is important for another dialectical reason: namely to account for the notions that feature as primitives in rival accounts of times, notably the tensed notions of the A-theory, and the earlier than relation of the B-theorist. This is what I will focus on. I will first discuss the truth conditions assigned to tensed sentences, and then those assigned to sentences that seem to refer to times.

In order to state the various truth conditions succinctly, it will be helpful to define notions that extract a bare temporal order from the passage facts. The following three suffice:

‘A *while* B’ \Leftrightarrow_{df} ‘ $A \circ B \leftrightarrow \neg(A \circ B)$ or $\neg(A \circ B) \leftrightarrow A \circ B$ or $A \circ B \leftrightarrow A \circ B$ ’.

‘A *before* B’ \Leftrightarrow_{df} $(A \leftrightarrow \neg A) \circ (\neg B \leftrightarrow B)$ or $(A \leftrightarrow A) \circ (\neg B \leftrightarrow B)$.

‘B *after* A’ \Leftrightarrow_{df} $(A \leftrightarrow \neg A) \circ (\neg B \leftrightarrow B)$ or $(A \leftrightarrow A) \circ (\neg B \leftrightarrow B)$.

A’s being the case *while* B is the case is understood as them co-obtaining as well as being *in* time. We do simply define ‘A *while* B’ as ‘ $A \circ B$ ’ because

²³See e.g. Mellor (1998) and Sider (2001).

²⁴This is not an ad-hoc way for B-theorists to avoid constraints on their theorizing; as the semantic framework is plausible for independent reasons. That is not say that there are nevertheless worries that the A-theorist may insist on. For further discussion, see the papers in Oaklander and Smith (1995).

the latter also state facts that are not properly said to be in time at all (such as the fact that $2+2 = 4$ or the fact that $A \leftrightarrow B$). The offered definition requires not just that the facts co-obtain, it requires that they co-obtain and are subject to passage, and only what is *in* time is subject to passage. Note that A 's being the case *before* B is the case, is understood as the passing away of A insofar as B comes to be the case *or* as the passing of A into A *insofar as* B comes to be the case. The latter disjunct captures cases where, for example, I continue to stand while you are first sitting and are then standing up. We want to say in such cases that I am standing before you are standing. The same applies to the notion of something being the case *after* something else is. Each definition ensures that the obtaining of temporal order implies that some passage occurs. It's the passing of certain facts that makes them be in time, as it were.

Using the notions of temporal order, we can assign schematic truth conditions to the utterances of various kinds of tensed sentences:

an utterance of 'it is now the case that A ' is true iff the utterance is made *while* A .

an utterance of 'it was the case that A ' is true iff A *before* the utterance is made.

an utterance of 'it will be the case that A ' is true iff A *after* the utterance is made.

These truth conditions share features with B-theoretic truth conditions in referring to particular utterances. This means that, for example, the utterance of a present tensed sentences is concerned only with facts that co-obtain with the utterance in time.

Present tensed utterances need to be sharply distinguished from atemporal utterances. We may not always be able to read from the surface structure of a given sentence whether an utterance of that sentence is present tensed or not, that is to say, whether it is concerned with what obtains with it in time, or whether is simply concern with what obtains (not necessarily in time, nor at all times). We saw an example of this in Fine's argument from

truth, in §2.2. If you ask me now whether the utterance I made yesterday of ‘KF is sitting’ is true, I will answer Yes, but I do not mean to say that the utterance is true now. Speaking without any focus on any time, it’s simply the case that the utterance is made *while* KF is sitting, and hence true. And so the assumption that Fine called *Truth stability*, the assumption that the utterance is always true, even now, is taken to be false. The argument from truth does not apply.

Atemporal utterances also need to be sharply distinguished from those that concern all times, and those that concern *some time or other*. *Just as with the definition of ‘while’, ‘before’ and ‘after’, we can introduce ‘always’ and ‘sometimes’, directly into the object language. They are naturally defined as follows:*

‘it is *always* the case that A ’ \Leftrightarrow_{df} ‘ $A \leftrightarrow A$ and $\neg(\neg A \leftrightarrow A)$ and $\neg(A \leftrightarrow \neg A)$ ’.

‘it is *sometimes* the case that A ’ \Leftrightarrow_{df} $A \leftrightarrow A$ or $\neg A \leftrightarrow A$ or $A \leftrightarrow \neg A$.

Something is always the case when it remains to be the case as time passes. Given that something is only always the case if it’s involved in passage, only what is *in* time can be said to be always the case. Similarly, something is sometime the case, when some passage of time either takes us there or takes us away from there.

These defined notions can then straightforwardly employed in the truth conditions of utterances of sentences involving them:

an utterance of ‘it is always the case that A ’ is true iff it is *always* the case that A .

an utterance of ‘it is sometime the case that A ’ is true iff it is *sometimes* the case that A .

Note that an utterance that concerns what is always the case is still taken to have a kind of restricted focus, namely its truth is sensitive only to everything that is in time.

Needless to say, there are far more complicated tensed construction in the English language. But this should suffice to show that the passage theory seems to at least match the expressive resources of a tensed theory of time in which the above tensed notions feature as primitives.

So let me now turn to a different kind of temporal language, namely those that ‘locate’ matters in time using dates. Thus far we have not accepted times in our ontology, and one may think we need times in order to account for the truth of these kinds of sentences. This is a choice point where quite a range of options can be explored. We could follow the B-theorist and simply countenance times, and various locative expressions, of something existing or having a property ‘at’ a time. Or, to name just one other option, we could follow those A-theorists who offer ersatz views of times: identifying times with sets of sentences or propositions meeting certain conditions, and allowing *them* into our ontology.²⁵ I think there is however another plausible understanding of what the truth of dated sentences consists in.

The passage theorist might want to deny that a sentence whose truth condition takes the form ‘*a* is F at *t*’ is true at all. Such sentences are false because there are no objects that are times. Accounting for sentences that have truth conditions of the form ‘*a* is F at *t*’ should however be sharply distinguished from accounting for the truth of ordinary talk in which dates are used. It’s plausible to think that such sentences do not really have truth conditions of the form ‘*a* is F at *t*’ and hence *can* be true. Take such an ordinary sentence as ‘the tree is bent at 12pm on 1/1/15’. The surface structure of the English sentence is of course ‘*a* is F at *date*’ but surface structure does not dictate the structure of the truth conditions of the sentence. The fragmentalist is under pressure not to attribute rampant systematic error to ordinary talk and so will assign truth conditions to the sentences that differ from their surface structure - as we often do.

First of all, the fragmentalist naturally reads a sentence such as ‘the tree is bent when it is 12pm on 10/01/2015’ as ‘the tree is bent *while* it is 12pm on 10/01/2015’. If it is that case that the tree is bent while it is 12pm on

²⁵For ersatzism about times, see Prior (1967: 79-82), Fine (1977/2005), Crisp (2003), Bourne (2006: Ch. 2), and Meyer (2013: Ch. 6).

10/01/2015, this is what accounts for the truth of the ordinary sentences. The second step is to explain what it is for it to be 12pm on 10/01/2015.

We use dates to locate facts in history in an objective way. There is no such thing as it being 12pm on 10/01/2015 *as such*; there is only such a thing as it being 12pm on 10/01/2015 according to a certain time standard. There are many time standards, and they work with different sets of ‘reference facts’ and it’s plausible that the locating of facts works by relating them to facts that are in a fixed set of ‘reference facts’, namely those associated with the time standards that the dates belong to.²⁶ What the relevant reference facts are depends on the time standard that the dates are associated with. In the past, we long used time standards based on the revolution of the Earth around the Sun, the various versions of solar time. For it to be 12pm on 10/01/2015 according to such a time standard is for the earth to be, say, a certain way into its *n*th revolution. In that case, ‘the tree is straight *while* it is 12pm on 10/01/2015’ is true in virtue of the fact that the tree is straight *while* the earth is at the relevant point in its *n*th revolution. Since 1972, however, the time system that is most commonly used, the Coordinated Universal Time (UTC), is based on what is known as International Atomic Time (TAI), which in turn is based on the states of about 270 atomic clocks in different locations (*Encyclopædia Britannica* 2015). For it to be 12pm on 10/01/2015 in UTC, for example, is ultimately for the collection of 270 atomic clocks to be in such and such a state. In general then, the following are natural truth conditions for sentences that locate events at certain dates:

an utterance of ‘*a* is *F* at *d* in time standard *TS*’ is true iff *a* is *F* *while* the reference facts obtain that are associated with *d* by time standard

²⁶Another option is to understand a sentence of the form ‘it is the 1st of December 2014’ as structurally similar to such a sentence as ‘it is raining’, whose structure doesn’t have to be understood as the attribution of a property to an object (see Strawson 1959: §6.2, §7 and Turner 2011), rather we understand the obtaining of these facts as the world’s being of a certain age (cf. Taylor 1960). Such facts may be attractive to anyone who wants to make room for a passage of time without any change in the objects that exist in time because the world could then age without there being any qualitative change of objects in time (see arguments in Shoemaker 1969 and Le Poidevin 2010; for a reply to these, see Warmbrod 2004). The understanding of the date facts proposed in the body of the text seems much more natural to me, and I’m not myself moved to make room for passage without change.

TS.

The sentence ‘the tree is straight *while* it is 12pm on 10/01/2015 in UTC’ is true in virtue of the fact that the tree is straight *while* the 270 atomic clocks are in the relevant state. To situate facts in history is to situate the obtaining of the fact with respect to the obtaining of some fact in a designated class of facts, namely those that function as the objective standard via which we can ‘locate’ facts in history. And so there is no need to introduce times into our ontology, at least not to account for the truth of ordinary sentences that use dates.

3.6 Appendix: translations into the A-theory

I have used semantics to elucidate and precisify the inferential role of the notion of passage and more general, the metaphysical picture proposed by the passage theory. Another way of elucidating a theory is by offering a translation into a more well-known theory; this is for example how Jaśkowski introduces his non-adjunctive logic (1948/1969). One might think that there is an easy translation of the passage-theory into claims of a standard A-theoretic view, which an adherent of that view can use to gain a hold on the passage theory. I have not been able to find an easy translation however, and I think it may be helpful to see why certain tempting translations fail.

Let ‘S’ (‘it is sometime the case that...’) be the sometimes-operator as it’s standardly defined in tensed framework: $SA \Leftrightarrow_{df} PA \vee A \vee FA$. Let i be a translation function that takes sentences of the passage theory (the P-theory, to follow the tradition of beautiful labels in the philosophy of time) into sentences of a standard A-theory as follows:

$$\begin{aligned}(p)^i &= Sp \\ (A \wedge B)^i &= S(A)^i \wedge S(B)^i \\ (\neg A)^i &= \neg S(A)^i \\ (A \circ B)^i &= S((A)^i \wedge (B)^i) \\ (A \leftrightarrow B)^i &= S((A)^i \wedge F(B)^i)\end{aligned}$$

The translation may seem quite intuitive. To take the case of atomic sentences, it's indeed the case that p is true if and only if $S p$ in a standard A-theory. Similarly, $p \wedge q$ is true when $S(S p \wedge S q)$ is true in the A-theory. So the translation seems simple and promising, and suggestive of the following translation claim:

Translation claim 1: A is true according to the P-theory iff $(A)^i$ is true according to the A-theory.

The claim is however false if time takes a certain course. Consider the following time series (with columns representing times, and underneath the sentences true at the relevant times):

t_1	t_2
p	$\neg p$
$\neg q$	q

Here: $p \circ q$ is false according to the passage theory, but it's i -translation, $S(S p \wedge S q)$ is true according to the A-theory.

A different translation function j and a different translation claim:

$$\begin{aligned}
 (p)^j &= p \\
 (A \wedge B)^j &= S(A)^j \wedge S(B)^j \\
 (\neg A)^j &= \neg(A)^j \\
 (A \circ B)^j &= ((A)^j \wedge (B)^j) \\
 (A \leftrightarrow B)^j &= ((A)^j \wedge F(B)^j)
 \end{aligned}$$

Translation claim 2: A is true according to the P-theory iff $S(A)^j$ is true according to the A-theory. (Note the use of the S-operator on the right hand side here).

This time, we no longer have the problem of the previous translation. Given the time series:

t_1	t_2
p	$\neg p$
$\neg q$	q

$p \circ q$ is false according to the passage theory. It's j -translation is just $p \wedge q$ and so according to the translation claim $S(p \wedge q)$ must be false according to the A-theory, which is indeed the case.

And yet the translation claim is false again. The following is false in the passage theory: $\neg p$. It's j -translation is $\neg p$, and so according to the translation claim $S\neg p$ should be false but it's true according to the A-theory.

Consider yet another translation, k , which is the same as before except for the translation of negation, for which we use the always operator A ('it is always the case that...') defined as follows: $SA \Leftrightarrow_{df} PA \wedge A \wedge FA$:

$$\begin{aligned} (p)^k &= p \\ (A \wedge B)^k &= S(A)^k \wedge S(B)^k \\ (\neg A)^k &= A\neg(A)^k \\ (A \circ B)^k &= (A)^k \wedge (B)^k \\ (A \leftrightarrow B)^k &= (A)^k \wedge F(B)^k \end{aligned}$$

Translation claim 3: A is true according to the passage theory iff $S(A)^k$ is true according to the A-theory.

This solves our previous problem: $\neg p$ is false in the above time series according to the P-theory, it's k -translation is $A\neg p$, and so according to the translation claim $S(A\neg p)$ should be false, which it is.

But it is not hard to see it fails for other sentences. For example, $p \circ \neg q$ is true according to the P-theory in the above two time series. It's k -translation is: $(p \wedge A\neg q)$, and so $S(p \wedge A\neg q)$ has to be true according to the A-theory of the above time series, but it isn't.

And so it goes: whatever translation I have tried to come with, it was easy to find a possible course of time that invalidated the translation claims. What does this show? It shows at the very least that one should not understand the fragmentalist's descriptions according to these translations, which I think can be tempting if one is used to think of time in tensed terms. Perhaps someone can come up with a successful translation; but, if it turns out to be very complicated or gruesome, it's not clear whether it will be of much help in elucidating the passage theory of time.

Chapter 4

Phenomenality in a Fragmented World

When you experience the world, there is something it's like to undergo this experience, that is, certain phenomenal facts obtain. How should we think of these phenomenal facts? This chapter explores a no-subject view of phenomenal facts. When you interact with the world, it comes to appear in certain ways, it manifests itself as including various objects that have various qualities. These appearance facts are not relativized to a subject: when you causally interact with the world, certain things come to appear to be the case simpliciter.

The main objective of this chapter is to formulate a no-subject view of phenomenal facts that is more plausible than existent formulations and, in particular, to show that a plausible no-subject view requires fragmentalism in order to withstand an otherwise decisive objection. The chapter consists of the following sections. Section 1 discusses three phenomenological considerations in favour of a no-subject view. Section 2 sketches an account of the subject-less phenomenal facts. Section 3 discusses the main problem that, I think, any no-subject view runs into. Section 4 shows how this problem is solved when we admit fragmentation in the phenomenal facts. And finally, section 5 discusses the possibility of fragmentation in the experiences of a single subject. Along the way, I will point out advantages of the no-subject

view over the two more standard views of phenomenal facts: the qualia view and representationalism but, as before, it's not my intention to offer conclusive objections to them. Though Fine also shows how fragmentalism can be employed to develop a no-subject view of experiences (see Fine 2005b: §12), I will not discuss Fine's account in any detail; I will merely point out the main point of divergence between his approach and the approach favoured here.

4.1 Three considerations in favour

Various philosophers have been drawn to some version of the no-subject view - such as Hume, Schlick, Wittgenstein and, more recently, Hare, and Johnston.¹ And I think rightly so. Here I will discuss three phenomenological considerations in favour of the no-subject view of phenomenal facts (also known as a no-ownership or no-self view).

Hume expressed a simple and direct phenomenological consideration in favour of the thought that experiences are not in any essential sense *had* by a subject or self: we simply do not find such a subject or self when we reflect on our ongoing experience of the world. Hume expressed it in these often quoted words:

There are some philosophers, who imagine we are every moment intimately conscious of what we call our SELF; that we feel its existence and its continuance in existence; and are certain, beyond the evidence of a demonstration, both of its perfect identity and simplicity. To attempt a farther proof of this were to weaken its evidence; since no proof can be deriv'd from any fact, of which we are so intimately conscious; nor is there any thing, of which we can be certain, if we doubt of this.

For my part, when I enter most intimately into what I call myself, I always stumble on some particular perception or other, of

¹See Hume (1739/1975), Schlick (1936), G. E. Moore (1955) and Kripke (1982) on Wittgenstein, Johnston (2007) and Hare (2009).

heat or cold, light or shade, love or hatred, pain or pleasure. I never can catch myself at any time without a perception, and never can observe any thing but the perception. [...] If any one, upon serious and unprejudic'd reflection thinks he has a different notion of himself, I must confess I can reason no longer with him. All I can allow him is, that he may be in the right as well as I, and that we are essentially different in this particular. He may, perhaps, perceive something simple and continu'd, which he calls himself; tho' I am certain there is no such principle in me. (Hume 1739/1975: Bk.1 Ch.4 §6).

I agree with Hume. When I attend to my experience, there is the content of my experience, and there is the fact that the content of my experience appears to obtain or manifests itself as obtaining. I do not encounter the experience as qualifying or determining a certain thing, a subject.

I do not think that Hume's consideration should be taken to concern personal identity.² The no-subject view that is of concern here is not the view that *I* am not a thing or do not exist, it's a view concerning phenomenal facts, namely that they are not relativized or specific or tokened to a subject. I will assume that, at any time, I'm a physical organism that causally interacts with the world and nothing other than that. Thus, when I use 'I', it always and only picks out this physical organism and when I consider the experience that I find when I introspect, it's the physical organism, ML, that is 'finding' these experiences. We can remain silent on issues concerning diachronic identity - about what makes me be this organism at this time and that organism at that later time. We will solely focus on the nature of phenomenal facts that obtain *at a time*.

More fully then, I take the Humean consideration to be this: an experience does not appear as being instantiated by a physical organism, nor by some entity, a subject, distinct from that physical organism. If I attend to that which is sometimes supposed to be the bearer of experiences, I come up

²I do not mean to deny that Hume might have taken it to bear on questions concerning personal identity, I only mean that it will not here be taken to imply that I am not anything.

empty.

This is not to deny that there is some sense in which the experiences strike me as being ‘mine’. But the sense of mine-ness consists in the fact that the contents of the experiences are centred around this physical organism (there is an experience with *my* two hands sticking out and the outer borders of *my* glasses, etc.), and there is an integration of these experiences with each other, with the availability of memories and intentions that all concern me (i.e. the physical organism that I am), and with their availability to higher order reflective acts which, when they occur, are themselves integrated with my experience of the world (cf. Johnston 2007: 258). Above all else, however, the sense of mine-ness is based in the dependence of the experiences on the organism that I am: when I am affected by the environment, in particular, when my perceptual apparatus causally interacts with the environment, the experiences I find will change. But any experiencing that goes on is not something that seems to consist in a property of some thing, a subject. We will return to this below.

A second phenomenological consideration in favour of a no-subject view concerns the purely qualitative nature of the phenomenal facts, their apparent transferability. Take my experience of a red apple. There is this experience, and then there is the fact that it is ‘had by me’ in a sense that seems extrinsic to the experience; instead of its dependence on me (i.e. the physical organism that I am) and its integration with my other experiences (i.e. those experiences that depend on me), there could have been a dependence on a different organism and an integration with different experiences.³

³Note this talk of ‘integration with *my* experiences’. Do we not illegitimately need to appeal to the sense of mineness that is denied by the no-subject view? Strawson argues that the no-subject view is incoherent because it needs to make such an appeal:

It is not coherent, in that one who holds it is forced to make use of that sense of possession of which he denies the existence, in presenting his case for the denial. When he tries to state the contingent fact, which he thinks gives rise to the illusion of the ‘ego’, he has to state it in some such form as ‘All *my* experiences are had₁ by (i.e. uniquely dependent on the state of) body B’. For any attempt to eliminate the ‘my’, or any expression with a similar possessive force, would yield something that was not a contingent fact at all. (Strawson 1959: 96-97).

This objection fails. By ‘all *my* experiences are had₁ by (i.e. uniquely dependent on the

If one rejects the no-subject doctrine, the experience of the red apple consists in the fact *that I experience the apple*. Next to the object of experience (the apple's being red) and its being experienced, its being experienced *by me* is considered to be part of what makes it the experience that it is. The experience is essentially mine, and could not be anyone else's. Someone else could have a distinct experience with the same content but they could not have *this* experience. This strikes me again as phenomenologically inadequate: when I reflect on the very experiences that I have, the experiences seem purely qualitative, I *can* imagine that someone else undergoes these very experiences, that someone else is in my seat as it were.

Vice versa, I can imagine being in the seat of someone else. Consider Lewis's thought about poor Fred:

Here am I, there goes poor Fred; there but for the grace of God go I; how lucky I am to be me, not him. Where there is luck there must be contingency. I am contemplating the possibility of my being poor Fred, and rejoicing that it is unrealized. (Lewis 1999: 397).

Lewis is contemplating a certain kind of possibility, described as 'the possibility of being poor Fred'. Lewis finds himself lucky that he isn't poor Fred and claims that where there is luck there must be contingency. I think he is right. But where there is contingency there is a difference of fact: we have contingency when things that are this way could have been that way instead.⁴ What are the ways things could have been in this case? It's not

state of) body B' we can perfectly well mean 'all the experience that depend on body B depend on body B'. Contrary to what Strawson maintains, this *is* a contingent fact: all the experiences that (actually) depend on body B could depend on a body C (distinct from body B). It's contingent in the very sense in which 'my neighbour is my neighbour' is contingent: the person who is my neighbour could have lived somewhere else. So-called 2D frameworks (see, e.g., Davies and Humberstone 1980) acknowledge a sense of necessity in which the claim could be deemed necessary, but that is not the sense of necessity that is in question here.

⁴Lewis himself backed down from this and developed the centred worlds framework (see his 1979) especially so that there can be different possibilities without differences of fact. The possibility of adopting such a framework does not take away from there being a prima facie reason to think that we are confronted by differences of facts.

Lewis's literally 'being identical to someone else': no thing can be identical to a thing distinct from it. Identity is necessary. And, in any case, the possibility seems to involve something unpleasant, something for which it makes sense to rejoice that it isn't realized, and bare identity facts aren't like that. No, when I imagine 'being poor Fred', I imagine having *the very* experiences that Fred has. I do not just imagine experiences with the content of Fred's experiences; this is the possibility of 'being *like* poor Fred'. When I imagine '*being* poor Fred', I imagine the world as it is experienced by Fred, including the fact that their being so experienced depends on an organism in the experienced scene, an organism that is Fred.⁵ This contingency in who has a given experience is not easily captured if one rejects the no-subject view. One might insist that when I imagine the world as it is experienced by Fred, I implicitly imagine the world being so experienced by *me*, that *I* experience the world this way, but since I am not Fred, it seems I would thereby be unable to truly imagine *being Fred*, and it seems that this *is* possible. The experience seems purely qualitative and transferable, and not specific to a particular individual.

A third phenomenological consideration in favour of the no-subject view, effectively strengthens the previous two considerations. This consideration concerns the much-discussed transparency of experience. The transparency of experience is best appreciated in contrast with the influential qualia view of phenomenal facts. According to the qualia view, the phenomenal character of an experience is captured in terms of properties *of* a (token) experience (or of the subject's 'state' of experiencing something).⁶ These properties are so-called phenomenal properties, or qualia, and are typically taken to be introspectively accessible, nonrepresentational properties of those experiences

⁵Shoemaker remarks: 'if I imagine the battle of Cannae as it might have been experienced by Hannibal, I do not thereby imagine being Hannibal' (1994:17). That is right. But I *can* also imagine the battle of Cannae as it might have been experienced by Hannibal *and* also being Hannibal by imagining that the world's being so experienced depends on the physical organism (in the imagined scene) that is Hannibal. This is a natural way of understanding Nagel's idea that there is something it's like to *be* a certain organism (e.g. 1979:166).

⁶Representatives include: Block (2003), Levine (2001), Peacocke (1983), Chalmers (1996), as well as Jackson (1982) and Shoemaker (1982) though the last two authors turned to defending representationalism in later papers.

(i.e. not the property of *being an experience that represents that A*). An ancestor of this is the view that there are sense-data with their own properties.

The qualia view is in tension with the apparent transparency of experience: what we are aware of when experience things are precisely those properties that objects in the world appear to have and not any properties of the experiences themselves.⁷ Harman describes this phenomenological fact as follows:

When Eloise sees a tree before her, the colors she experiences are all experienced as features of the tree and its surroundings. None of them are experienced as intrinsic features of her experience. Nor does she experience any features of anything as intrinsic features of her experiences.(Harman 1990: 667).

Similarly, Shoemaker:

I am looking at a book with a shiny red cover. The property I experience its surface as having, when I see it to be red, is one that I can only conceive of as belonging to things that are spatially extended. How could *that* property belong to an experience or sensation? (1994: 25).

Experiences are not like particular images of the outside world that have their own properties. We don't represent what properties are instantiated in the world via our awareness of certain properties of some mental picture of the world; our experience is of the very properties that appear to be had by things out there and that are the sorts of properties whose instantiation we represent as obtaining.

Note that the transparency of experience leaves little room for the idea that experiences are 'tokened' as it were by being had by a particular subject. If experiences are not things with self-standing properties, then they *ipso facto* cannot have the properties that can make for a token distinction

⁷Transparency features in arguments for a range of different conclusions: for a defence of functionalism in Harman (1990), against qualia-realism in Tye (2014), for naive realism in Kennedy (2009), and for Russelian contents of experience in Speaks (2009).

between them. Say that we both look at Shoemaker's book with the shiny red cover. We attend to the experience and see right through it, as it were, attending to a kind of manifestation of the content of the experience: the book with the shiny red cover. But, given that the content of our experiences is the same, we attend to the same thing, the same book manifesting the same colour. As the content of the two would-be token experiences is the same, the token-ess needs to lie in properties of the vehicle of the content, but we never encounter such a vehicle, we encounter just a manifestation of the content.⁸

Note secondly that the transparency of experience fits the claim that there appears no subject to which the experiences attach. Again, when I attend to the experience of the shiny red book, I find the shiny red book, appearing to be red. If I were to also find a subject, I would have to be experiencing it not as standing in some relation to the experience, but as standing in some relation to the red book. And that is not, on this occasion, what I find. Whenever I *do* find that things stand in some relation to me, such as when I find that the book lies in front of me for example, this just means that the content of my experience is of the book lying in front of me, not that the *experience* of the book stands in some relation to me. Given transparency, there does not seem any experience for a subject to be related to, and hence no room for any subject *having* the experience in an essential sense.

These then are three phenomenological considerations in favour of a no-subject view: we simply do not find a self or subject having the experience, the experiences seem purely qualitative and transferable to others, and experiences do not seem to be things having properties or standing in relations to anything. Whatever reason there is for positing a subject over and above the body or for attributing experiences to the body, must be a theoretical reason, and not a phenomenological one.

⁸Fine appears to accept such token experiences in his discussion of first-personalism, writing that 'It might be supposed, for example, that there are token experiences that only I can have [...] Each of my token experiences would then give rise to an egocentric fact, the fact of the experience being experienced' (Fine 2005b:318). Just as positing token events seemed to be in tension with tense realism (§2.3), positing token experiences seems to me to be in tension with the motivations that drive a no-subject view.

4.2 Phenomenal facts as appearance facts

If experiences are not self-standing ‘things’ with properties, the talk *about* experiences that we occasionally resorted to above cannot be quite right. Similarly, the talk of a ‘manifestation’ of the content of experience cannot really be talk of a thing that is the manifestation of an experienced content. Given the three phenomenological considerations, it also cannot consist in an attitude *of* a subject. So how are we to understand the subject-less phenomenal facts of the no-subject view? The above three considerations naturally lead to a particular understanding of the subject-less experiences. Here I want to make this understanding explicit. Along the way, I will draw some further distinctions that I think are crucial if the no-subject view is to be plausible.

When we experience things, it should be uncontroversial that those things come to appear in various ways, that when we see things they look certain ways, when we hear things they sound certain ways, when we touch things they feel certain ways, etc. It’s a natural thought that the way things appear when a subject interacts with the world and the fact that they do so at all is ‘what it is like’ for a subject to undergo that experience. When a tulip looks red when I see it, it’s appearing red *just is* what it’s like to undergo the experience of the red tulip. Call this the *phenomenality = appearance principle*:

Phenomenality = appearance: The phenomenal character of an experience consists in the content of that experience coming to *appear*.

Furthermore, the ways things appear when we interact with the world are naturally taken to be ways things appear *to be*, i.e. properties things appear to have. Following Shoemaker, call this the *ways = properties principle* (Shoemaker 2006: 461):

Ways = properties: The way things appear when interacting with the world consists in some property or relation that things appear to instantiate when interacting with the world.⁹

⁹The qualia view rejects the *ways = properties principle*, i.e. rejects that how things

This principle squares with the truism that things can look to be different from the way they really are, that is, things can fail to be the way they appear to be. The question of the veridicality of the way things appear is a question of whether things really have the properties they appear to have. Without the *ways = properties principle*, appearances would be incomparable to the way things really are; but they are comparable.

The notion of appearance is most naturally understood as a type of modality, that is, in terms of *it appearing to be the case that A*. Formally, the phenomenal facts are best described using a sentential operator ‘ \mathcal{A} ’ with which to form sentences such as ‘ \mathcal{A} (the tulip is red)’, saying that *it appears to be the case that the tulip is red*.¹⁰

What, then, is meant with something ‘appearing to obtain’? I will take appearance to be the primitive notion of the no-subject view. By way of homing on the relevant sense of the notion though, we might say that when you undergo experience, a certain qualitatively rich scene manifests itself as obtaining, discloses itself. That is, I understand ‘appearing’ in the way that Johnston understands ‘presence’:

Consider presence, the variety of ways in which real or ostensible items, be they objects, qualities or whatever, disclose some aspect of their nature. Perhaps the best way to bring presence into view is to begin with perception. When one sees one’s dogs running in the front yard, the whole content of the perceptual experience is of the dogs and their running being present in a certain way, a way that discloses something of the nature of the dogs and their running. THERE the dogs are, immediately available as objects of attention and demonstration, and as topics of one’s further thought and talk. (Johnston 2007: 233).

appear in experience consists in properties that the experienced objects out there appear to instantiate. They distinguish between ‘how things appear’ and ‘how things appear *to be*’ and claim that how things appear (or seem, or look) when we experience something consists in the instantiation of properties by those experiences.

¹⁰This proposal should therefore be distinguished from the ‘theory of appearing’ as it is defended in Alston (1999) and Langsam (1997). The core primitive of this theory is the relation of an object *o* appearing to be *F* to a subject *s*; whereas the core primitive of the proposed account is that of some fact appearing to be the case simpliciter.

When something appears to be the case, there is a presence of the involved objects and properties.

It is an obvious characteristic of what appears to be the case that it may be misleading. What appears to be the case doesn't have to be the case; it doesn't even have to involve anything that actually exists. This means that appearance statements aren't factive, nor admit of an analogue of the Barcan Formula:

$$\mathcal{A}A \not\equiv A$$

$$\mathcal{A}\exists x(Fx) \not\equiv \exists x\mathcal{A}(Fx)$$

If it appears to be the case that a dragon flies in front of my window, we clearly shouldn't infer from this that there is a dragon flying in front of the window, nor that there is something in the world that appears to be a dragon flying in front of my window, as there doesn't need to be anything in the world that I experience *as* the dragon.

Beside the possible non-veridicality of appearances, there are less obvious elucidatory questions that need to be settled. A first question concerns the content of appearances. I will assume that what appears to be the case does not just concern colour patches in certain configurations (in the case of visual appearances) but also concerns identified particulars. There can be an appearance of *this book* being red. And similarly for the ways things feel, sound, smell, and taste; there are particular things that smell or sound or taste certain ways. I will also assume that things appear in ways that discriminate much more finely than one is able to do conceptually in judgements and mere beliefs and in ways that go beyond what one can memorize or process cognitively at any point in time. We may not have words or concepts for every single way things appear to be. These assumptions are however not essential to the no-subject view as such, and could be dropped.¹¹

A second important question is whether the relevant notion of appearance should be understood as being essentially first-personal, or not. Adherents

¹¹For further discussion, see Tye (2006) and Byrne (2001:202). Tye (2002) offers a defence that the phenomenal facts consist in certain facts appearing to be the case (to a subject - in his view).

of a no-subject view typically distinguish between such sentences as ‘I am sitting’ or ‘I have a matchbox in my hands’ and sentences such as ‘I see a red book’ or ‘I am in pain’ (G. E. Moore 1955:13-14 reports that Wittgenstein drew this distinction during a certain period; see also Anscombe 1975:61). ‘I am sitting’ or ‘I have a matchbox in my hands’ are straightforwardly understood as attributing properties to a subject: the physical organism that I am has the property of sitting, and has a matchbox in his hands. ‘I see a red book’ or ‘I am in pain’ are not understood in this way; they do not straightforwardly attribute properties to an organism. Many adherents of the no-subject view take a sentence such as ‘I am in pain’ to be equivalent to the subject-less ‘it is paining’, where the latter is taken to be implicitly first-personal: it is the case simpliciter that it is paining if and only if it is the case from *my* experiential perspective that it is paining.¹² This is analogous to the tense-logical convention that it is the case simpliciter that it rains if and only if it is the case *right now* that it rains.¹³ Following Fine, we can call this species of the no-subject view ‘first-personalism’.

An adherent of the no-subject view is in no way forced to adopt a first-personalist understanding of the phenomenal facts however. Indeed, it seems to me plausible to think that, if I say, ‘I am in pain’, this may - in appropriate contexts - be equivalent to ‘it appears to be the case that there is pain and this appearance depends on ML’. Similarly, ‘I see a red book’ can be understood as saying that ‘the book appears to be red and this appearance depends on ML’. In general, such sentences may be understood as expressing that things appear a certain way and that those appearances depend on a certain subject. The appearances themselves are not here understood as being first-personal in some implicit way - they are truly anonymous, and the use of ‘I’ is taken straightforwardly as referring to the subject using it. We do not distinguish between two senses of ‘I’: one where it picks out a subject (as in ‘I am sitting’) and a sense where it disappears as it were (as when ‘I am in

¹²This view is inspired by Lichtenberg’s famous critical remark concerning Descartes’ cogito: ‘We should say *it thinks*, just as we say *it lightens*. To say cogito is already to say too much as soon as we translate it I think.’ Lichtenberg (1804/1990: K18).

¹³For an exploration of the analogy, see Prior (1968/2003a) and Prior (1977/2003), Fine (2005b: §6 and §12), and Hare (2009: Ch.3).

pain' is read as 'it is paining'). In contrast, 'I' always and uniformly picks the one who uses it. From here on, I will refer to this as the no-subject view and distinguish it from first-personalism. I will only explore the no-subject view.¹⁴

A third important question is how appearances bear on the information that a subject has about the world. It will be helpful to address this question in the context of representationalism, the main rival of qualia views. According to representationalism, the qualitative character of an experience is determined by the representational content of that experience.¹⁵ Representationalism is minimally committed to a supervenience claim: there can be no difference in the qualitative character of two experiences without a difference in what they represent to be the case.¹⁶

Contra representationalism, however, an adherent of the no-subject view should distinguish sharply between representational facts and appearance facts. We describe a subject as representing something in order to capture that she interacts with the world informed by some body of information about her surroundings. Representational states thus qualify a certain item in the world, namely a physical organism. The appearance facts, in contrast, do not capture the informational interaction of a subject with the world, but what things appear like when this subject experiences things. An adherent of the no-subject view should accept that representational states *are* instantiated by subjects (or are attitudes had by subjects). We ascribe representational states to subjects on the basis of how they seem disposed to act: when a subject is disposed to act (and talk) as though it is the case that *A*, then we have reason to ascribe a representational state to the subject according to

¹⁴First-personalism and the no-subject view are typically conflated, and so it is difficult to find unequivocal adherents of the no-subject view as understood here. It is plausible to think that Hume had a no-subject view as understood here, never mentioning any special use of 'I' (as far as I know) and admitting the bundles of experiences that - we say - depend on different subjects.

¹⁵Proponents include Harman (1990), Armstrong (1999), Tye (1995), Dretske (2003), Byrne (2001), and Jackson (2004).

¹⁶I will leave it open how we understand the representational state. One might require a special experiential propositional attitude: *s* experiences that *A* (see Byrne 2009). Or one might require that the relevant representational states need to satisfy certain conditions for their content to count as phenomenal (see Tye 1995).

which it is the case that A .¹⁷ For example, if someone says that a tulip is red, groups it together with other red things, and seems disposed to continue to do so in a reliable fashion, then it is plausible that the subject represents the world as being one in which the tulip is red. This is part of the information that this organism possesses and acts on.

Given the distinction between representations and appearances, there is a question concerning the relation between them. Is it the case that a subject represents that A if and only if it appears to be the case that A and it appearing to be case that A depends on that subject?

There are reasons to deny this. We should allow that the contents of appearances and representations can come apart. Consider the case of an inverted spectrum due to Shoemaker (1982). This thought experiment is normally invoked to show that representationalism makes the wrong predictions about phenomenal character but shows more generally that the content of appearances should not be identified with the representations of the subjects on which these appearances depend. The case involves someone, call him Nonvert, who experiences red things exactly the way we experience them and someone else, Invert, who experiences red things exactly the way we experience green things. Both have been like this since birth, and make the same discriminations between things. And both are part of the same linguistic community, describing the same things as red, green, etc.

If Invert and Nonvert both make the same colour discriminations when interacting with the same coloured things, and both describe for example the same tomato as red, then both Invert and Nonvert represent this tomato as being red. Externalism about representational content leads us to say that their experience must represent the tomato in the same way. And so, if the way things are represented by experiences is all there is to the phenomenal character of experiences (as representationalism maintains), externalism forces us to claim that there simply is no difference in the phenomenal char-

¹⁷This weak claim about reasons for attributing beliefs is not to be confused with a claim about the nature of those beliefs. The weak claim should be compatible with dispositional accounts of belief (as in e.g. Braithwaite 1932, Marcus 1990 and Schwitzgebel 2002), with functionalist accounts (as in e.g. Stalnaker 1984 and Armstrong 1973), and with interpretationist accounts (as in e.g. Dennett 1991 and Davidson 2001).

acter of Nonvert and Invert, which seems wrong. There is a phenomenal difference between Nonvert's and Invert's experience of the tomato - given that Nonvert experiences it as we experience red things and Nonvert experiences it as we experience green things.¹⁸

There are therefore independent reasons to believe in a distinction between phenomenality and representations, and that the representations of a subject may have a different content from the appearances that depend on that subject. Both Nonvert and Invert represent the tomato as being red, given the environment that they interact with. But when Invert experiences the tomato it appears to be green, and when Nonvert experiences the tomato it appears to be red. The difference between them lies in what things appear to be like when they interact with coloured objects and not in what things are represented to be like. When we consider Invert's experience, things do not appear as they are represented to be by him whereas in the case of Nonvert they do. More generally, there seem to be good reasons to think that, whereas representations of a subject are externally determined, what appears to be the case when a subject interacts with the world is internally determined.

We could consider further characteristics of appearances. There is however a major issue that we need to discuss first, and whose solution will constrain how we think of appearance. To summarize the main characteristics of the subject-less experiences or, more precisely, of the phenomenal facts concerning experiences: they consist in subject-less appearance facts, that are non-factive, that concern particulars having various (sensible) properties, are not inherently first-personal, and that should be distinguished from representations, which *do* inhere in subjects (i.e. physical organisms).

¹⁸Whereas in the standard case of the inverted spectrum, representationalism predicts no phenomenal difference where there intuitively should be one, there are also cases where representationalism predicts that there is a phenomenal difference where there shouldn't be one - the inverted earth cases. See Block (1990).

4.3 The problem of conflicting perspectives

We have seen three phenomenological considerations in favour of the no-subject view; and we proposed a more theoretically precise picture that is in line with these considerations. There is however also a phenomenological consideration that spells trouble for the no-subject view of phenomenality. It's very simple. When I reflect on my experience, I find that various things appear in various ways and that everything that appears to be case depends on the physical organism that I am. I do *not* find anything else appearing to be the case, in particular, I do not find the appearances that you find when you reflect on your experience. Each and every appearance depends on how things stand with me, indeed, this allows me to identify myself via the appearances: I am the one on which *the* appearances depend. This requires that the appearances that depend on you do not obtain. So, although this sounds megalomaniac, there is something phenomenologically right about the thought that there is something it's like when I undergo experience and nothing it's like when anyone else does so.

Each one of us encounters a tension between a subjective and objective view on the world. The opposition between these different views of the world is well-known from the writings of Nagel. He made the following insightful observation regarding it:

The opposition looks like a stalemate because each of the points of view claims dominance over the other, by virtue of inclusion. The impersonal standpoint takes in a world that includes the individual and his personal views. The personal standpoint, on the other hand, regards the deliverances of the impersonal reflection as only a part of any individual's total view of the world. (Nagel 1979b: 205-206).

Each view adds facts to the world according to the other view. The subjective view adds to the content of the objective view that there is something it's like when one subject undergoes experience. The objective view adds to the world according to the subjective view that things are like that *for any other subject*.

In the context of a no-subject view, the move to an objective view leads to incoherence. To see this, we can reformulate Nagel's tension into a paradox, or rather a schema for a paradox that each will have to formulate for him or herself. In my own case, the paradox revolves around the following two principles:

Egocentricity: Something appears to be the case when ML interacts with the world and nothing appears to be the case when any other subject interacts with the world.¹⁹

Egalitarianism: Whatever general principles hold of the phenomenal facts involving one subject hold of the phenomenal facts involving any other subject that undergoes experience (cf. Hellie (2013)).

The paradox arising from this is straightforward.²⁰ According to *Egocentricity*, it's only the case that something appears to be the case when ML interacts with the world and nothing appears to be the case when anyone else does so. Or, to put it in more general terms, there's only something it's like when ML interacts with the world and nothing it's like when anyone else does so. The world is disclosed in a way that depends on exactly one physical organism - me. Now given that *Egocentricity* is a general principle concerning the phenomenal facts involving the experiences of ML, *Egalitarianism* tells us that similar facts should hold when we substitute for ML any other arbitrary subject that undergoes experience. So take someone else, TN. If we substitute TN for ML in *Egocentricity*, this means that something appears to be the case when TN interacts with the world and nothing appears to be the case when any other subject interacts with the world, including ML. But this means that something appears to be the case and nothing appears to be the case when ML undergoes experience, and that something appears to be the case and nothing appears to be the case when TN interacts with the world, and so on for any other subject. Contradictions arise all over.

¹⁹For further discussion supporting the thought expressed in the *Egocentricity* principle, see Hare (2009), Hellie (2013), and Merlo (Dissertation).

²⁰Note the similarity with McTaggart's argument as it is understood in Fine (2005b), revolving around a conflict between reasons to privilege one subject and thinking we are all on a par.

Egalitarianism describes an aspect of what we think the world is objectively like: whatever general principles hold of the phenomenal facts involving one subject hold of the phenomenal facts involving any other subject that undergoes experience. It captures the truism that I'm just one amongst many subjects and that everyone's experience is on an exact par. This seems undeniable to me, after all, we have exactly the same kind of brains operating in the same world under the same physical laws. If there were such general principles, true of one subject but not of others, the world order would somehow be oriented towards one subject. That cannot be right.

The problem does not arise for the standard view because of the relativization to subjects: there is something *to ML* when ML interacts with the world and there is nothing it's like *to ML* when TN interacts with the world. And correspondingly for TN: there is something it is like *to TN* when TN interacts with the world and nothing it is like *to TN* when ML interacts with the world. No contradictions arise. But this solution is not available on a no-subject view of the phenomenal facts. The relativization is exactly what the no-subject view denies, and goes against the phenomenological considerations we saw in §4.1. In fact, it also goes against the phenomenological facts that lies at the heart of this problem because, with the relativization, we do not explain how the phenomenal facts allow us to single ourselves out: I cannot say 'I am the one that has *the* experiences', I can only say 'I am the one who has *my* experiences' but for this to single me out, I already need to know what experiences are *mine* and which ones *aren't*. And yet, if I consider the world as it appears to be, then I can single myself out.

An adherent of the no-subject view is thus naturally led to some form of solipsism: something appears to be the case when *I* interact with the world; and nothing appears to be the case when *you* interact with the world.²¹ And this, in turn, fits neatly with the first-personalist understanding of the subject-less phenomenal facts, according to which 'there is pain' and 'I am in pain' are equivalent. Borrowing from the A-theoretic framework of tense-

²¹Alternatively, one might think we have to live with two incommensurable ways of understanding the world, a first person way and a third person way. This was Nagel's response to the problem (in his 1986; see also Harman 2007).

operators, one might introduce an operator ‘from someone else’s point of view it is the case that...’ that allows one to capture something of the experience of others. If I feel no pain but you do, I will deny that there is pain but allow that *from someone’s point of view*(it’s paining) (see Hare (2009), Merlo (Dissertation) and Hellie (2011)). In this case, you will disagree and say instead that there is pain but allow that *from someone’s point of view*(it’s paining). So you and I come to offer disagreeing descriptions of the world. The awkward question is of course: who could be right? Really just one of us? It’s here that the fragmentalist framework allows us to offer a more plausible understanding of things.

4.4 Fragmentation in the appearance facts

Each of us is able to switch attention to phenomenal facts that allow us to identify a subject in the world because they *all* depend on that subject. Let us simplify things for the moment and imagine, instead of the phenomenal facts, a big neon arrow floating above a single subject that it points out. Now consider again two toy-models of the world. The first toy-model is a gigantic picture where the neon arrow either points to ML, or to TN, but cannot point to both ML and TN. Pointing to ML and pointing to TN are incompatible facts, and this model leaves no room for such facts both to obtain. We have to choose who the arrow points at, or we have to leave the arrow out of the picture altogether, or we have to give everyone their own arrow (in which case I cannot identify myself as the one who *the* arrow points at).

Now consider a fragmentalist toy-model of the kind we have seen before (§1.1). It consists not in a single gigantic picture but in a collection of such gigantic pictures, where there is one picture in which the big neon arrow hovers above ML and one where it hovers above TN. The arrow points at ML *insofar as* it doesn’t point at TN, and it points at TN *insofar as* it doesn’t point at ML. Each picture shows a single subject as singled out. If we were to mesh all the pictures together into one, we would arrive at something like a picture with many arrows, one for each subject. In the meshed-together

picture no one is singled out. So we don't mesh the pictures together, indeed, to assume that they need to be meshed together in order for them to depict the world is to operate on the basis of the unity assumption.

Say that we are engaged in the metaphysical project of describing everything that is case. We start describing common-or-garden facts including objective facts concerning the organism that is me, ML:

Snow is white ◦ the earth revolves around the sun ◦ Boston lies 190 miles away from New York ◦ ML is a human being ◦ ML's brain is in state S_1 ◦ TN is a human being ◦ TN's brains is in state S_2 ◦ ...

As we describe these objective facts, we come to add representation facts for each and every subject:²²

... ◦ ML represents that the sky is blue ◦ ML represents that white clouds are passing by ◦ ML represents that a crow is flying through the sky ◦ ML represents that ML's foot itches ◦ ... ◦ TN represents that the walls of his office are white ◦ TN represents that there is a coffee mug on his desk ◦ TN represents that birds are chirping outside ◦ ...

Assuming further that certain things appear to be the case when ML interacts with the world, we further extend our description with certain appearance facts and their dependence on ML:

... ◦ *something appears to be the case when ML interacts with the world* ◦ it appears that (the sky is blue) ◦ it appears that (white clouds are passing by) ◦ it appears that (a crow is flying through the sky) ◦ it appears that (ML's shoulder itches) ◦ ... ◦ if ML were to close his eyes (or change his state in similar ways), it would not appear to be the case that the sky is blue and that

²²Though the framework is neutral with regard to what representation consists in, it may help to think of representations in terms of dispositions, such that roughly speaking ML's representing that A consists in ML's being disposed to act as though it's the case that A .

white clouds are passing by \circ if ML were to scratch his shoulder,
it would not appear to be the case that ML's shoulder itches \circ ...

Let us abbreviate the long co-obtainment description as 'Frag₁'. It describes the co-obtaining of a range of facts, whose obtaining constitutes the world as including one subjective perspective on it, namely that which depends on ML. This co-obtainment sentence, Frag₁, captures: (1) the common-or-garden objective facts about the world, (2) how *all* subjects experientially represent the world as being and hence, how all subjects are disposed to behave and respond to that world and, finally, (3) the subjective perspective of ML in terms of a certain collection of subject-less appearance facts; and (4) the dependence of those appearance facts on the state of a particular organism, ML. It's the third set of facts - the appearance facts - that make for the phenomenal character of ML's experience, they describe a single phenomenal field.

Note that the dependence of the appearances on ML has been represented by counterfactual facts of the form $\neg F(s)\Box\rightarrow\neg\mathcal{A}(A)$, where F is some property of the relevant subject s (ML in the case of Frag₁) with which the relevant appearance, $\mathcal{A}(A)$, counterfactually co-varies. The counterfactual claim is only true *insofar as* the relevant appearances obtain, and not as such. For example, take the fact that if I were to close my eyes, it would not appear to be the case that the sky is blue (i.e. ML does not have his eyes open $\Box\rightarrow$ the sky does not appear blue). The counterfactual claim is as such false, at least as long as there are other subjects who look at at a blue sky. But it is true within my subjective perspective on the world, here it seems that if ML closes his eyes, the sky no longer appears to be blue.

It's an interesting question whether the use of counterfactuals really captures the type of dependence that is at play here and, if so, how we should understand the variation in the truth-values of counterfactual sentences across fragments. These are open questions that I cannot settle here and need to leave for future work. All we need for the current discussion is *a* suitable notion of dependence and it *prima facie* seems that counterfactual dependence is the required notion.

Thus far we only have a description of all the objective facts and of the appearances that constitute a single subjective perspective. Of course, we find that another co-obtainment description is also true, capturing the same objective facts but this time replacing the previous appearance facts with those that capture the subjective perspective of TN:

... \circ *something appears to be the case when TN interacts with the world* \circ it appears that (the walls of TN's office are white) \circ it appears that (there is a coffee mug on TN's desk) \circ it appears that (birds are chirping outside) \circ ... \circ if TN were to close his eyes, it would not appear to be the case that the walls of TN's office are white and that there is a coffee mug on TN's desk \circ ...

Let us abbreviate this co-obtainment description as 'Frag₂'. This sentence agrees with 'Frag₁' on the common-or-garden facts (e.g. that Boston lies 190 miles away from New York) as well as on the facts about the representational states of subjects (e.g. that ML represents that a crow is flying through the sky); it only disagrees with 'Frag₁' about what appears to be the case.

We continue this process for each and every subject who has a subjective or first person perspective. We arrive at our overall conception of the world through the conjunctive sentence assertion that Frag₁, Frag₂, Frag₃, etc. each obtains:

$$\text{Frag}_1 \wedge \text{Frag}_2 \wedge \text{Frag}_3 \wedge \dots$$

Note the use of conjunction here, instead of co-obtainment. We arrive at a conception of the world in which the appearance facts do not all co-obtain. We have it for example that $\neg(\text{Frag}_1 \circ \text{Frag}_2)$. The world appears to be this way *but only insofar as* it doesn't appear to be that way. The fragments are each the case in the world described from a third person view, without being the case together or forming one unified world. In particular, the appearance facts that depend on one subject do not co-obtain with those appearance facts that depend on someone else.

To engage in this type of description is to describe the world from an objective view precisely to the extent that it incorporates all the first person

points of view. But there being one total objective conception of the world should be sharply distinguished from it being a conception of a unified world, which it clearly isn't and couldn't be.

With this fragmentalist turn in the no-subject view, it's able to solve the problem discussed in the previous section. According to *Egocentricity*, it's *only* the case that something appears to be the case when ML interacts with the world and nothing appears to be the case when someone else does. The negative aspect of this claim, the denial that something appears to be the case when anyone else interacts with the world is read as the negation of a co-obtainment claim. It's the fact mentioned in the previous paragraph, namely the fact that it's not the case that there is something it's like when ML interacts with the world insofar as there is something it's like when TN interacts with the world. So ML's experiences exhibit phenomenality *insofar as* no one else's experience does, and hence ML is singled out *insofar as* these facts obtain.

The appearance facts arguably capture the first person or subjective character of undergoing experience precisely because they aren't indexed to subjects. To imagine that a tulip appears to be red for example *is* to imagine what it's like to experience the tulip as red. Imagining that the appearance facts obtain seems to make one imagine experiencing the world a certain way (cf. Martin 2002). To imagine that a tulip appears to be red *to you* is arguably *not* to imagine what it's like to you to experience the tulip as red, precisely because it brings you into the picture. To take a more elaborate example, consider everything that Hannibal must have experienced when marching with his troops onto the battlefield at Cannae. Now imagine that those things appear to be the case, that there is a battlefield full of troops marching forward, that legions of Roman soldiers are emerging on the horizon, that it smells of sweat and horses, etc. When you imagine all of this appearing to be the case, you are hereby imagining what it must have been like according to Hannibal to undergo his experiences of the battlefield, attending to both the qualitative and first person character of undergoing that experience. We thereby imagine that the world is manifested the way it is within the experience of that subject, and hence imagine undergoing

the experience of the subject and what that is like. With this understanding of the phenomenal facts, we can understand how their co-obtainment can constitute a subject's first person or subjective perspective as well as how those facts can all fit within the same objective world.

Nagel remarked that the subjective view can claim dominance over the objective view of the world because the ways things appear from a subjective view seem to be missing from the way we objectively take the world to be. This abstract issue has now been addressed. When we engage our subjective view of things, we no longer step outside of the fragmentalist's objective view of the world given that this already acknowledged the subjective exactly as we find it when we reflect on our experience. The subjective view on the world is just a proper part of the objective view on the world, nothing we find in the former is missing in the latter. And so fragmentalism resolves the stalemate between our subjective and objective view of things.

4.5 Co-consciousness and split brains

So we have seen how the no-subject view can offer an account of the distinction between the subjective perspectives of different subjects in terms of fragmentation across appearances. In this final section, I want to discuss the so-called 'unity of consciousness' and the possibility of intra-subjective fragmentation in split brain cases.

Current discussions of the unity of consciousness focus on phenomenal unity, which is often discussed in terms of a relation of co-consciousness (due to James 1909/1967: 221): in addition to being conscious of two things, it's said that we can be co-conscious of those things.²³ Schechter offers the following helpful gloss on co-consciousness:

[W]hen I see a red balloon while simultaneously seeing a blue balloon, the two experiences are co-conscious, and thus there is

²³Experience exhibits various kinds of unity, such as object unity, spatial unity, gestalt unity, and so on. For a discussion of how these differ from phenomenal unity, see Tye (2003: 11-16) and Bayne and Chalmers (2003: 25-35).

something it is like for me to see the two balloons at once: it is not just that there is one thing it is like for me to see a red balloon and another thing, totally separate in my experience, that it is like for me to see a blue balloon: there is also something it is like for me to see a red balloon *while* seeing a blue balloon, something it is like for me to have these experiences *together*. (Schechter 2013: 673).

When I interact with the world, there are not just the individual manifestations of individual facts, there is a single phenomenal field, an overall manifestation of a unified scene within which things are certain ways.

To understand the phenomenal unity of consciousness we need to understand what it is to be co-conscious of bits of world. Within the proposed view, co-consciousness is naturally understood as the co-appearance of certain facts, which in turn is understood as the co-obtainment of those appearances: for s to be co-conscious of A and B is for the appearance that A to co-obtain with the appearance that B (i.e. for it to be the case that $\mathcal{A}A \circ \mathcal{A}B$) and for this to depend on the state of s .

Note however that the co-obtainment of appearances captures phenomenal unity only if the co-obtainment of appearances implies the appearance of the co-obtainment of the content of those appearances. If it could be that $\mathcal{A}A \circ \mathcal{A}B$ without it being the case that $\mathcal{A}(A \circ B)$, the former arguably does not capture how the contents of individual appearances coalesce into one overall unified content of experience. The unity in appearing should make for a unity in what appears. This suggests that we should accept that the co-obtainment of appearance is equivalent to the appearance of the co-obtainment:

$$\mathcal{A}(A) \circ \mathcal{A}(B) \equiv \mathcal{A}(A \circ B)$$

The co-appearance of A and B is (logically) equivalent to the appearance of the co-obtainment of A and B .

An implication of this is that any incompatibilities in the contents of appearances bar those appearances from co-obtaining. Some argue that we can

experience contradictory things, for example, Crane argues that the world can come to appear in incompatible ways when we undergo the so-called waterfall illusion: if we stare at a waterfall for some time and then turn to look at a stationary object, it will seem to be moving and, yet - it is claimed - it will also seem to remain in the same place relative to the background (Crane 1988). I'm not convinced that there is a single experience with a contradictory content. It rather seems to me that the stationary object seems now to be moving, and now to be in the same place relative to the background. There seems to me to be a flipping between contradictory contents, not a single contradictory content. Similarly, when I look at Escher's lithograph 'Waterfall', it seems to me that the stream now appears to be going upward, and now to be going downward, I do not (I think) truly see it as moving upward *as well as* downward *at once*. And this is in line with the equivalence between unified appearances and appearances of unified contents: as the contents cannot co-obtain, neither can they co-appear, forcing a flipping back and forth between two appearances that do not co-obtain. If I'm mistaken in this however, and contradictory things *do* appear to be the case; then it can after all be that $\mathcal{A}(\text{the object is moving} \wedge \text{the object is not moving})$ and that $\mathcal{A}(\text{the object is moving}) \circ \mathcal{A}(\text{the object is not moving})$. This should not in that case imply that $\mathcal{A}(\text{the object is moving} \circ \text{the object is not moving})$. *If* the contents of an appearance can be contradictory (I am not convinced they can), we should deny the above equivalence between unified appearance and appearance of unified contents.

A closely related question is whether the conscious experiences of a single subject are necessarily unified in one phenomenal field of consciousness, one subjective perspective on the world, simply in virtue of depending on the same subject. Or, to reformulate the question for the no-subject view: do the appearances that depend on the same subject necessarily co-obtain? The assumption of necessary experiential unity was an important assumption amongst philosophers of the 17th century, and is defended to this day (see Bayne and Chalmers 2003).²⁴ There are however well-known empirical find-

²⁴Whether modern philosophers are concerned with the unity of consciousness as we now understand it is an exegetically sensitive question. It does certainly seem like it. Descartes:

ings that suggest that appearances that depend on the same subject can be disunified, namely findings concerning subjects whose brain has been ‘split’.

The human brain consists of two halves, called hemispheres, connected by a ‘bridge’ of neural connections, called the corpus callosum.²⁵ To a large extent, each hemisphere receives partial information, from one eye, one ear, one hand, etc. and, oddly, the information is received ‘cross-laterally’, for example, the information from the left eye goes to the right hemisphere, and that of the right eye to the left hemisphere. Subjects who suffer from extreme epilepsy can be helped by cutting the corpus callosum that connects the two hemispheres, keeping an epileptic seizure in one half from spreading to the other. The operation, called commissurotomy, was first performed by Sperry (see Sperry 1968; Gazzaniga 1970 provided a detailed study of the psychological implications of commissurotomy).

Such subjects exhibit very interesting behaviour. Consider the following case (due to Sperry 1968; the presentation here is based on Tye 2003: Ch.5):

Case 1: A subject SB, whose corpus callosum was cut, stares at the center of a screen on which two words are flashed, say ‘pen’ to the left of the fixation point and ‘knife’ to the right of it. The words are flashed so quickly that no eye movement is possible from one word to the other, ensuring that each word can only be registered by one eye only. So, given that optic nerves cross: the right hemisphere receives information of seeing ‘pen’ and the left hemisphere of seeing ‘knife’.

Data in case 1: When asked verbally, SB does not linguistically report seeing ‘pen’ (this is thought to be because brain areas responsible for speech are situated in the left hemisphere, and not the right). When

‘When I consider the mind, that is to say, myself inasmuch as I am only a thinking thing, I cannot distinguish in myself any parts, but apprehend myself to be clearly one and entire.’ (Descartes 1641/1970: 196). Leibniz: ‘[I]n natural perception and sensation, it is enough for what is divisible and material and dispersed into many entities to be expressed or represented in a single indivisible entity or in a substance which is endowed with genuine unity’ (Leibniz 1686-87/1989: 339). Kant: ‘[A]ll appearances stand in a thoroughgoing connection according to necessary laws, and hence stand in a transcendental affinity of which the empirical affinity is the mere consequence’ (Kant 1781/1997: A113-114).

²⁵This paragraph draws on information from Gazzaniga, Ivry, and Mangun (2014: Ch.2).

asked to pick the seen object using the left hand (wired to the right hemisphere), SB will pick out pens and reject knives. When asked to sort through objects with both hands, his left hand (wired to the right hemisphere) will pick out a pen and the right hand (wired to the left hemisphere) a knife, with the two hands working independently.

So SB is disposed to behave differently, depending on whether he is forced to rely on information fed to the right hemisphere, or on information fed to the left. Now of course, we do not know with any certainty how things come to appear when SB interacts with the world, but this has nothing to do with the case in particular; I cannot know with any certainty how things appear when anyone else interacts with the world and yet I trust that things *do* come to appear when others interact with the world, and I use a range of facts to gain a sense of another's subjective perspective. In the case of SB, the facts suggest that there is fragmentation in the way things appear when SB interact with the world. After all: if you only tell me that there are two hands, possibly of one subject possibly of two different subjects, and you tell me that, when the subject or subjects are asked to pick out the object they saw, one hand picks out a pen and the other hand a knife, I will conclude that that the hands belong to different subjects because there must have been two phenomenal fields, one with a pen in it, and one with a knife in it. When you, surprisingly, tell me that both hands *do* belong to the same subject, the plausible conclusion is that, when the subject (i.e. the organism that is SB) interacts with the world, this gives rise to two phenomenal fields (i.e. two subjective perspectives, i.e. two sets of co-obtaining appearances).

With respect to the no-subject view, the empirical data suggests that if it appears that *A insofar as* this depends on a subject *s*, and it appears that *B insofar as* this depends on the same subject *s*, we should not infer that *A* and *B* co-appear, i.e. that $\mathcal{A}A \circ \mathcal{A}B$. So, next to the inter-subjective fragmentation we saw in the previous section, we can accommodate intra-subjective fragmentation: it's possible that there arise distinct phenomenal fields that depend on the same subject, such as SB. Two subjective perspectives on the world can emerge from a single organism.

Many have drawn a similar conclusion from the data and deny that the experience that are based in a single organism are necessarily unified.²⁶ And yet, Bayne and Chalmers have argued that denying the thesis that experiences are necessarily unified (what they call the unity thesis) results in something that is inconceivable or even incoherent:

We think that there is a strong prima facie case that the unity thesis is true. This prima facie case is brought out by the fact that there seems to be something inconceivable about phenomenal disunity. It is difficult or impossible to imagine a subject having two phenomenal states simultaneously, without there being a conjoint phenomenology for both states. [...] This prima facie inconceivability - whether it takes the form of unimaginability or apparent incoherence - gives at least some reason to believe that cases where phenomenal unity breaks down are impossible, so that the unity thesis is true. (Bayne and Chalmers 2003: 38).

The fragmentalist has a straightforward way of explaining the apparent inconceivability away. It's indeed inconceivable that there is one subjective perspective *insofar as* there is the other. When SB is asked to sort through a bunch of objects, his left hand will not only pick out pens, it also *rejects* knives. In other words, in the subjective perspective that depends on SB's right hemisphere, there is not just the apparent presence of a pen, there is also the apparent absence of a knife. That is, there appears to be a pen *insofar as* there does *not* appear to be knife. And yet there appears to be a knife, and there does not appear to be a pen *insofar as* there appears to be a knife. The perspectives conflict, and hence we cannot - in imagination - adopt the one perspective *insofar as* we adopt the other. But, according to the fragmentalist, our inability to conceive the co-obtainment of all the appearances that depend on SB is no evidence that things do not appear those ways. We have it that $\mathcal{A}(\text{there is a pen}) \circ \mathcal{A}(\text{there is no knife})$ and that $\mathcal{A}(\text{there is no pen}) \circ \mathcal{A}(\text{there is a knife})$, to imagine that this is the case is not to imagine that the one is the case *insofar as* the other is the case. This

²⁶See, amongst many others, Sperry 1968, Nagel 1979a, Lockwood 1989 and Tye 2000.

is just another case of the distinction between imagining a co-obtainment and imagining a mere conjunction.

As in other cases, it may be difficult and perhaps impossible to imagine the conjunction of incompatible conjuncts without lapsing into imagining their co-obtainment. With the current picture of appearances, we can see a possible explanation of *why* it is difficult to imagine a mere conjunction without also imagining the conjuncts as co-obtaining. Imagining something as being the case is plausibly understood as making something appear to be the case - or at least to be in a state that often gives rise to certain (coarse-grained and rough) appearances. If this is right, and one makes it appear as if a certain conjunction obtains, then this makes the conjuncts co-appear, and this implies the appearance of the co-obtainment of the two conjuncts - given the above mentioned equivalence between co-obtaining appearances and the appearance of a co-obtainment). And so, to imagine facts that fail to co-obtain, one requires fragmentation in the acts of imagination themselves: one needs to imagine the first without the second and the second without the first. To imagine a worldly fragmentation one needs to bring about a fragmentation in imagination. This is difficult, and perhaps impossible to do simultaneously. The closest we can come to this is by imagining the two conjuncts sequentially: *first* the one, and *then* the other.²⁷

Note also that the data concerning split brain subjects puts pressure on a possible motivation for thinking that phenomenal objects essentially involve a subject. One might think that experiences are essentially had by subjects because subjects just are in some sense the single subjective perspectives or, at least, are individuated by such subjective perspectives.²⁸ On such

²⁷Cf. Tye: '[I]n one sense we can project ourselves into the mental lives of split-brain patients: we know just what it is like to experience each experience a split-brain subject undergoes. What we cannot do is to imagine ourselves undergoing both the visual experiences a split-brain subject experiences at any given moment while the experimental controls are in place and simultaneously being aware, *all in one introspective act*, of the content of each.' (Tye 2003: 120, my italics).

²⁸Different thinkers have drawn different conclusions with respect to the number of subjects vis-à-vis the number of subjective perspectives. To mention just a few, Sperry (1968) concludes that commissurotomy gives rise to two subjects, each with their own mind (and own free will); Parfit (1984) concludes that there is one subject who has two streams of consciousness that remain divided in all circumstances after the commissurotomy; Marks

a view, an experience is had by a subject if and only if it falls within the phenomenal field that individuates the subject. In split brain cases, however, this understanding of subjects, as being or corresponding to the reification of subjective perspectives, leads to substance dualism. If the subjects of experience (or phenomenal facts) go with single subjective perspectives, then there are arguably two subjects in the case of SB. Call them SB-L and SB-R. SB-R sees a pen and doesn't see a knife, SB-L sees a knife and doesn't see a pen. As SB-R and SB-L are non-identical, neither can be identical to the physical organism we called SB, whose corpus callosum is cut, given that identity is transitive. And so we are led to substance dualism: at a given time, the subject of experience is distinct from the physical organism that gives rise to the relevant experiences. If one wants to avoid substance dualism and yet one believes that there are two distinct subjective perspectives in the case of SB, then one must believe that subjects of experience cannot be individuated in terms of subjective perspectives.

Note also that - even on the view that the experiences of SB are all *had* in some essential sense by the same physical organism, viz. SB - there is some pressure to resort to fragmentalism. After all, the subjective perspectives conflict and are nevertheless thought to inhere in the same organism: take the phenomenal field that depends on the information fed to the left hemisphere of SB. If we imagine adopting it, the experience of a pen is absent from it. That is, given what there is and isn't in this phenomenal field, it seems clear that the relevant subject does *not* experience a pen. The experience of a pen is absent in the very same way that the experience of my keys (I'm sure) is absent from this phenomenal field. But if we now consider the phenomenal field that depends on the information fed to the right hemisphere of SB, it seems clear that SB *does* experience a pen. But, then, how do we capture the apparent absence of the experience of a pen in the perspective where the

(1980) and Tye (2003) conclude that there is one subject whose stream of consciousness becomes divided only in special circumstances, such as those created in the experiments; and Nagel (1979a) concludes that there is an indeterminacy in how many subjects there are in split brain cases and that this throws doubt on the very concept of a 'subject of experience'. For a discussion of the interaction between subjective perspectives and subjects of experience, see Schechter (2013).

knife appears to be present?

One might propose we distinguish between experiencing that *not A* and *not* experiencing that *A* (cf. Tye 2003:116). SB experiences that there is a pen and experiences that there is no pen, where these experiences are not co-conscious. This way we avoid saying that SB both has and does not have the experience of a pen. Even if we make this distinction however, given the one phenomenal field, there is still as much reason to deny that SB experiences a knife as there is reason to deny that SB experiences my keys, and it seems clear that SB does not experience my keys. There is absolutely no phenomenal differences supporting the claim that, in the one case the absence consists in *the experience of an absence* and in the other case in *the absence of an experience* of something. And hence there is reason to resort to fragmentalism here: SB experiences a pen *insofar as* SB does not experience a knife just as SB experiences a pen *insofar as* SB does not experience my keys. SB also experiences a knife *insofar as* SB does not experience a pen, just as SB also does not experience my keys *insofar as* SB experiences a knife. If an adherent of the subject view of phenomenal facts were to adopt fragmentalism, then the subject view appeals to the conceptual machinery that - I have argued - is required by the no-subject view.

Returning to the no-subject view, there is another matter worth emphasizing. Since co-consciousness is understood in terms of co-obtainment, and co-obtainment is non-transitive, the view implies that co-consciousness is non-transitive. This means that the view implies a standpoint on a controversial topic: some argue that co-consciousness should indeed be non-transitive but others that it should be transitive. Again, however, there are empirical findings which, when taken at face value, suggest that co-consciousness is indeed non-transitive (first proposed as such in the neuroscientific work of Trevarthen 1970).

The two hemispheres of a brain are jointly connected to an area that is known as the brainstem, which is thought to be intimately related to our emotions. As the two hemispheres are connected by the brainstem, one expects that emotions can be shared between the hemispheres of split brain subjects. And this is indeed what was found. Consider the following case

(based on Gazzaniga 1985):

Case 2: A subject VP, whose corpus callosum has been cut has a frightening scene of a fire presented to her right hemisphere only (via the left eye).

Data in case 2: When asked verbally what she saw, VP reports not seeing anything interesting and certainly not a fire. When asked verbally how she feels, however, VP said 'I don't really know why but I'm kind of scared. I feel jumpy. I think maybe I don't like this room, or maybe it's you. You're getting me nervous.' (Gazzaniga 1985: 77).

Since it's the left hemisphere that is responsible for verbal reports but the frightening scene was presented to the right hemisphere, it seems that the emotion triggered by information in the right hemisphere is available to the left hemisphere. There thus seems to be an emotional bridge between the two hemispheres, as it were, provided by the brainstem connecting them. This case and similar such cases, suggest an overlap in the phenomenal fields.²⁹

As far as I can see, the non-transitivity of co-obtainment is exactly right. Why would anyone think different, and suppose that phenomenal fields cannot overlap? One reason is theoretical: if one assumes that the distinct subjective perspectives make for distinct subjects *and* that experiences are tokened in terms of the subject that has them, then the phenomenal fields do not overlap, as there are distinct token experiences (of the same type) had by distinct subjects. But neither assumption is made by the no-subject view, and I have already mentioned a worry for the view that there are distinct subjects for each distinct subjective perspective: it leads to substance dualism. Without the assumption that there are two subjects residing in split brain cases, even the assumption that there token experiences does not help

²⁹Lockwood (1989), who also defends that co-consciousness is non-transitive, argues for it on the basis of a thought experiment: what if the corpus callosum is cut, fiber by fiber. Do we think that there is some one fiber which, when cut, makes a second phenomenal field pop into existence with token copies of experiences that are found also in the other phenomenal field? This seems implausible. - I'm not sure what to think of this argument: it seems to me that how the phenomenal field is affected by such gradual changes is an empirical matter, about which we can hardly expect to have any reliable intuitions.

to avoid the overlap between phenomenal fields, as there would not be two different subjects that make for the two distinct token experiences. In this case too, the natural conclusion is that co-consciousness is non-transitive.³⁰

To sum up. We saw in the previous section that the no-subject view accepts a fragmentation in the appearances across subjects to account for the conflicts across multiple subjective perspectives. The appearances are not relativized or tokened to subjects - as on standard views. But those standard views need to account for a similar sort of fragmentation in the case of split brain subjects. One option here is to posit a distinct subject for each of the distinct subjective perspectives that arises in split brain cases, so that we can account for the conflicts across those subjective perspectives in the same way the ordinary view accounts for the distinct subjective perspectives that arise from distinct subjects: through relativization. If so, we need to posit, implausibly, two subjects, one for each subjective perspective in the split brain case - even though there is only one organism. For the no-subject view, the fragmentation that arises in split brain cases is phenomenally the same as the fragmentation that arises across subjects. So, there is inter-subjective fragmentation when appearances that depend on distinct physical organisms fail to co-obtain, and we have intra-subjective fragmentation when appearances that all depend on the same physical organism fail to co-obtain. The no-subject view does not only have independent phenomenological considerations in favour, it also allows us to interpret the empirical data at face value.

4.6 Summary and concluding remarks

As you find yourself in the world and interact with it, those things appear in various ways. The appearances that constitute your subjective experiential view do not obtain together with the phenomenal facts that constitute my subjective experiential view on the world: the one set of co-obtaining appearance facts obtains insofar as the other set of co-obtaining appearance

³⁰For a critical discussion of non-transitivity, see Dainton (2000: Ch.4) and Hurley (1998: 175). For a defence of it, see Tye (2003: §5.6) and Schechter (2014).

facts doesn't and vice versa. The account that arises from this harmonizes an objective view of the world with the deliverances of our subjective views on the world. The consequence is that the appearances make for a metaphysical fragmentation of the world.

Needless to say, the account is far from a complete theory of experience. Many more details need to be filled in and many questions need to be answered. The notion of appearance that makes for phenomenality needs to be further elucidated and there is the - thus far bracketed - question whether any of this clashes with physicalism. These and other questions will need to be addressed at some other occasion. Rejecting that the world is a unified place may seem drastic. But conscious experience is ill-understood. We should all be open to the possibility that, as this chapter maintains, the unity assumption stands in the way of a better understanding of conscious experience. I hope I have said enough to place a view on the table that can be critically engaged with, and that - perhaps - allows us to see consciousness in a light that will prove to be illuminating.

Chapter 5

Colours in a Fragmented World

Are things really coloured, out there in the world? Colours, I will assume, are amongst the ways things appear to be in your visual experience. For a thing to be coloured is for it to be at least the way it appears to be in your visual experience of it. So, the question is, are things really the way they appear to be in your visual experience of them, out there in the world?

There is a well-known problem that led many to think that things cannot be coloured out there in the world. Objects can seem to have incompatible colours in different circumstances, without any of these circumstances being privileged over the other. I will discuss this problem, offer a fragmentalist solution to it and discuss alternative views.

This is a short chapter. The reason for this is that the paradox of conflicting appearances is well-known and the solution offered here straightforward. To make clear why the fragmentalist view of colours is interesting, I will criticize some of the more popular alternatives again. But, as before, the criticisms are meant only to express dissatisfaction with the status quo, not sufficient reason to overthrow it.

5.1 Conflicting appearances

When discussing colours, we need to sharply distinguish between two questions: what are colours? And is the world coloured? Because there are

different views of what colours are, there is a risk of confusion when we discuss whether the world is coloured. Coloured in what sense?

I will call ‘colours’ the coloured ways things appear to be when we experience them. It’s the way they are in your visual experience of things, such as the redness that appears to sit right there on the surface of the red ball that you are experiencing, or the redness that seems to permeate the red wine in your glass.

The colours of things are manifest in your experience in the very same way in which the shapes of things are manifest in your experience. When you experience a round ball, roundness is manifested in a way that makes it clear how something that is round is more similar to oval things than to square things, and how nothing can be both round and square. The same goes for colours: the redness of the ball is disclosed in a way that shows how something that is red is more similar to orange things than to blue things, and how something cannot be red and, say, green all over.

As in the previous chapter, I will assume that the way things appear when we experience things are ways things appear *to be*, i.e. properties that things appear to have. This I called the *ways = properties principle*. and implies that colours are properties. Indeed, like the shapes we experience, colours are intrinsic *sui generis* properties.¹ And there is a wide variety of objects that have such colours, besides the surfaces of middle-sized everyday objects, colours appear as permeating materials or atmospheres (such as water and glass), or as qualifying the light emitted by some light source.

The qualitative nature of colours exhibits various systematic patterns, which are used to ‘order’ the colours.² The hues of colours divide first of all into two classes, the achromatic hues such as black and white, and the chromatic hues such as red, yellow, orange, blue, pink, etc. Amongst the chromatic hues, there is furthermore a distinction between unique hues and binary (or ‘mixed’) hues. There are unique hues of green, yellow, red and blue (i.e. there is a ‘unique green’, a ‘unique yellow’, etc.) and any other hue

¹This is known as the simple view of colours. Realism about colours so understood is defended in Campbell (1993), Yablo (1995), Gert (2008) and Watkins (2005).

²This paragraph draws on Hardin (1988: Ch.1).

is a binary hue: there is a blueish green, a reddish yellow, etc. These hues are all mutually incompatible: when a patch of some surface appears unique yellow, it thereby does not appear to be reddish yellow. Part of what it is for the patch to appear unique yellow is for it *not* to appear reddish in the way that reddish yellow appears reddish, and part of what it is for the patch to appear reddish yellow is for it *not* just to appear yellowish but also reddish. Hence the patch cannot appear both unique yellow and reddish yellow, and similarly for all other hues of colours. They are all incompatible properties.

Even though colours are incompatible, there is a large range of cases where something seems to instantiate incompatible colour properties when observed under different conditions. There are many kinds of such circumstances (see Cohen 2009: Ch.2):

Colour blindness. There are typically three types of cones in our retinas, that differ as to whether they are sensitive to short wavelengths, medium wavelengths and long wavelengths. But there is a considerable percentage of people who have less than three cone-types, or whose cone types are anomalous. We call these people ‘colour blind’ because they are unable to make certain colour discriminations that a larger percentage of the population are able to make. They do not see the colours we see.

Goldfish. There are many other types of animal that have more cone types than humans, whilst otherwise processing colours in much the same way as we do. The retinas of goldfish for example, have cone-types that are sensitive to the near ultraviolet and are thus sensitive to difference in wavelengths that we humans aren’t sensitive to (see Byrne and Hilbert 2007:94, and the references therein). Such animals see colours we do not see.

Soap bubbles. So-called iridescent surfaces, such as those of soap bubbles, DVDs, butterfly wings and insects, exhibit rainbow like colour patterns that shimmer and change as the viewing angle changes.

Backgrounds. When we take two squares that appear to have the same shade of grey and place the one in front of a dark background and the other in front of a lighter background, the first square (against the dark background) will now appear a lighter shade of grey than the second square (against the lighter background). (See Cohen 2009: 20-22, and the references therein).

Sunglasses and illumination. The sky can seem orange when seen through sunglasses and a banana can look green when seen in blueish light.

These are some of the many ways in which variations can arise in the way things seem to be coloured in different circumstances, which include being experienced by subjects with different visual systems or by the same subject under different viewing conditions. There are even reasons to think that normally sighted subjects viewing the same object in the exact same external conditions will not typically see that object to have the exact same hue of colour. It has been found, for example, that two normal random human subjects will typically differ in what look unique green to them. The wavelengths that people report as unique green varies from 490 to 520nm. This is a significant variation, as differences of 15nm either way make for a distinctly bluish or yellowish look.³

These colour variations give rise to a well-known paradox. Assuming that we are two normally sighted human subjects, imagine that we both look at the same patch of surface of a soap bubble but under slightly different angles. Because of the different angles, and the fact that the surfaces of soap bubbles are iridescent (i.e. such as to look different when seen under different angles), the same patch will appear green to you but will appear blue to me. But *being green* and *being blue* are incompatible: the patch cannot be both all over. At the same time, there is no angle that is the one true angle from which the soap bubble must be observed to see its true colours. And there is also no reason to think that our experiences are non-veridical simply by

³See Byrne and Hilbert (2003: 16) and the references mentioned there; see also Block (2003: 188).

virtue of occurring from specific angles. So, we have here the following three facts:

1. The same patch p of the surface of a soap bubble appears to be green all over to Norm and appears to be blue all over to Norma.
2. The only salient difference between Norm and Norma is that they view the patch from different angles, and this is no reason to think that either one, or both, of the experiences fail to be veridical.
3. The same patch p cannot both be green all over and blue all over.

Something appears to instantiate incompatible properties from different perspectives, and we have absolutely no reason to privilege one perspective over the other.

This argument has a long history and has been used to motivate various views (see Burnyeat 1979). First and foremost, however, this argument is one of the prime motivations for antirealism about colours. Our experiences, it is claimed, are after all not veridical, not because the patch appears in different ways when viewed under different angles, but because nothing is coloured at all. The world is colourless; nothing is coloured in the way it appears to be in experience, and to that extent the experience of coloured things isn't veridical at all. As Mackie puts it, 'no colour as we see colour [...] is literally in or on the things we call coloured' (Mackie 1976: 14).

This antirealist position about colours leads to a stark asymmetry in the way the contents of appearances are treated: things can be shaped the way they appear to be in our experience, yet they cannot be coloured the way they appear to be in our experience. It's as if our visual system takes a spatial print of the world, and then colours it in.

Many have searched to avoid antirealism about colours. For the remainder of this section I want to briefly discuss the main views.

If colours are ways things appear to be when we experience them, and things appear in incompatible ways, it seems that things cannot be coloured. Perhaps then, many have thought, colours are not the ways things appear to be in our experience of things. The conception of colours I adopted above is

known as primitivism. It's part of that view that colours cannot be defined in non-chromatic terms, and that they are not relations to the perceiver, reflectance properties of surfaces, or dispositions. These are all instantiated all right, but they are not colours. This view of colours stands opposed to various kinds of reductive analyses of colours, which identify what we describe in colour terms with something that is not described in colour terms.

One reductive approach identifies colours with *reflectance properties* of the surfaces of objects. Objects have various capacities to absorb, emit, reflect, or scatter light. The precise degrees to which some surface does any of these things determine the *surface reflectances* of the objects. The reflectance type is roughly the proportion of light that the object reflects at each wavelength in the visible spectrum. Some identify the colours with these reflectance types or, alternatively, with the disjunctions of the particular reflectances that are of the relevant type. When we experience things to have different colours, our experience represents things to have different reflectance types.⁴

I should note that this view isn't motivated in terms of the paradox of conflicting appearances and the reason for this should be clear: it does not help solve it. Different reflectance properties are incompatible properties, for something to reflect *this* proportion of light at each wavelength is in part for it not to reflect any other proportion of light at those wavelengths. Something reflects either *this* proportion of light at each wavelength or *that* proportion, and it cannot reflect both proportions. So, as we experience different colours, we must be representing the soap bubble as having different reflectance properties. Who is right? Proponents of the view bite the bullet: someone must be, and the other must be mistaken.⁵ Given that the only relevant difference between us is the angle at which we view the soap bubble, there would turn out to be one true angle from which to view the bubble in order to see its true colours. This is quite implausible, as Berkeley asked:

And now tell me, whether you are still of the opinion, that every

⁴Reductive views like this are defended in Armstrong (1969), Jackson (1996), and Tye (2000). I focus on the statement of the view given in Byrne and Hilbert (2003).

⁵See Byrne and Hilbert (2003: §3.4).

body hath its true real colours inhering in it; and if you think it hath, I would fain know farther from you, what certain distance and position of the object, what peculiar texture and formation of the eye, what degree or kind of light is necessary for ascertaining that true colour; and distinguishing it from apparent ones. (Berkeley 1713/1979: 22).

These are awkward questions. The reflectance view reduces colours to physical properties, for which the same paradox arises as for the colours understood as the way things appear to be in experience.

Besides this issue, things simply do not *appear* to have specific reflectance types, they appear to have colours, simple intrinsic properties. So the fact that they do have such reflectance types out there in the world does not constitute a realism with regard to the ways things appear to be in our visual experience of them - whether we call the reflectance types colours or something else. But it is only the ways things are like in our visual experience that are at question.

Another prominent reductive approach identifies colours with the disposition to produce certain kinds of experiences in subjects in certain circumstances. Here the paradox of conflicting appearances does not re-arise. The same patch of the soap bubble can be disposed to look one way when viewed under one angle, and be disposed to look another way when viewed under a different angle.⁶ The same goes for a close cousin of the disposition view, according to which colours are identified with triadic relations holding between object, perceiving subject and the circumstances in which the experience occurs. For example, the orange of a clementine is identified with the relation of *being orange-for-subject-S₁-in-circumstances-C₁*.⁷

Even though the problem of conflicting appearances does not arise for the dispositional and relational views, these views still do not help solve that problem. When we experience the soap bubble, the patch appears in incom-

⁶Dispositionalist views are defended in McDowell (1985), Peacocke (1984), McGinn (1983) (who later retracted the view, in his McGinn (1996)) and Johnston (1992).

⁷Relationalist views are much more recent than the other views. They are defended in Averill (1992), McLaughlin (2003) and, most extensively, in Cohen (2009).

patible ways to us. Our experiences disagree about what that patch is like. But the dispositions and relations that they register are not incompatible. Indeed, that is precisely how these views solve the problem of conflicting appearances. But there is incompatibility in the way the object appears in our experience of it. Even if the incompatible ways in which the object appears to be aren't the colours (because colours are said to be some disposition or relation), they remain qualitative ways things seem to be, and remain just as problematic, no matter how they are labelled. That is to say, there remains the question whether objects are those ways out there in the world, even if we use the colour terms for various relations that the object bears to perceiving subjects and even if things really do have those dispositions or really do stand in those relations. Talking about the dispositions or relations merely draws our attention away from those properties that are problematic; it merely changes the subject matter. We can all agree that there are such relations, and we can name anything any way we want, but the question was simply whether objects are the incompatible ways they appear to be in our experience of them. And the relational and dispositional views remain silent on this.

These views aren't silent on the question at issue if they insist that the ways things appear to be *are* the relations or dispositions. Just as things appear to be shaped certain ways, they appear to be disposed to look certain ways. Thus we would have the flatfooted denial that the patch of soap surface appears in incompatible ways, that we have misinterpreted the phenomenological data. It doesn't confront us with the incompatibilities that we thought it did. But this solution is really implausible. For example, in our discussion of change we saw that something only appears to change when it instantiates now one property, and now an incompatible one. So say that we take a yellow banana and then shine a blue light on it. If all that is going on is that it now appears to have one relation and now a relation compatible with it, it doesn't appear in incompatible ways and hence doesn't appear to change in terms of colour. But surely, when I take a yellow banana and shine a blue light on it, it does appear to change in terms of colours.

In short then, the appearances of things are incompatible when the things

appear to have distinct colours. The dispositions and relations aren't incompatible. That things instantiate such dispositions and relations doesn't speak to the question whether things are the way they appear when they appear in incompatible ways under different conditions. We call the incompatible ways things appear colours. Since the relations and dispositions aren't incompatible in those circumstances in which things appear in incompatible ways, they aren't colours.

A recently proposed view, color pluralism, reverses things: though the colour-appearances *are* incompatible, the colours themselves aren't. In the case of conflicting appearances, the object really has the colours they appear to have, even though the object cannot appear to have them both.⁸ Harman has claimed that we overlook the possibility of the colour appearances being incompatible without the colours themselves being incompatible:

On reflection, one must allow that red-green incompatibility may be an illusion. Suppose an object looks red from one angle and green from a slightly different angle. In that case, might we not say that the object is red all over and also green all over, even though one cannot see both colors at the same time? In imagining cases, one forgot about that possibility; one tried to imagine something that looked both red all over and green all over at the same time. But something could be both red all over and green all over at the same time without looking both red all over and green all over at the same time. (Harman 2001:661).

According to Harman, something may be unable to *appear* red and green all over, and yet be red and green all over.

The crucial assumption that Harman makes is that when we try to imagine something that is both red all over and green all over, we are really trying to imagine something that *looks or appears* red all over and green all over; that we cannot do so, only points to the incompatibility of something *appearing* red all over and green all over. But why think that when we try

⁸See Kalderon (2007).

to imagine something that is both red all over and green all over (i.e. try to imagine an object being a certain way), we are really trying to imagine something that looks or appears red all over and green all over? Surely, what we are really imagining in such cases is just something having two properties, not something looking or appearing to have those properties. We can dissociate the *property* something appears to have from its *appearing* to have that property. That is, the ways in which things appear when we experience them are ways things appear *to be*, i.e. properties things appear to have. Given an appearance, we can imagine objects to really be as they appear to be, and this is just to imagine the object to have the property it appears to have. We try to imagine something really being red all over and green all over, we fail, and we take this to be evidence showing that *being red* and *being green* are incompatible, the properties themselves, not just the appearance of them. Similarly, when we imagine a square, we simply imagine a square object, and not something appearing to be square.

Now of course someone might continue to insist that what we are really doing in such cases is to visualize things, and that this is to imagine something appearing to be the case, which only provides evidence for the appearances to be incompatible. And, indeed, there is no qualitative element in the imagined case in which something appears red, that we can point to in order to distinguish it from the imagined case in which something is red. (There is only a difference in the instruction of what we try to imagine: Imagine something being red! Now imagine something appearing red!) But there is another more theoretical reason to think that the properties themselves are incompatible. To be unique yellow, we saw above, is to be yellowish and *not* reddish, and to be reddish yellow is to be yellowish and reddish, and *not* just yellowish. For something to be unique yellow as well as reddish yellow, it has to be both reddish and not reddish. This is a simple contradiction and we do not have to imagine anything to ensure that it doesn't obtain. We do not need to imagine any scenarios to rule this out; we have evidence for this simply on the basis of the way the properties themselves can be systematically characterized.

The question is whether things can be the way they appear to be. The

following all seem truisms:

Colour realism: Objects are coloured.

Colours = ways: Colours are amongst the ways things appear to be when we experience them.

Colour variance: An object can appear to have incompatible colours from different perspectives, none of which is privileged.

Things appear to have incompatible colour properties across different conditions. No collection of mutually compatible properties or relations can be plausibly identified with the relevant properties. Nor can we plausibly say that things are only the way they appear to exactly one subject in exactly one condition. If we cannot say that things have both incompatible properties, nor that they have just one of them, the inevitable conclusion is antirealism.

5.2 A fragmentalist account of colour

The fragmentalist framework helps us make sense of a coloured world. The starting point is that, when things appear to have incompatible properties in circumstances none of which is known to be either privileged or distortive, we have defeasible evidence that the object really has those incompatible colours. Something can be green all over and blue all over but, given that these are incompatible properties, nothing can be one colour *insofar as* it is another. Insofar as an object is coloured one way it isn't coloured the other way, and vice versa.

Say that we start with a description of the physical properties of the soap bubble, until at some point we come to add the colour of the soap bubble's surface:

The bubble is made of soap ◦ the soap bubble has a round shape
◦ the bubble has reflectance property F ◦ ... ◦ *the bubble is green.*

Call this description C_1 . We have another such description of the bubble, overlapping with the previous one yet differing in the colour ascribed to the bubble's surface:

The bubble is made of soap \circ the soap bubble has a round shape
 \circ the bubble has reflectance property F \circ ... \circ *the bubble is blue*.

Call this description C_2 . Both C_1 and C_2 obtain, but they do not co-obtain:

$C_1 \wedge C_2$...

Indeed, to be fully explicit, we say that \neg (the bubble is green \circ the bubble is blue), even though the bubble is green, and the bubble is blue. An object can have a range of colour-facets or guises, exhibiting the one only insofar as it doesn't exhibit the other.

With this conception of coloured objects comes a certain conception of what it is for us to experience a coloured world and the conditions that bear on what we experience. Why do we, on some occasion, experience one colour and not the other? Because, under such conditions, our experience co-obtains with the object having that colour, and doesn't co-obtain with the object having the other colour. So, to return to our experience of the soap bubble, we have it that:

The bubble is green \circ you visually experience the bubble \circ it appears to be the case that the bubble is green.

And my case we have it that:

The bubble is blue \circ I experiences the bubble \circ it appears to be the case that the bubble is blue.

Why don't you experience the bubble as green? Because insofar as you experiences things (under a specific angle), the bubble *is* blue and not green. The conditions under which things are experienced to be coloured are conditions that determine what colour the experiencing co-obtains with. Insofar as you experience the bubble under one angle, the bubble is objectively blue

and not any other colour. As you tilt your head and come to see the bubble from a different angle, you make your experience co-obtain with different colours of the bubble, and if your experience happens to be of those colours it co-obtains with, it is a veridical experience.

This consideration answers an objection that might apply to other views which allow things to have multiple colours. Consider this (long) passage by Chalmers:

[I]magine that the very same apple normally causes phenomenally red experiences in one community and normally causes phenomenally green experiences in the other. We can now ask: when a subject in the first community has a phenomenally red experience of the apple, and a subject in the second community has a phenomenally green experience of the apple, which of these experiences is veridical? Intuitively, there is a case for saying that both experiences are veridical. But this is an unhappy answer for the primitivist. On the primitivist view, any phenomenally red experience attributes perfect redness, and any phenomenally green experience attributes perfect greenness. If both experiences are veridical, it follows that the apple instantiates both perfect redness and perfect greenness. The argument generalizes. For any phenomenal color, it seems that there is a community in which the apple normally causes experiences with that phenomenal color. Taking the current line, it follows that the apple instantiates every perfect color! The choice of an apple was unimportant here, so it seems to follow that every object instantiates every perfect color. It follows that no color experience of an object can be illusory with respect to color. Whatever the phenomenal color of the experience, the object will have the corresponding primitive property, so the experience will be veridical. This conclusion is perhaps even more counterintuitive than the conclusion that all color experiences are illusory. (Chalmers 2010: 400).

Chalmers argues that for any object and any colour, we can imagine that

the object appears to have that colour in standard viewing conditions to some community. This, Chalmers think, has two implausible consequences: any object has any colour and, because of this, any colour experience of an existing object is a veridical experience. There is no possibility for error.

First of all note that, from the fact that for any two colours, we can imagine cases in which objects appear to have those colours, it doesn't follow that objects actually have those colours. It only follows that for any two colours we can imagine an object having those colours. Perhaps, we can imagine a scenario in which any object is experienced to have any colour by some community. But, what this shows is just *the possibility* of objects having all the colours, not that any actual object actually has all colours. An actual object has just the colours it has, and this can be one colour, three, 10. We, on our side, attribute those colours to an object for which we have reason to think that it has those colours. The mere possibility of an object having all colours is not a reason to think that it actually has all colours.

Moreover, even in the possible case where an object has all the colours, it is not the case that any colour experience of that object is veridical, at least not on the fragmentalist account. (Of course, Chalmers didn't raise the problem for the fragmentalist account but for a simpler version of colour primitivism). To take the case of the soap bubble again, it may be the case that the patch is green *insofar as* Norm experiences the patch as blue (say because Norm is viewing the patch under the angle under which it is green whilst having taken a hallucinogenic drug). In such a case, Norm's experience of the patch as blue isn't veridical, even though the patch is blue. The patch doesn't just need to be blue, it needs to be blue *insofar as* Norm experiences it as blue in order for that experience to be veridical. Even in the imagined scenario in which any object has any colour, there is still a possibility for error.

Note also that, if I take some hallucinogenic drug and start to see a wild display of colours on my normally white ceiling, there is nothing in the above conception that forces us to say that the wall really is all of these colours, simply because they appear in my hallucinating condition. We know of conditions under which we experience things that make us see things that

aren't the case. With those same hallucinogenic drugs my ceiling may seem to grow vines hanging down even though there are no such vines. We know that this is a hallucinating condition because if I try to climb up using the vines that appear to be there, there is nothing to hold on to and I will fall down. And what goes for the vines goes for the colours: when we know that some condition is one under which we hallucinate, the way things appear do not then offer us a reason to think that the world is a certain way, the condition distorts the way things appear and stands in contrast to conditions in which things appear as they are. The fact that we take multiple conditions seriously as giving us a reason to think that an object is coloured a certain way doesn't mean that we have to take all such conditions seriously as providing us such evidence. Seeings things under a certain angle is not a distortive condition; seeings things when under the influence of drugs *is* a distortive condition.

This conception of the conditions of perception supports a conception of the colours as mind-independent properties of objects. The angle under which the bubble is perceived is not responsible for the colour of the bubble, it's only responsible for the co-obtaining of someone's experience with the bubble's being of that colour. In the limit case in which there is no experience, the bubble is still the colours that it is, only no experience co-obtains with any of these colours, which is just to say that none of the colours is experienced. Because the conditions under which the appearance of things differ only bear on the experience of things, there is no reason to assume that the colours are dependent on them.

It should be clear how the paradoxes of conflicting appearances is resolved on this conception of things. Things are the way they appear to be, even those incompatible ways that they appear to be from different perspectives. We maintain a realism about colours, without needing to identify colours with things they are not, and without needing to deem some specific perspective on objects to be the One perspective where things show their true colours.

5.3 Summary and concluding remarks

The problem of conflicting appearances is old. The paradox of conflicting appearances was of concern to pre-Socratic thinkers such as Democritus, Heraclitus and Protagoras, is discussed in detail in Plato's *Theaetetus* and features heavily in the philosophy of the 18th century, particularly in Berkeley's argument for idealism, (see Burnyeat 1979). Historically, it has been a major factor in the development of various views. It plays a role in drawing a veil of mental particulars between us and the world, stimulating the thought that, since an object cannot be red all over and green all over, and yet redness and greenness seem instantiated in the world, it must be 'ideas' or 'sense data' that instantiate these properties. It also seems that the perspectival variance is one of the main motivations for the distinction between the so-called primary and secondary properties.

The fragmentalist view is simple and secures a coloured world. Objects are really the way they appear to be in our visual experience of the world. An object can be red all over and blue all over, though they cannot be red all over *insofar as* it is blue all over. Depending (nomonologically) on certain conditions obtaining, our experience co-obtains with the perceived object having a certain colour, and is veridical when it is an experience of that very colour. There is no asymmetry in the way the contents of appearances are treated: things can be coloured the way they appear to be in our experience in the very same way in which they can be shaped the way they appear to be in our experience. There is no need to distinguish between the primary and secondary properties.

Chapter 6

Special Relativity in a Fragmented World

Do objects have the lengths, shapes and masses they appear to have? And do objects have constant velocities? Do events have durations? And do they occur in a certain temporal order? Just as the paradox of conflicting apparent colours threatens the view that the world is genuinely coloured, similar arguments threaten the view that objects have lengths, shapes and masses, and that events have durations and occur in a certain kind of temporal order. For each of these, certain facts related to special relativity show that things instantiate incompatible lengths, masses, etc. from the ‘perspective of different frames of reference’, all of which are on a par.

As the relevant arguments, and the physical facts that underlie them, are more complicated than in previous sections, I will introduce them more slowly, and in two steps.¹ After having introduced the relevant facts, I will argue that they really do seem to have antirealist implications when combined with the unity assumption, and that they have indeed been understood in this way. I then offer a fragmentalist interpretation of the relevant facts, showing how it does not have the antirealist conclusions concerning the various properties under discussion. Finally, I will revisit the passage theory of time, and show how the proposed account of relativity is compatible with

¹The information as well as the manner of presentation relies upon Maudlin (2011), Maudlin (2012), Dainton (2010), Sartori (1996) and Einstein (1920/2001).

the passage theory of time offered in §3.2 yet not compatible with a standard A-theory.

6.1 From Newtonian to Galilean spacetime

Imagine that you wake up and find yourself in a spacesuit, hanging still in space. You only see a constellation of white dots in the far away distance, not in any way moving. After some time, you suddenly see me, also in a spacesuit, floating by. From your perspective, you seem at rest and I seem to be moving by at a constant velocity. - But I also woke up some hours ago to find myself in a spacesuit hanging still in space. The only thing I see is a (different) constellation of white dots in the far away distance. And after some time, I suddenly see you, floating by. From my perspective, I seem to be at rest and you seem to be moving by at a certain velocity. We make radio contact at the moment of passing and waste our few minutes of social interaction arguing who is at rest and who is moving by.

Could one of us be right? Is there any answer to a dispute like this? Newton thought there was. Newton took space to be a three dimensional spatial realm that endures through time (Dainton 2010:195). Any spatial location at one time is identified with a spatial location at a different time, and which location is which is an entirely objective matter. This implies that there is a single right answer to our dispute: if you remain located at the *same* location throughout the interval in which I seem to pass by, then you are right, you are at rest and I'm passing by. But if I remain at the *same* location throughout time, however, then I'm right, you move at a constant velocity and I'm at rest. A final possibility is of course that both of us are changing locations, so that neither of us is right.

But how can we figure out which of these three possibilities obtains? That is, how can I figure out whether my current location is the same as my location of some minutes ago? Here a difficulty arises. We cannot straightforwardly re-identify the locations across time. They do not come with little labels or special intrinsic properties, telling us which point is which. And we also cannot identify one location in virtue of its distance from another particular

location. For such relations to be of any help, we need to have already identified the second location in order to identify our location as being a certain distance from it, but the question is precisely how to re-identify any point at all at another time. Given that the points themselves are of no help, we naturally turn to the objects that have their trajectories through space; if we can determine directly which object is at absolute rest and which has an absolute velocity, we can use that to say which points are identical to which: if an object is located at x at t_1 and at y at t_2 and the object is at absolute rest from t_1 to t_2 , then $x = y$. But this only shifts the problem. We cannot tell from an object whether it is at absolute rest or has a constant velocity. Galileo brought this realization to light:

Shut yourself up with some friend in the main cabin below decks on some large ship, and have with you there some flies, butterflies, and other small flying animals. [...] With the ship standing still, observe carefully how the little animals fly with equal speed to all sides of the cabin. [...] When you have observed all these things carefully (though there is no doubt that when the ship is standing still everything must happen this way), have the ship proceed with any speed you like, so long as the motion is uniform and not fluctuating this way and that. You will discover not the least change in all the effects named, nor could you tell from any of them whether the ship was moving or standing still. [...] [T]he butterflies and flies will continue their flights indifferently toward every side, nor will it ever happen that they are concentrated toward the stern, as if tired out from keeping up with the course of the ship. (Galileo 1632/1967: 186-187).

Considerations like Galileo's have been taken to show that a system that is at constant velocity is intrinsically just like a system that is at rest.² The perspective that treats the ship at rest and the shore moving, and the perspective that treats the ship as moving and the shore at rest, each agree on

²Acceleration can of course be discerned from the way things are in the ship: when the ship sets off, things are pulled backwards. We are not concerned with acceleration, only with constant velocities.

what things are like in the ship *except for the velocities attributed to things*. So we arrive back at the very situation that we are trying to resolve: from your perspective I seem to be moving and from my perspective you seem to.

Each of our perspectives only differ with regard to what velocities are ascribed to things, and between us, there are equally good reasons to think that you are at rest and that you are not at rest. But *being at rest* and *having a constant non-zero velocity* are incompatible properties (just as having distinct constant velocities are incompatible properties), and so it's commonly assumed that something cannot have both. Given that we have equally good reasons to attribute either property to the same thing, we are led to a view on which things are neither, and hence a view that denies that locations at one time are genuinely identical to locations at another time. On this view, locations do not endure through time, contrary to what Newton assumed.

More formally, any two models of the positions of things across time that disagree only on which object is at rest and which moves (with a constant velocity) are deemed to be equally legitimate in the sense that we take there to be no substance to the disagreement between such models. It's important to be clear on how this approach of equating models works, so let us go through it slowly.

Let us consider first the case of a two dimensional Euclidean plane, a flat spatial surface. If we want to describe this plane, we first need to be able to refer to the locations on the plane. Of course, we want to name them in a convenient, systematic manner. For this we use Cartesian coordinate systems: the familiar grids with an x -axis and y -axis that are orthogonal to each other. Each location on the plane can then be associated with a point on the coordinate system. Such a point is itself just a pair of numbers, consisting of an x coordinate and a y coordinate. We refer to locations in the plane by way of these pairs of numbers. We could call a particular way of associating locations on the plane with points in the coordinate system a 'model' of that Euclidean plane, but it is more common to call them *reference frames* or simply *frames*.

It's important to sharply distinguish the Euclidean plane from a reference frame for that plane, to distinguish between the world and particular ways of

describing the world. There are many ways of placing a Cartesian coordinate system on an a plane, i.e. many different ways of associating locations on the plane with points in the coordinate system. This is to say that there are many equally legitimate reference frames for the plane.

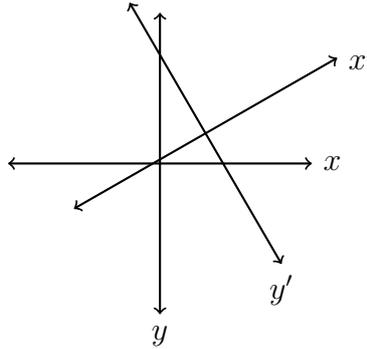


Figure 6.1: Two coordinate systems. In the figure, the $x - y$ frame differs from the $x' - y'$ frame on where the origin lies (i.e. which location in the plane is ‘named’ $(0, 0)$ in the reference frame) and they differ in their orientation (i.e. which straight paths in the plane lie parallel to the x-axis and y-axis).

We don’t think that there is one true way of laying a coordinate system over a plane, i.e. one true frame for a plane. The reason for this is simple: we do not think that anything in the plane itself corresponds to the definite orientation of the coordinate system, or that any location should be assigned one pair of numbers instead of another, that locations have ‘one true name’. These are all features of the relevant frame, not of the plane that the frame describes. And so we say that a statement describes a feature of the plane itself only when that statement holds just as well when we translate it into a description that uses another reference frame (i.e. into a description formulated in terms of another way of placing the coordinate system on the plane). For frames that differ in orientation (captured by the sine and cosine of the angles) and that are displaced with respect to each other (captured in terms of the constants, the n s), these translations are as follows:

$$x' = x \cos \theta + y \sin \theta + n_1$$

$$y' = -x \sin \theta + y \cos \theta + n_2$$

So, for example, if the angle θ between the axes of the x - y -frame and the x' - y' -frame is 30° and the location that is named $(0, 0)$ in the x' - y' -frame is named $(3, 2)$ in the x - y -frame, the translations are as follows:

$$x' = x \cos 30 + y \sin 30 + 3$$

$$y' = -x \sin 30 + y \cos 30 + 2$$

Using this translation scheme, one can take any point in the x - y -frame frame, and plug its x and y coordinates into these translations to find the corresponding x' and y' coordinates.

Given a class of frames that we deem to be equally legitimate descriptions of the same plane, we possess a set of translations between them, and we can use this set of translations to distinguish between statements that remain to hold true across translations, the so-called *invariant* statements, from those statements that do change with translation, the *frame-dependent* ones. To illustrate with an example from Maudlin 2011:31, the statement of a difference between two points in one frame, a and b , is not preserved when translated into another frame. If the x -coordinate of a point a in a frame is 2 ($X(a) = 2$) and the x -coordinate of another point b is 3 ($X(b) = 3$), then the difference along the x -axis between a and b , written as Δx is the following: $\Delta x = X(b) - X(a) = 3 - 2 = 1$. But say that we now translate the x and y coordinates to another frame and that the x' -coordinate of b is 6 and the x' coordinate of a is 4. *Their* difference is $\Delta x' = X'(b) - X'(a) = 6 - 4 = 2$ and not 1. So $\Delta x \neq \Delta x'$: when we translate the x and y -coordinates to x' and y' coordinates, their differences are not preserved. This means that the difference between points in a frame is a so-called frame-dependent notion. In contrast to difference, it can be checked that the quantity of $\sqrt{(X(b) - X(a))^2 + (Y(b) - Y(a))^2}$ remains the same when it is translated into the x' - y' -frame. This is the distance between points stated in terms of the Pythagorean theorem. The attribution of this quantity to a pair of points is an example of an *invariant* or *frame independent* statement.

So what is frame-dependent and invariant is relative to a set of translations, which are in turn just those translations that hold between those

frames that we have deemed to be equally legitimate. The relevant set of translations includes what it does *because* of what we deem to be equally legitimate reference frames. If we had thought that frames with different orientations cannot be equally legitimate because we think there is a true orientation to the world, then the translations would be different, fewer frames would be deemed equally legitimate and more quantities would turn out to be invariant.

Let us return to things seeming to have incompatible velocities across perspectives. When, instead of describing a simple plane, we describe things in space across time, we find the same decision points. To describe the motion of things, we need to describe the way things are in three spatial dimensions over time. So we move to four-dimensional coordinate systems, with one dimension representing time. We have one time axis, the t -axis, and three spatial ones, the x , y and z axes. As before, we decide which ways of associating locations in spacetime with points in the coordinate system are equally legitimate. On the Newtonian view (and the neo-Newtonian view we will look at shortly) there are various constraints on the alignment of the t -axis, most importantly, any two events occurring at the same time are to be associated with the same t -coordinate: when one frame assigns the same t -coordinate to two events and another frame assigns two different t coordinates, then these frames are not equally legitimate ways of describing things.

The points in the frame with different t -coordinates also stand in spatial distance relations to each other. In particular, given a point with a certain t and x , y and z coordinate, there is a point with a different t coordinate but the same x , y and z coordinate, which is to say that the second point lies at zero spatial distance from it. As there are distance relations between the points at different times, a vertical path through a frame (or better, a path that is parallel to the t -axis) represents something remaining at the same spatial location (i.e. having the same x , y and z coordinate) across time. See figure 6.2. A vertical path through a frame describes something at rest and stands in contrast to a diagonal path, which describes something changing location at a constant rate across time, i.e. represents something

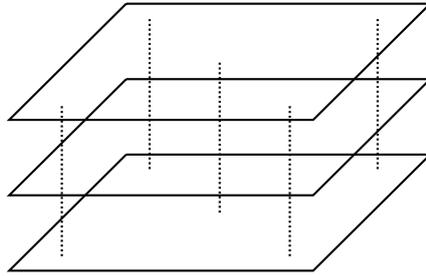


Figure 6.2: Newtonian spacetime. One spatial dimension is ignored in this diagram. The t -axis runs vertically. The stacked planes that lie orthogonal to the t axis are so-called hyperplanes and represent instants of time. The vertical dotted lines mark out points that lie at no spatial distance from each other (i.e. that have the same x and y coordinates, and only differ in t coordinates).

with a constant velocity. See figure 6.3.

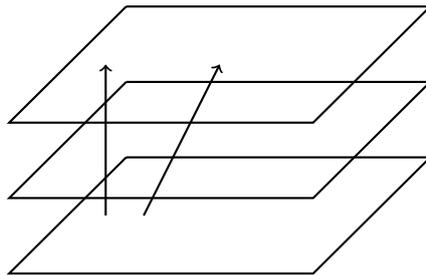


Figure 6.3: Paths in Newtonian spacetime. Vertical paths through hyperplanes describes something as remaining at the same spatial location through time, whereas a diagonal path describes something that changes locations.

If we think that there is an objective distinction between rest and constant velocity, one will think that grids that disagree about whether some object traces a vertical or diagonal path cannot both be legitimate. But, in light of the Galileon considerations we saw above, many think that there is no such objective distinction between rest and constant velocity. This has led to the view that frames that only disagree about whether a path is vertical or diagonal are equally legitimate - in much the same way in which frames that disagree about where the origin lies are equally legitimate.

As before, our class of equally legitimate frames provides us with translations, from one frame to another frame that represents things in an equally legitimate way. If we ignore orientations for the sake of simplicity, but allow displacements of the origin (captured by the cs) and let v_x , v_y and v_z be velocity in the x , y and z -direction respectively, and the cs capture displacements of the origin, the translations are as follows:

$$t' = t + n_1$$

$$x' = x - v_x t + n_2$$

$$y' = y - v_y t + n_3$$

$$z' = z - v_z t + n_4$$

Of course when we do not ignore different orientations of the frames, these translations will look more complicated.

Within our class of equally legitimate frames, we find frames that differ about whether an object traces out a vertical or diagonal path, i.e. about whether that object is at rest or has a constant velocity. That is to say, if we translate one frame to another using the above translation, a vertical path can be translated into a diagonal one. That some object is at rest is not an invariant statement. Rest and constant velocity fail to be invariant, as desired by those who think there is no distinction between a world in which an object is at rest and one in which it has a constant velocity, only different ways of describing it.

When we have this class of equally legitimate frames, we can distil the structure that remains invariant between them, which we will think of as the structure of spacetime itself, instead of the structure of a frame that happens to be used to describe that spacetime. It's as if we describe things in terms of reference frames with superfluous structure, for which we then need to correct by deeming pluralities of such frames equivalent in the relevant respects. The intrinsic structure of spacetime itself, according to the neo-Newtonian view, is one in which there are no spatial distance relations across time. Although there is no distinction between straight and vertical paths

(i.e. rest and constant velocity), there is still a distinction between straight paths and curved paths through time, corresponding to non-accelerated and accelerated trajectories. So although we distinguish between straight and curved paths within the structure of spacetime itself, we do not distinguish between different straight paths (such as between the vertical and diagonal ones, corresponding to absolute rest and constant velocities). We cannot draw a diagram of this structure as it is in itself, as this will unavoidably draw some path as vertical and others as diagonal (that is, our diagram is unavoidably done within only one of the equally legitimate frames). When we associate the locations in a Neo-Newtonian structure with points in our Cartesian coordinate system, we have an arbitrary choice to make about which path to align parallel to the t -axis in just the same way in which we have a free choice where to put to origin of the coordinate system.

Given this conception of the intrinsic structure of spacetime, we cannot say that an object really has a velocity of 10 kmh^{-1} , or that it traversed 10 km , or that it moves faster than some other object, or that it is at rest. All of these statements refer to frame-dependent features. Within spacetime itself, according to the neo-Newtonian conception of it, it's only the case out there that an object is non-accelerated (i.e. has a straight path), or is accelerating through spacetime (i.e. has a curved path). Of course, these are precisely those quantities that do make for intrinsic differences across systems and hence can be verified: the birds in Galileo's ship will all be pushed towards the stern to the extent that the ship is accelerating.

6.2 From Galilean to Minkowski spacetime

Imagine now that you're standing in the exact middle of the carriage of a moving train, and that you throw two balls at the same time with the same effort in both directions. As one ball hits the front of the carriage, the other hits the rear. I stand outside the train and observe the whole thing. From your perspective it seems that the two balls have the same velocity, since they had the same distance to travel and arrive at the same time. But from

my perspective, outside the train, it seems that they must have different velocities. After all, since the front of the carriage is itself moving forward, away from the ball that is thrown in its direction, the forward-moving ball has to travel more distance to arrive at the front of the carriage. As it still arrives at the same time as the other ball, it must have gone faster, that is, the forward-moving ball must have inherited the speed of the moving train when it was thrown. So, it seems, we have another case of the kind we saw in the previous section: the velocities that things seem to have differs between your perspective and mine.

But things aren't so simple. It was discovered that the speed of light is observed to be the same (roughly 300,000 kilometres per second) by any observer regardless of their own apparent velocity. This well-established physical fact is captured in the so-called light postulate:

Light postulate: the velocity of light is a constant, c , which it's measured to have regardless of the speed of the source emitting the light, and regardless of the speed of the one measuring it.

This is of course a very surprising fact. If you shoot a bullet in some direction, and chase it at great velocity, it will seem to move slower and slower the faster you are moving alongside it. If you managed to go just as fast, you would measure the bullet's speed as being zero. But not so in the case of light. Emit some light and it will move at the same velocity in all directions. Start 'chasing it' in one of these directions. No matter how fast you go, you will always measure the light to have the speed c , and it will not seem to go slower or faster, depending on how fast you go. This is surprising.

The constant speed of light has had a dramatic effect on the way we conceive of space and time. To bring this out better, consider another case. Instead of throwing two balls, you now flip the switch of a laser that sends light to the front and back of the carriage. From your perspective, the light has to travel the same distance to the front and back, and since it has a constant speed c , must arrive at the same time at both ends. But now, from my perspective, the forward moving light has to travel more distance, and since it still has the same constant speed c , arrives later than the light sent

towards the back. As before, there is no reason to privilege either one of our perspectives (indeed, this equality between perspective is typically stated as the second postulate of the theory, called *the relativity postulate*). And so we have a case of variance, not concerning the speed of light, but concerning the simultaneity of the light arriving at, respectively, the front and back of the carriage. As it turns out that it's a mistake to think that the velocity of light varies between our frames, we have the surprising result that things vary around it: the constant velocity of light shows the simultaneity of events to vary between our two perspectives on the matter.

This shows that something is wrong with the neo-Newtonian conception of spacetime. In one respect it takes too many things to be invariant, and in another respect it takes too few things to be invariant. As we saw in the previous section, the neo-Newtonian conception of spacetime accepts a definite division of the spacetime into hyperplanes. Two events that have the same t coordinate cannot legitimately be transformed into one where they have different t coordinates; rendering simultaneity between events invariant across equally legitimate frames. On the other hand, the neo-Newtonian conception of spacetime allows too many transformations given that under the neo-Newtonian conception no velocity is invariant across frames, rendering *all* constant velocities variant, including that of light, and that is wrong.

The velocity of an object is the distance traversed divided by the time taken, that is, if something travels only in the x -direction, from a to b , then the speed is $\Delta x/\Delta t$. In the case of light (and light only), we want this value to remain unchanged when we translate one frame into an equally legitimate one. For it to remain unchanged under translation, there are only two options: either Δx and Δt must both remain unchanged, or they must change in such a proportion to each other that $\Delta x/\Delta t$ remains the same (Maudlin 2011:41). We can rule out that both Δx and Δt remain unchanged under translation given that for example, the forward moving light covers more distance from within my preferred frame than it does in your preferred frame. So, if we imagine the light to travel in the x -direction, then Δx , the distance travelled, differs between equally legitimate frames. For the speed of light to remain nonetheless constant, there must therefore

be a proportional difference in the duration of the light's journey, i.e. a proportional difference in Δt . That is to say, different frames disagree about the duration of a process or journey.

For certain pairs of events, different frames can disagree about whether they are simultaneous or not, and when not simultaneous, how much time lies in between them. They cannot disagree in this way on all pairs of events however, only for some of them. Take a certain event, e , and take a frame such that e lies at the origin of the frame. Say that the event is the emission of some light. The surface of the light emitted moves in all directions, expanding with time, at the speed of light. As the surface of the light expands through time, it draws out a cone shape through time. We call this the future light cone of the point at which e is located; there is a similar sort of cone for the light incoming at e 's location. This divides the frame up into different areas for each point, which stands in crucial relations to each other. Points that lie on the trajectory of the light travelling from e are said to be 'light-like' separated from e . Points inside the cone are said to be 'time-like' separated from e . We can think of these as being far away in time whilst close apart in space. Points that lie outside the light cone are said to be 'space-like' separated from e . We can think of these points as relatively close in time whilst at a large spatial distance.

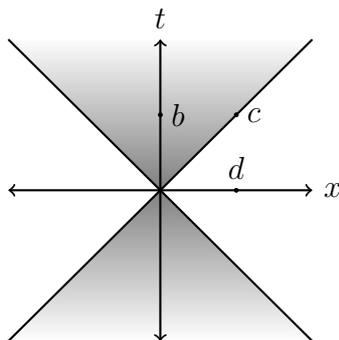


Figure 6.4: Lightcone. The points in the top cone are in the absolute future of the origin, the points in the bottom cone in the absolute past. The origin and b are time-like separated; the origin and c are light-like separated; and the origin and d are space-like separated.

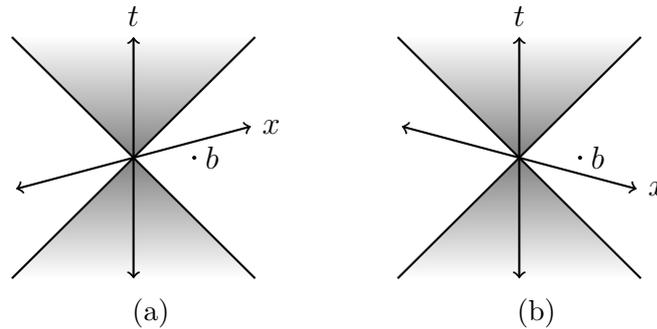


Figure 6.5: Equivalent frames. According to frame (a), what happens at point b occurs before what happens at the origin. According to frame (b), what happens at b occurs after what happens at the origin.

Using these three relations amongst points - light-like, time-like and space-like separations - we can (roughly and crudely) specify what sorts of frames are equally legitimate, and which ones aren't. First of all, for any two space-like separated events, two frames that only disagree on whether they have the same t -coordinate or different ones, are equally legitimate. Equally legitimate frames can differ in how they draw simultaneity planes through the area outside the light cone. (See figure 6.5).

Secondly, for any two events that are time-like separated events (two events that lie within the light cone), two frames that disagree only on whether they have the same spatial coordinates or different ones, are equally legitimate.

Thirdly, any two events that are time-like separated in one frame must also be time-like separated in any frame that is equally legitimate, i.e. agreement on which events are time-like separated is a necessary condition for two frames to be equally legitimate. Equally legitimate frames cannot disagree about where one event lies with respect to the light cone of another event.

As before, this gives us a class of equally legitimate frames, and using this class of frames, we can determine translations that take one frame into one that is equally legitimate. Say for example, that one frame reflects a mode of description naturally associated with the perspective of having someone having a velocity v in the direction of the x -axis as seen from a second frame,

then the following translations hold (where c stands for the velocity of light):

Where $\gamma = \sqrt{1 - v^2/c^2}$:

$$t' = \gamma(t - vx/c^2)$$

$$x' = \gamma(x - vt)$$

$$y' = y$$

$$z' = z$$

These are known as Lorentz translations. As in the case of the neo-Newtonian framework, with the class of frames (or, alternatively, with the translations between them) we have a distinction between quantities that are invariant and those that are frame-dependent. Two important invariant matters are the following:

Absolute time order is invariant: given an event e , it's an invariant matter that all events in its future light cone come after it, and all events in its past light cone come before it.

Spacetime interval is invariant: the quantity given by $I^2 = \Delta x^2 + \Delta y^2 + \Delta z^2 - c^2 t^2$ is invariant. This is known as spacetime interval. The interval I^2 between any two light-like separated events is 0.

Many more quantities are now frame-dependent however:

Simultaneity is not invariant: two events that are space-like separated can be simultaneous in one frame and non-simultaneous in a different but equally legitimate frame.

Duration is not invariant: an event can be assigned different durations by equally legitimate frames.

Velocity of anything but light is not invariant: although the velocity of light is kept constant, different velocities are assigned in different frames.

Length is not invariant: the faster an object goes in a frame, the shorter it is.

Shape is not invariant: as one frame can move in only one direction relative to another frame, objects are shorter in that direction only, thus differing the shape attributed to the object.

Inertial mass is not invariant: The faster an object goes in a frame, the more mass it has.

In the case of length and mass, the use of ‘the faster an object goes ...’ can be misleading. All it means is that, in the case of length, in a frame that assigns a greater velocity to an object than another frame, the object will also be assigned a shorter length. It does not mean that the object’s going faster causes the object to be shorter, it just means that its description as going faster (than it does in another frame) goes together with its description as being shorter (than it is in the other frame). In fact, all the frame-dependent properties vary with each other across frames, velocity is in no way special.

Given the specified class of equally legitimate frames, we can again distil the structure that remains invariant between them, which we again take to be the structure of spacetime itself, instead of the structure of a frame that happens to be used to describe that spacetime.³ When we do this, we arrive at a structure that is known as Minkowski spacetime.

6.3 Unity, invariance, and antirealisms

In our discussion of colours we encountered the paradox of conflicting appearances: something can appear to have incompatible properties from different

³Minkowski himself followed the so-called ‘Erlangen program’ (introduced by Felix Klein), where we classify and characterize geometric structures in terms of the ways in which they remain the same under sets of transformations; see Gray (2005). Nowadays, it is more common to specify the Minkowski structure in terms of a collection of vectors, a vector space with certain properties. The vector presentation makes it much more difficult to see the various motivations that generate the Minkowski conception of spacetime, so is not very suitable for a discussion of a view that aims to unearth these motivations and resist them.

perspectives, while none of these perspectives is privileged. The arguments that drove physicists towards the Minkowski conception of spacetime are of the same form and similarly depend on the unity assumption, the assumption that all facts co-obtain. Because of the emphasis on frames, however, it can be hard to see what metaphysical conclusions are drawn within a Minkowski conception of spacetime. In this section I will argue that they are indeed the same sort of conclusions as we saw in our discussion of colours, in particular, they are antirealist conclusions concerning the properties that fail to be invariant across frames.

Consider again the argument concerning rest and constant velocity. To you, you appear to be at rest whereas to me, I appear to be at rest. *If* there is one true answer to who is at rest and who moves at a certain speed, we saw that we have no means of finding out what it is. To conclude from this that there is no such thing as absolute velocity and absolute rest is of course rather quick. At the time the relativity theory was developed, verificationism reigned widely amongst working physicists (in particular Einstein). We are now less worried about unobservables if they fit an explanatory niche. Whether absolute rest and constant velocity fit such an explanatory niche is a complicated question, depending on what sort of things need explaining, what sort of explanations count, and whether there are other explanations that require a distinction between rest and constant velocity in the background. But, in addition to the verification issue, there is the arbitrariness worry: to resist the move to a Minkowskian conception of spacetime (and yet hold on to the unity assumption) we would have to arbitrarily privilege one of our perspectives over the other.

Note also that *absolute* rest and *absolute* velocity are unobservable should not be confused with the idea that no rest or velocity is ever observed. On the contrary, the whole argument runs on the basis of the fact that what appears to have a velocity to you appears to be at rest to me. We observe rest and velocity. The problem is not that things do not appear to differ in terms of their velocities whilst Newton predicts that things do differ in such ways, the problem is that incompatible velocities are observed without there being any non-arbitrary reason to choose between them.

It's also important to note that things do not merely seem to move or be at rest within human experience. We can measure velocities using all kinds of instruments. If you measure the distance travelled by the ball you threw on the moving train (the distance from the middle of the carriage to the front) and I measure the distance travelled by the ball (the distance from the middle of the carriage of the train + the distance travelled by the front of the train between the moment you threw the ball and the moment it arrived there), our measurements will have different results about the distance covered by the ball. The readings of these measurements do not in essential way depend on human experience however; they would be the same if we weren't there to do the measuring. The conflicting verdicts here are not due to some quirk of human perception. Again, we have perfectly good reasons to think that things are at rest, or move with certain speeds; we only find equally good reasons to attribute these incompatible states to objects. That is the real problem motivating the neo-Newtonian conception of spacetime.

We find a similar situation when we turn to the arguments that motivate the Minkowskian conception of spacetime. Consider again the case in which light is emitted from the middle of a train. From your point of view in the train, the light reaches the two endpoints simultaneously; from my point of view, the light reaches the front before it reaches the rear. We can measure this. Your measurements will show the events to be simultaneous, my measurements will disagree. It's not that things fail to appear to have any temporal order from our perspectives; rather, they appear to have incompatible ones. And similarly for the other properties: we assign incompatible lengths to things, incompatible shapes and incompatible masses to objects, and incompatible durations to processes and journeys. It's not that we have been given no reason to think anything to have any of these properties; the situation is that we have equally good reasons to attribute a whole range of incompatible properties to things.

And it's here that the unity assumption kicks in. Something cannot be at rest insofar as it's moving, and hence - because it's assumed that any two facts that obtain co-obtain - rest and velocity are demoted to frame dependent properties of things, together with the length of things, their mass, shapes,

etc. - properties that are deemed on a par with the origin (0,0) assigned to a point. The unity assumption is at work to devastating effects in the background of the Minkowskian conception of spacetime; squeezing out any property that admits of variance across the frames we deem to be equally legitimate.

But what metaphysical status do the frame-dependent quantities have exactly? Minkowski himself famously pronounced that ‘henceforth, space by itself, and time by itself, are doomed to fade away into mere shadows, and only a kind of union of the two will preserve an independent reality’ (Minkowski 1964/1908: 297). But what is it for space and time to fade into mere shadows, and what is for their union (spacetime intervals) to have an ‘independent’ reality?

It can seem that the Minkowskian conception is ambiguous about its antirealist implications. One finds this ambiguity in many discussions of relativity. Consider this passage by Maudlin:

A test of how well one has internalized the spirit of Relativity is how one is tempted to answer the question: what would it be like to travel at 99.99 percent of the speed of light? [...] The truly relativistic answer: right now you are travelling at 99.99 percent the speed of light – in some perfectly legitimate inertial frame. It is no more correct to say that you are now at rest than that you are infinitesimally close to moving at light speed. Indeed, except for things travelling at light speed, it makes no sense to attribute any absolute velocity to any object. Not being invariant, sublight velocities are not real properties of objects. There can therefore be no dynamical effect of ‘traveling near the speed of light’ because there is no such objective state as travelling near the speed of light. Things don’t shrink or slow down. Rather, there are always infinitely many ways of expressing space-time intervals in terms of distance in space and elapsed time. None of these ways, represented by the various inertial frames, is any more valid, or less valid, than any other. (Maudlin 2011: 53).

Within this truly relativistic answer, the story wavers. You are to say that ‘right now, you *are* travelling at 99.99 percent of the speed of light in some perfectly legitimate frame’ but at the same time that ‘sublight velocities are not real properties of objects’ (i.e. that you do not *really* move at any velocity at all, not even at a zero velocity). Indeed, you are to say that ‘except for things travelling at light speed, it makes no sense to attribute any absolute velocity to any object’ but then, on the other hand, attributing such velocities is meant to be one of the ‘infinitely many ways of expressing space-time intervals’, which *are* real properties of things.

So, is it false to say that something is at rest or has a certain sublight velocity, or not? The ambiguity arises because we do not discuss the world directly, but rather different models of the world, namely the different frames. When we model things, we distinguish between those aspects of the model that are representative of the world, and those aspects that aren’t. Say that, to model a car crash, I use a bunch of wooden toy cars. The shape of these toy cars is taken to be representative of the real shapes of the cars, yet the fact that they are made out of wood is clearly not meant to represent that the cars themselves are made of wood. The model doesn’t simply consist in the wooden cars, it consists of them plus a set of intentions about which aspects of the cars represent things about the situation out there and which aspects don’t. So we cannot even say that for the statement ‘the cars are made of wood’, our car model represents a world in which it is true. The wooden car model doesn’t represent a world in which the cars are made of wood, it simply represents a world in which cars have a certain shape and stand in certain configurations over time, and it is otherwise representationally silent, like an incomplete specification of a possible world. Of course, we could *decide* to use those same cars to now represent a world in which the cars are indeed made out of wood, and in such a model, the statement ‘the cars are made out of wood’ is indeed represented as being true. And only in that case can the model fail to represent things as they are. So could it similarly be that when, within a certain frame, something has a straight path, the Minkowskian conception merely requires that we take this aspect of the frame not to be representative of how things are? Just as we can model the

same car crash using wooden cars, aluminium cars, big cars, small cars, so we can model the same spacetime with a model in which you have a vertical path, one in which you have a diagonal path, etc. Your vertical or diagonal path in a frame isn't meant to be representative of anything. Any such frame is simply taken to be representationally silent in that specific respect. Such a convention falls short of an explicit metaphysical antirealism about the sublight speeds of things - one might think. The Lorentz translations are in that case merely ways of precisifying what aspects of the frames are representative of the world and which aspects aren't.

Though this might be a sensible way of metaphysically deflating what the physicists are doing, it's ultimately of no help with regard to the metaphysical issues themselves. Even though, when I use wooden cars to model a car crash, their being made of wood is not representative of the cars and hence my modelling doesn't imply anything with regards to the make up of the cars, the question remains whether the cars are made of wood or not. Leaving a question open doesn't mean that it doesn't have an answer. Even if we say that the verticality of paths in a frame is simply not representative of anything, we are still left with the question whether objects really have sublight velocities or not.

Maudlin remarks in the quoted passage that 'it makes no sense to attribute any absolute velocity to any object'. Perhaps then, one might think, the question whether an object does or doesn't have sublight velocities simply fails to make sense in a Minkowski spacetime. What this requires though, is that the attribution of sublight velocity results in utter gibberish that is not truth evaluable, and I don't see any good reason why this should be so. We experience things to have sublight velocities, and it is clear that this is not to experience things as having some invariant property - given that we experience different things from our perspectives. When I sit in a train, I can experience the other train as moving away and then, after a Gestalt switch, experience my own train as moving away. There is a difference in the way I experience things before and after the switch (hence it cannot be an invariant quantity that I am experiencing), and for each of these ways I experience things, it makes perfect sense to imagine things as they appear to

be. Are they this way or that way? This is a perfectly sensible question that distinguishes between possible worlds on the basis of the sublight velocities of things. To deny that the question makes sense is to hide our heads in the sand.

As in the case of colours, one might try to resist the full blown antirealist conclusion that things don't really have real lengths, shapes or masses, etc. by relativizing the various properties. But what are we to relativize to? The obvious candidate would seem to be frames. But we have to remind ourselves at this point that frames are just descriptive tools, enabling us to describe trajectories and other properties using coordinate systems. To say that something can still have a certain length- L -in-frame- f is a little like saying that the rose is red-relative-to-English but not red-relative-to-German (instead being rot-relative-to-German). The relativized properties end up denoting mere ways things happen to be described within some descriptive system. It offers no way of avoiding the antirealist conclusions enforced upon us; just as the fact that the point at the tip of my nose lies at the-origin-of-frame- f does not confer any substance on the property of *being the origin*, it means that the point is *described* or *referred to as* the origin (0,0) in some frame and not in some other frame.

So, again, do objects have sublight velocities, or not? In Minkowski spacetime, the only answer can be that objects don't have sublight velocities, nor lengths, nor masses, nor shapes, and that events do not have definite durations or occur in a definite order. Things only have properties that are invariant, such as spacetime interval. If they did really have any of the properties that differ across frames, then any two frames that assign me a different velocities, different lengths, etc. cannot be equally legitimate, and hence fewer quantities would turn out to be frame dependent. Though we can treat the frames as equally legitimate by treating the frame-dependent quantities to be bits of structure that isn't representative of how things are, this is ultimately because - on the Minkowskian conception of spacetime - there is not anything corresponding to that structure out there. The only metaphysical conception that the Minkowskian spacetime leaves room for is one that has wide-ranging antirealist implications and takes the world to be

radically different from how it appears to be.

This is deeply problematic, and not just because it denies that which we have ample empirical reason to believe in. Some of our very best theorizing about the world, indeed, almost all of our other successful theories about the world presume that things have definite lengths, shapes and masses, and that events occur in a definite temporal order. Examples are legion outside of physics: biology, geography, chemistry, psychology, economics, etc. all treat in objects of certain shapes and sizes interacting with each other in a patterned way over a definite course of time. It might be tempting to think that the Minkowskian can say that all of these are approximately correct or that they are correct enough because they deal with things that have very low velocities. But the Minkowskian cannot really say that. The Minkowskian conception of spacetime implies that all these theories presume a background theory of properties that just fail to be instantiated. There just aren't objects such as those presumed by any of these theories and there just isn't the temporal progression assumed by them. This is even in tension with successful theories found within physics itself: almost all interpretations of quantum mechanics, whether they be collapse based or non-collapse based interpretations, fail to be invariant across frames because of the role played by simultaneity in them.⁴ Quantum mechanics is no less successful than the special theory of relativity, and yet it is naturally understood to predict incompatible results across different frames. The antirealism with regard to most of the content of the manifest image isn't just problematic because it denies what is outright manifest to us and what we have good reason to believe in but, more importantly, it is problematic because our most successful theories only apply to a world that is largely the way it manifests itself in our experience. And yet, for all this, it is an established fact that light has a constant velocity, and undeniable that incompatible quantities are assigned in different frames of reference, meaning that things simply *cannot* have ve-

⁴For clear discussions, see Putnam (2005:631) and Maudlin (2011) and Cushing (1994: §10.4.2). I say '*almost* all interpretations' because Tumulka (2006) has proposed a version of the GWR spontaneous collapse theory – 'flashy GWR' – that reconciles relativistic space-time structure with the violation of Bell's inequality. For a critical discussion of this, see Maudlin (2011: Ch.11).

locities, and be simultaneous, etc. - at least, not *if* the unity assumption holds.

6.4 A fragmentalist account of relativity

It should be clear how the fragmentalist framework invalidates the arguments used to motivate the Minkowskian conception of the world: if there is no good reason to think that some object has only one of a pair of incompatible properties (such as incompatible velocities), it does not follow that it must have neither of them, at least not if we reject the unity assumption. Of course, that the argument is not valid does not of itself mean that the conclusions drawn are mistaken. But since the conclusions drawn are implausible and conflict with most of our other successful theories of the world, we have strong evidence in favour of a fragmentalist account.

Start again with the simplest case, used to motivate the Neo-Newtonian view. You experience me moving and yourself at rest and I experience you as moving and myself at rest. Between us, we have good reasons to think that you are *both* at rest and moving by. You are not at rest *insofar as* you are moving by, that cannot be the case given the incompatibility of these facts, yet you are at rest *and* you are moving at some non-zero velocity (indeed, you are moving at a whole range of velocities). The same goes for me: I am at rest and moving, however, you are at rest *insofar as* I move, and I'm at rest *insofar as* you move.

Similarly in the case in which light is emitted in a moving train: insofar as you are at rest in the middle of the carriage next to the light source, the light reaches the front and back of the train simultaneously. Insofar as you are moving and I'm at rest observing things outside the train, you have a constant velocity and the light reaches the front before it reaches the back. So the light reaches the front and back, both simultaneously and one after the other. The light reaches the front and back simultaneously *insofar as* it doesn't reach them one after the other. And the same goes for the length, shape and mass of objects, and the duration of events. So the following sort

of facts are the case in the moving train scenario (where you are observer A and I'm observer B):

A is at rest $v_A = 0$ ◦ the train is at rest $v_{train} = v_A = 0$ ◦ the light reaches the front and back simultaneously ◦ the time taken for the light to reach the front is 1 clock time ◦ the distance travelled by the light 1 measuring rod ◦ A has mass m_1 ◦ B is moving by the train at a velocity v_B towards the back of the train ◦ ...

A moves at a velocity $v_A \neq 0$ ◦ the train is moving at a velocity $v_{train} = v_A$ ◦ the light reaches the front before it reaches the back of the train ◦ the time taken for the light to reach the front is less than 1 clock time ◦ the distance travelled by the light that reaches the front is less than one measuring rod ◦ A has mass $m_2 \neq m_1$ ◦ B is at rest $v_B = 0$ ◦ ...

Call these two descriptions $Frag_1$ and $Frag_2$. Each of these corresponds to the way things are described within one of the many equally legitimate frames of reference. Our metaphysically complete description of a certain scenario is the conjunction of these and other such descriptions:

$$Frag_1 \wedge Frag_2 \wedge \dots$$

We describe a certain pattern of properties as co-obtaining, and a different pattern of properties incompatible with the first, and so on.

As before, it is not implied that my being at rest 'makes' or causes things to be one way and my moving makes or causes things to be another way. All we have are collections of properties instantiated together in patterned ways: your being at rest goes together with things having certain velocities, lengths, shapes and masses whereas your moving at some velocity goes together with things have certain other velocities, lengths, shapes and masses.

To be absolutely clear about the implications of this view: you are, right now, really travelling at 99.99 percent of the speed of light, where this is understood to be a full-blown property of you. This velocity of yours is not some mode of presentation, it is really one of the ways you are. Of course,

you do not move at 99.99 percent of the speed of light *insofar as* things have the velocities that you currently experience things as having. Things have those velocities *insofar as* it is entirely false to say that you are moving at 99.99 percent of the speed of light. And yet you are moving at that speed, and *insofar as* you do, the things have different velocities from the ones they have insofar as you are moving at lower velocities. This is not to say that you have just *any* velocity; in particular, no ordinary object has a velocity exceeding the speed of light. The fact that light has velocity c co-obtains with you having any of the velocities that you do have. Note also that there is no longer the odd asymmetry of light having a certain velocity, whilst nothing else has a velocity at all. Everything moves. And everything has a length, shape and mass, and is involved in events of a certain duration, occurring in a definite order. The surprising fact that we learn from *the light postulate* is not that things aren't any of these ways; the surprising fact is that they are more than one such way.

Things also still have the quantities that are invariant on the Minkowskian conception, such as spacetime interval. These quantities have the nature of complex ratios of quantities which are such that, as the quantities differ across frames, their ratios remain the same. Given that, on the fragmentalist conception, things genuinely instantiate the properties that the ratios are ratios of, they straightforwardly instantiate the ratios themselves. Indeed, the objects have the invariant quantities *because* they have the varying quantities in patterned, law-governed ways.

Measurements in general are individuated in terms of what is measured. If I move at great velocity (insofar as you are at rest), and we use the same measuring rod and clock (or rods we earlier compared, and clocks we synchronized) to measure some distance and some other velocity, we obtain different readings. This is just to say that the fact that these things are measured to be these ways is itself is just one of the varying facts: this measurement co-obtains with these facts, and that measurement co-obtains with these other facts. To the extent a measurement is accurate, it co-obtains with the velocities, distances, etc. that things are measured to have.

On the Minkowskian picture we make a certain measurement and then ex-

tract the invariant quantities from those readings, denying that the measured velocities, lengths, etc. have themselves any reality. In the fragmentalist picture we make a measurement, and take things to be exactly as measured. From this, using the Lorentz translations, we can predict all the other velocities, lengths, etc. that things must have *as well*. The Minkowskian uses the translations to arrive at invariant quantities from the measured quantities, leaving these measured quantities behind; the fragmentalist uses the Lorentz translations to arrive at predictions of further facts, on the basis of the facts observed.⁵ The Minkowskian treats the class of equally legitimate frames as a specification of an intrinsic structure out there; the only thing that is really out there. On the fragmentalist conception, different descriptions in terms of different frames are just incomplete descriptions, capturing only a subset of the facts, namely those that correspond to one of the ways things appear to be to a possible observer in space time. I think it has always already been tempting to understand the frames in this way.⁶

⁵To be sure, when we use the Lorentz translations to predict other facts, we do not take these other facts to be mere redescriptions of the facts we observe. Consider this passage by Sklar:

If we now associate the real (for an observer) with the simultaneous for him, we must, accepting the conventionality of simultaneity, accept as well a conventionalist theory of ‘real for’. It is then merely a matter of arbitrary stipulation that one distant event rather than another is taken real for an observer. (Sklar 1981/1985: 297).

If it’s merely conventional how we take spacetime to be foliated, then it’s merely conventional what we take to co-obtain with the experience of a certain observer and hence whether the experience is veridical or not. But, of course, the fragmentalist does not accept the conventionality of simultaneity, nor the conventionality of co-obtainment. What co-obtains with what are contingent facts like any other, and special relativity theory is taken to be concerned with certain law-governed patterns in them.

⁶Take for example Prior:

I suspect that the infinity of local proper times which figure in relativistic physics amount simply to what appears from the various points of view, or in various ‘frames of reference’, to be the course of events. And given how the course of events appear from a certain point of view, your relativistic physicist will be able to calculate how it appear from certain other points of view. He can also indicate what features of the course of events (what temporal orderings of those events) will be common to all points of view. (Prior 1968/2003b: 136)

Prior, himself a presentist, thinks physicists are only able to calculate how things appear

The Minkowskian conception builds the speed of light into the geometry of spacetime and explains things in terms of this structure. The fragmentalist conception does not attribute this structure to spacetime. It thus offers different explanations of things, competing with those of the Minkowskian. Given this difference in explanation, one may worry that this is scientifically and not just metaphysically revisionary. But although the accounts compete in terms of the explanations given of things, the formalisms are all the same, and the central postulates are accepted by both. As far as I can see, the conflict can only be settled in terms of what explanations ultimately turn out to be better.

One might object that the Minkowski conception of spacetime is superior in precisely this regard. The Minkowskian has the speed of light built into the very structure of spacetime in such a way that the structure explains the constant velocity, whereas the fragmentalist must accept the brute fact that light always has a certain speed insofar as other things have different speeds. But such an objection is misguided. The Minkowskian does not so much explain the constant velocity of light by some geometrical fact as simply understanding the constant velocity of light to *be* a geometric fact. And it's a similarly brute fact that spacetime has the geometric structure that it has, and not some other structure. The light postulate is a brute fact for everyone, no matter how it is understood.

When it comes evaluating explanatory values of theories, moreover, one cannot myopically stare at one single theory but needs to consider the overall conception of the world in which it is to fit, its integration with other successful theories. If the Minkowskian picture undermines what other successful theories presume, for example that things are genuinely simultaneous, it undermines our explanation of things elsewhere. The fragmentalist account takes things to be as they are presumed to be in our other theorizing about the world and can therefore combine with such theories in making sense of things. The standard Minkowski conception conflates the 'merely linguistic' difference between different ways of naming the locations in a space, and the

from different points of view and not how they truly are; on the fragmentalist account, the physicists *are* thereby calculating how things truly are.

qualitative differences of objects between different perspectives. The principle that only what is invariant across such perspectives is real works as a metaphysical filter, filtering out any properties that differ across perspectives, with the result of crudely thinning out the properties that objects instantiate and that we can legitimately resort to in theories about the world.

The fragmentalist account does not contradict any of the physical facts that support the Minkowskian conception, such as the constant speed of light, it only contradicts the metaphysical conception of the world that is suggested by it, and which depends, not just on physical observations, but on metaphysical background assumptions, in particular the unity assumption.

6.5 Relativity, tense and passage

The fragmentalist account of the passage of time that I proposed in §3.2 implied one foliation of spacetime into a set of hyperplanes. The A-theory also implies such a foliation. The Minkowskian understanding of spacetime is inconsistent with both the A-theory and the passage theory, given that it takes spacetime not be intrinsically foliated in any way. The fragmentalist conception of spacetime, on the other hand, is in tension with the standard A-theory but perfectly compatible with the passage theory.

According to the passage theory of time, reality consists in the passage of facts. When one state passes over into another state, it's not the case that all facts that are part of the one state co-obtain with the facts that are part of the next state. And this suffices for a metaphysical foliation of spacetime into a collection of states, which can be represented by Newtonian hyperplanes.

According to the A-theory, only what is currently the case is the case at all. As I snap my fingers, everything that is the case throughout the universe at this very moment is all that is the case. Given the speed of light, we will only be able to know some time into the future what was the case at far-away places when I snapped my fingers, namely when the light of those far-away places finally reaches us; but that does not take away from the fact that

those things are the case right now. There is thus a fact of the matter which things are simultaneous with the event of your reading this sentence: they are simply all the facts that obtain. And so, the A-theory also implies a metaphysical foliation of spacetime into simultaneity planes.

But spacetime isn't foliated into simultaneity planes according to the Minkowskian conception of spacetime. As we saw above, given a certain location a in spacetime, there are those points that lie in its absolute future and those that lie in its absolute past, but then there are still all those points that are space-like separated from a . These points come neither at the same time as a , nor before, nor after. These points lie in the 'absolute elsewhere' of a , as it is sometimes called. As we discussed, any structure 'missing' from Minkowski spacetime is structure that isn't there according to that conception: it cannot be 'added' without abandoning the very structure attributed to spacetime on the Minkowskian conception. So the Minkowskian conception of spacetime is incompatible with any theory that implies a metaphysical foliation of spacetime into simultaneity planes, and that includes both presentism and the fragmentalist account of passage.⁷

Things are different when we turn to the fragmentalist account of relativity. Here we do not deny that spacetime is foliated, but we rather take it to be foliated in multiple ways (which do not, indeed cannot, co-obtain). If we combine this account of relativity with the passage theory, the fact that a certain state passes over into another state turns out to be just like the fact that things have a certain velocity, length, shape, mass, etc. Insofar as I am at rest, state A passes into this other state B, but insofar as I have a certain velocity, a different state A* passes over into this other state B*. Although A passes into B and A* passes into B*, A doesn't pass into B *insofar as* A* passes into B*. The combination of these views is thus a natural fit, requiring no revisions in our account of passage or our account of relativity.

⁷Putnam offered a slightly different argument, showing the inconsistency of presentism and special relativity, in Putnam (1967). The argument is criticized in a much-discussed paper by Stein 1968, and defended (to my mind successfully) against Stein's criticisms in Saunders (2002). Presentist-friendly views on the issue are offered in Prior (1970), Tooley (2000), Godfrey-Smith (1979) and Monton (2006). Particularly helpful general discussions of various present-friendly views are found in Hinchliff (2000) and Savitt (1998).

The passage of time are simply not amongst the invariant matters, just as simultaneity and duration aren't amongst the invariant matters.

But not so in the case of standard A-theories. If spacetime is foliated in many ways, this is just as incompatible with the A-theory as its being not foliated at all. The A-theory does not only imply that spacetime is foliated, it also makes the further claim that only one simultaneity plane is real, i.e. that just *these* events *are* the case, and just *those* events *were* the case, and so on. The A-theory implies that spacetime is foliated in exactly one way, namely the way on which there is exactly that one simultaneity plane, one definite collection of events that are now the case. If spacetime is indeed foliated in many ways, as it is according to the fragmentalist account of relativity, different simultaneity planes must be genuinely real, and this is straightforwardly incompatible with the standard A-theory of time.⁸

So, although the standard A-theory and the passage theory are in the same boat with regard to the Minkowskian conception of spacetime, they come apart in the case of the fragmentalist conception of relativity. With the fragmentalist conception of relativity, we have a way of squaring a genuine objective passage of time with relativity, but we do not thereby have a way of combining the standard A-theory with special relativity.

6.6 Summary and concluding remarks

Objects do have definite lengths, shapes and masses, and objects move at definite velocities. Events have durations, and occur in a definite temporal order. These properties are distributed over things in multiple states of the world, none of which obtain together. Within one such state the history of the world looks like this, within an another such state the history of the world looks like that instead.

Of course we have focussed solely on the special relativity, and different issues come into play when we discuss other theories, in particular the

⁸What about the fragmentalist A-theory proposed in §2.4? I believe that this *is* compatible with the fragmentalist account of relativity that is offered here. See Fine (2005b: §10).

theory of general relativity and theories of quantum gravity. It has been argued that general relativity has models that cannot be foliated at all (see Gödel 1949/1990). At the same time, it has been argued that quantum gravity has models with one preferred foliation (see Monton 2006). We should not think that physics offers us an unambiguous picture of the world which we can straightforwardly use to constrain our metaphysics. At the very least, I hope to have made clear how the fragmentalist framework changes the way in which the physics feeds into our metaphysical conception of reality, and changes the extent to which the scientific advances threaten to make a hash of the sort of objects and properties that are manifest in our experience of the world and that ultimately underlie a wide range of explanations of things.

Chapter 7

General Discussion: Perspectival Variance

We have now explored various ways in which the world might be fragmented. In this final chapter, I want to step back from the particular applications of the framework and discuss the bigger picture. The various cases discussed clearly share a common structure, which we can use to systematize the application of fragmentalism. When we make this general structure more explicit, we come to see two further reasons in favour of fragmentalism, and can show why certain general worries are unfounded.

The chapter is divided in three sections. Section 1 discusses what makes a phenomenon a case of perspectival variance and systematizes the various conceptual strategies that can be employed in response to such cases. Section 2 discusses whether there are further instances of perspectival variance and whether their resolution through fragmentation does not over-generate all kinds of implausible views. We will see that fragmentalism can in fact promise a uniform treatment of perspectival variance, whereas the standard strategies cannot. Section 3 discusses the worry that fragmentalism undermines the importance of objectivity in our inquiries concerning the world, or undermines certain lessons that we should learn from science. I will argue that these worries are unfounded. An appendix discusses what may seem a less costly alternative to fragmentalism, namely metaphysical relativism,

arguing that such a view is conceptually unstable.

7.1 Perspectival variance

It should be clear that the cases that we have discussed in the previous chapters share a common structure. Roughly, they are all cases where something seems to have incompatible properties from different perspectives even though neither perspective is privileged over the other. I will call this phenomenon *perspectival variance*. It can be defined more precisely as follows:

Perspectival variance: We have an instance of *perspectival variance* when: (1) under one set of conditions C_1 it appears to an observer that A , and (2) under another set of conditions C_2 it appears to an observer that B , where (3) ‘ A ’ and ‘ B ’ state incompatible facts, and (4) neither set of conditions, C_1 or C_2 , is worse than the standard conditions of observing things to be the ways they appear, nor is one set epistemically privileged over the other.

Typically the incompatible facts are such that the same object is observed to have one property under one set of conditions and an incompatible property under a different set of conditions, or they are such that a pair of objects is observed to have one relation and a relation incompatible with the first. By ‘conditions of observation’ I mean some fact that obtains and which is such that, if it didn’t obtain, the content of the observation would be different. If I observe a blue sky, the fact that I do so under certain lighting conditions, at a certain time, using a certain visual apparatus, and even the very fact that I’m observing the blue sky, are all amongst the conditions of observation. The sets of conditions will typically share members, for example, two observations might both be made at the same time, or by the same person, or be of the same object, etc. But when incompatible facts are observed, there will be some difference in the set of conditions under which the relevant facts are observed to be the case. This cannot fail to be the case given that the fact that the observation has a certain content is itself a condition of observation

(vacuously, if it were false of the observer that she observes what she does, then the observation would be of something else). I will also refer to the conditions of observation as ‘points of view’ or ‘perspectives’. Still following the *ways = properties* principle that we saw in §4.2, I assume that the ‘ways in which things appear to an observer’ are ways things appear *to be*, i.e. properties things appear to have when observed. This means that sentences ‘A’ and ‘B’ simply state that certain things have certain properties (namely those properties that things appear to have).

Note that when we are confronted with an instance of perspectival variance, we have no reason to believe the conditions of observation to be distortive, given clause (4). If one accepts the weak empiricist thesis that observations are a guide to how things are in the world, cases of perspectival variance constitute by definition cases in which we have *prima facie* reasons to believe that incompatible facts obtain. And the latter is of course typically assumed to be something that cannot be.

Each of the cases we discussed is an instance of perspectival variance. They fit the structure as follows:

Time. An object a is (1) straight when observed at one time, and (2) bent when observed at a different time, and (3) the fact that a is straight and the fact that a is bent are incompatible, and (4) no time is *the* time at which things are observed as they truly are.

Consciousness. (1) When it is S_1 who observes the world there is something it is like when S_1 experiences things and (2) when it is someone else S_2 who observes the world there is nothing it is like when S_1 experiences things, and (3) the fact that there is something it is like when S_1 experiences things and the fact that there is nothing it is like when S_1 experiences things are incompatible facts, and (4) no one subject is *the* subject for which things are observed as they truly are, i.e. you do not have to be a certain subject to observe the world as it truly is.

Colours. An iridescent surface is (1) blue all over when observed under one angle, and (2) green all over when observed under a different angle,

and (3) the fact that the surface is blue all over and green all over are incompatible facts, and (4) no one angle is *the* angle under which things are observed as they are, i.e. you do not need to look at the world at a certain angle to see how it truly is.

Special relativity. An object a is observed to be at rest, and (2) a is observed to have a constant velocity, and (3) the fact that a is at rest and the fact that a moves is are incompatible, and (4) there is no reason to privilege the observing of any object as being at rest or the observing of any object as having a constant velocity to be privileged over the other.

Needless to say, these are just some of the types of incompatible facts that can arise and only some of conditions of observations under which such observations can be made. Note that in the case of special relativity, the relevant condition of observation is just that the observation is of a certain fact obtaining. As Galileo urged, there is nothing about an observation of something as being at rest that correlates with this observation other than it simply being an observation of that object as being at rest. In other words, the different conditions of observation are just that the one is of the object as being at rest and the other of the object as having a certain velocity. I have argued that, on closer inspection, the other cases are like that as well. For example, in the case of conscious experience, I argued that the relevant sense of *it being the case that it is S_1 who makes the observation* just is it being the case that there is something it is like when S_1 observes the world. In the case of time, I argued that *it being a certain time when the observation is made* consists in the observed fact co-obtaining with the observation being made.

The phenomenon of perspectival variance is a widespread phenomenon. Given that it's so widespread, it's unsurprising that we possess a range of conceptual strategies for dealing with such cases. Different strategies give rise to different views of the world. We have encountered the following:

Multiplication: it's not really a that has those properties but two distinct objects, b and c that have respectively the one and the other

property. For example, in the case of change, when a stick appears straight at one time and bent at another time, we can argue that it is really distinct temporal parts of the stick that are straight and bent. Or, when the same rose appears red to one subject and green to another subject, one can argue that it is really distinct sense data or ideas or experiences that are respectively red and green. When it's no longer one and the same object that instantiates these properties, we no longer have incompatible facts.

Modalization: The conditions of observation are a modality, a way for something to be the case. For example, in the case of time, we can argue that, though *a* is *F*, it merely *will* be or *was* the case that it is *G*. Or in the case of consciousness, we can argue that there really is something it is like when I undergo experience, but that it is merely *from some other else's point of view* the case that there is something it is like when that person undergoes experience. This way, one of the two incompatible facts is not really the case.

Relationalization: The apparently incompatible properties are really relations to those different conditions or to entities involved in those conditions. For example, in the case of change, we can argue that *a* is not straight and bent simpliciter but bears the straight-at and bent-at relations to distinct times. Or, we can argue that colours are really three places relations of *being red-to-observer-...-under-conditions-....* Given that it is different pairs of relata that instantiate the relations, there are no longer any incompatible facts.

Dispositionalization: The relevant properties are really dispositions to look those ways under those conditions. For example, for an object to be red is for it to be such that *it would look phenomenally red if it were observed by ... under conditions*

Of course, there are certainly other strategies and subtle variations of the ones mentioned here. We make our choices to apply one or the other strategy to different cases of perspectival variance. For example, someone might apply

the multiplication strategy in the case of change and hence accept temporal parts, yet refrain from applying it in the case of colours and deny that there are sense data or similar such mental entities that have the properties we experience things to have. Different mixes of these and other strategies result in different conceptions of the world.

These strategies all have something in common, they all lead to antirealist conceptions of the apparent facts that are observed to be the case under the different sets of conditions. Each of these strategies consists in denying that the incompatible facts obtain, and offering up some substitute for the facts that appear to be the case. To the extent that the *incompatible* facts appeared to be the case, this is taken to be merely a way things seem to be or merely a tempting way of describing the facts, and not really the way things are out there. They are all instances of an anti-realist treatment of perspectival variance:

Anti-realist treatments: When some object a appears as F under some set of conditions and as G under a different (but equally good) set of conditions, a is neither really F , nor really G .

There are many discussions on how best to formulate an antirealist position about a certain subject matter.¹ Here I have described antirealism as the denial of something *really* being F . What is it for something to be *really* F ? This is a difficult question. It does not matter for our purpose whether the notion of something really being the case can be unpacked and, if so how. In the case of the antirealist strategies, what is observed under the different (but equally) good conditions is deemed to fall on the appearance side of the reality-appearance distinction, however one conceives of it.

Since the antirealist strategies deny the reality of that which is nevertheless observed to be the case, it naturally leads to the assumption that perspective is *always* some condition that affects the way we experience things, and that perspective is thus only responsible for different ways of experiencing or describing the same set of facts and not a matter of the facts themselves.

¹On a helpful discussion of how not to formulate antirealist conceptions of a certain subject matter, see Rosen (1994).

The perspectival variance is on this assumption itself only ever a variance in the experience or description of the world and not a variance in the world itself.² But seeing perspectival variance as a single phenomenon with the range of instances that we encountered in the previous chapters makes this assumption suspect.

First of all, there are instances of perspectival variance where it's clearly a mistake to think that the world remains exactly the same across the different conditions of observation. The case of time is a straightforward example of this. Here it is clearly the facts out there that, in some way or other, do not remain exactly the same as time passes. The variance is not just in our experience or description of the world, it is out there, in the facts experienced. It is *because* there is a variance in the facts that obtain out there across time that there is a variance in the facts that are observed to obtain across time.

Secondly, we saw that it's really any kind of instrument or measuring system that can record different facts under different conditions; the perspectival facts are not just 'recorded' within human experience. We saw this very clearly in the case of special relativity, where any measuring device records incompatible facts across different perspectives (frames), not just human experience. In the case of time too, it is not just human experience that shows the world to be built from incompatible facts across time, any measuring system will record different facts across time and different measuring systems will concur in how the facts differ across time. It's quite an assumption to believe that such a wide range of different kinds of measuring devices are 'affected by perspective' in the same way to produce the same kinds of distortions.

Cases of perspectival variance can be clearly distinguished from cases of illusion or hallucination which *are* a matter of certain factors specifically affecting our experience, and not due to certain a variance of the facts observed. When I hallucinate, what I observe does not square with what is

²For example, as Adrian Moore puts it: 'One and the same reality is represented from different points of view in different ways. [...] There is only a question of its being represented in one way from one point of view, and in another from another. A perspectival representation of things is not a representation of perspectival things.' (A. Moore 1997: 50).

registered by other types of measuring systems. There is a disagreement between different kinds of measuring systems because hallucination and illusion *are* a quirk of human experience - and specific to it. But it's precisely for this reason that the conditions that cause a hallucinating state are not on a par with conditions that do not cause such states, indeed, this is precisely why such cases are not instances of perspectival variance - violating clause 4, which said that neither set of conditions of observation is worse than the standard conditions of observing things to be the ways they appear.

Thirdly, there are instances of perspectival variance where the relevant condition that would have to be causing human experience to observe certain things turns out to be just the fact that those facts are observed. If we assume that it's the conditions of observation that cause the observations to have the content they have, and not the facts themselves, then in these sorts of cases it will be the fact that observations have a certain content that causes those observations to have that content, a circular and vacuous explanation. For example, what is the relevant condition for seeing things from the first person perspective of a certain subject? You have to be that subject. But the relevant sense of 'being that subject' just is for one to see things from the first person perspective of that subject. But we cannot say that, seeing things from the first person perspective of a subject is what *causes* or *affects* experience to be such that I see things from the first person perspective of the subject. This is a circular explanation. Similarly, in the case of special relativity, the only condition of observation of an observation of some object as being at rest is that it is an observation of that object as being at rest. Some of the observations involved in perspectival variance float free from other facts, and have a very narrow set of conditions under which the observation is made. There really is nothing out there to affect human experience to be a certain way other than the very facts observed.

The assumption that perspectival variance is a phenomenon *of* human experience or of human representation of the world is an implausible dogma of the current paradigm. The fragmentalist rejects this assumption: incompatible facts are registered across a wide range of observational and measuring systems and these systems all concur in their measurements despite their

mutual differences *because* perspectival variance is a feature of what is measured and recorded, the facts themselves, and not a matter of the recording or measuring of the world. Instead of seeing the perspectival variance as a quirk of human experience, a proper appreciation of the phenomenon and its range of instances shows again that the emergence of perspectival variance does not have its source in features of human experience. Rather, perspectival variance is a worldly phenomenon, independent of human engagement with that world.

This means that the antirealist approaches have little to say in support of them by way of independent reasons to think that they must be right. That is to say, the antirealist approaches are supported by little else than the unity assumption, the assumption that any two facts that obtain co-obtain: two incompatible facts cannot co-obtain, and so, given the unity assumption, the only option seems to be that neither obtains (or, on the modalization strategy, that merely one of them obtains).

With the possibility that the unity assumption does not hold, however, the antirealist treatment no longer has the status of a direct consequence of the incompatibility of the observed facts, on the contrary, that the observed facts are incompatible now constitutes a good reason to doubt the unity assumption and with it, the plausibility of the antirealist approaches. Cases of perspectival variance support a fragmentation in the relevant facts:

Fragmentalism: when some object a appears as F under some set of conditions and as G under a different set of conditions, a is really F insofar as it is observed under the set of conditions and a is really G insofar as it is observed under the second set of conditions.

Fragmentalism takes cases of perspectival variance at face value. It should be clear that fragmentalism stands in contrast to the antirealist strategies in offering a realist conception of the facts that are observed to be the case under the different conditions:

Realist treatment: When some object a appears as F under some set of conditions and as G under a different set of conditions, a is really F

and really *G*.³

If the world can be fragmented, the denial that the observed facts in cases of perspectival variance are real turns out to be unmotivated. The antirealist treatment is unmotivated because (1) the non-obtaining of the observed facts is assumed to follow from the incompatibility of those facts even though this does not follow, and (2) there are no independent reasons to think that the observed facts fail to obtain, on the contrary, given the observations of them as obtaining, there are good reasons to think that the relevant facts do obtain.

Is fragmentalism the only realist treatment on offer? One might think that another realist treatment would be a form of metaphysical relativism: according to which the observed facts *do* really obtain, only their obtaining is relative to the relevant points of view. In the appendix to this chapter (§7.5), I argue that metaphysical relativism is not in fact a conceptually stable response to cases of perspectival variance. I do not know of any other realist treatment and, if this is right, then any case of perspectival variance provides *prima facie* reasons to believe in the fragmentation of the world.

From the point of view of the fragmentalist, the paradigm that is based in the unity assumption is a paradigm in crisis. The antirealist consequences of the unity assumption slowly drive us to a Kantian conception of things, where we live in a world that is more and more of our own making. Cases of perspectival variance have long been exerting a pressure to extract content from the world as we find it, a process that became especially explicit from the 17th century onwards.⁴ But we gather more and more evidence that cases of perspectival variance run wide and deep, and then deeper still. Do we really want to say that, next to the colours, we somehow project the lengths, shapes and masses of objects, as well the temporal order and duration of events? This recurring of perspectival variance inspires a relentless thinning out of our conception of reality if we presume we must deny that perspectivally variant properties are ever genuinely instantiated.⁵ If we

³This is assuming that '*a* is F' and '*a* is G' are atomic sentences, and the predication of a complex predicate. Better, it is assuming that F and G are intrinsic properties of *a*.

⁴For discussion, see Burnyeat (1979) and B. Williams (1979/2005).

⁵The thinning out due to perspectival variance is to be sharply distinguished from the sort of thinning out that arises from the assumption that any concept whose content is

find such perspectival variances over and over, in the case of colour, shape, length, mass, temporal order and duration, then each time the plausibility of the thesis that the world is fragmented raises with the implausibility of the claim that things don't really instantiate the properties that are involved in perspectival variance. The moment we countenance the possibility of a fragmented world, the sheer range of cases of perspectival variance constitutes a solid body of inductive evidence that the world is indeed fragmented.

If the world is fragmented in the way proposed above, there is an irony in the mistake that we have made in our antirealist treatments of the phenomena. We thought that, every time we revised our conception of the world and deemed more and more properties to be merely secondary properties or merely ways things appear, we were freeing our conception of the world from a parochial anthropomorphism, taking the relevant properties to be due to particularities of our experience of the world. As it turns out, however, the anthropomorphism wasn't in the properties we observe things to have in our experience, but rather in the hidden assumption that any fact must obtain together with every other fact, a structural mould that we have tried to impose on a world that does not fit it and which forced us, again and again, to reconfigure the facts to varying degrees of contortion to make them fit. The antirealist treatments seem like contorted attempts to reshape the facts as they are given in our observations of the world, and reflects a failure to follow the observed evidence where it leads.

derived from our experience *thereby* does not apply to the world as it is in itself. Dummett criticizes the latter question as unintelligible:

How is the world to be described as it is in itself?' This formulation shows very clearly the contradictory objective of our quest. [...] we were seeking to attain a conception of the world not encapsulated in any description; for any description must employ a particular conceptual vocabulary, and any such vocabulary must reflect, and depend on, the particular way in which the world is apprehended by beings whose thoughts are framed within that vocabulary. But there can be no such thing; a conception of something can be mediated only by some manner of describing it. (Dummett 2006: 98-99).

Dummett may be right about any conception that filters out any concept whose content depends on the experience of a certain kind of thinker (though I doubt this assumption is really driving the quest for a conception of the world itself), it does not bear on any conception that only filters out properties that admit of perspectival variance.

7.2 Other instances of perspectival variance

The previous systematization of the proposed views around the phenomenon of perspectival variance leads to the question whether there are more such cases. I want to tentatively defend the following claim:

Uniform application thesis: For any genuine case of perspectival variance, the incompatible facts that are observed to obtain indeed obtain without co-obtaining. Any perspectival variance is a source of fragmentation.

Now opponents of the fragmentalist framework can pounce on this thesis and use it offer objections of the form ‘So, then, this other case is also one in which the relevant facts really obtain’ where the obtaining of the relevant facts is deemed to be wildly implausible. I will discuss a few examples of this overgeneration worry. This overgeneration worry is common in discussions of perspectival facts but, I think, should be less worrying than it is perceived to be.

Before looking into this overgeneration worry, why defend the strong claim that any case of perspectival variance is a case of fragmentation? Roughly the reason is that any case of perspectival variance cannot fail to be a case in which we have good reasons to accept the incompatible facts that are observed to obtain under different conditions, this is part of the conditions that render a case a genuine case of perspectival variance. As I pointed out in the previous section, the only reason why we resort to antirealist treatments is the fact that the relevant facts are incompatible and the fact we assume that all facts co-obtain. The possibility of a fragmented world undermines this reason to deny the relevant facts and so will undermine any good reason to still deny that those facts obtain. In short: there is no reason not to follow where evidence leads in genuine cases of perspectival variance. But, of course, if I turn out to be mistaken in this, this in no way imputes the application of the fragmentation strategy to the areas we have considered, it only imputes the existence of a simple formula telling us when to resort to fragmentation. If the uniform application thesis turns out to be mistaken, we have no choice

but to follow a more piecemeal approach, just as we typically follow a more piecemeal approach in the application of the antirealist treatments mentioned above.

The crucial element in the *uniform application thesis* is the condition that a given case is genuinely a case of perspectival variance. The importance of this becomes clear when we turn to the various overgeneration worries that can be raised against any theory that accepts some form of perspectival facts. I will discuss three examples.

The first example is the case of metaphysical possibilities. One can imagine someone objecting along the following lines: ‘We could say that from the perspective of this possible world I have no wings, but from the perspective of a different possible world, I do have wings. So are we now going to say that it is really the case that I have wings and that I am wingless, but that I have wings *insofar as* I am not wingless?’

Clearly not. It’s not the case at all, in any way, that I have wings, not even insofar as anything else is the case. And the problem with the line of reasoning should be obvious: the condition of observing me as having wings is the condition of *observing things in a certain possible world*. But observing things in another possible world is not on a par with observing things at the actual world. Or, better, possibly observing something is not at all on a par with actually observing something, only actual observations offer us reasons to think that something is really the case. As Stalnaker remarks:

We can grant that fictional characters are as right, from their point of view, to affirm their full-blooded reality as we are to affirm ours. But their point of view is fictional, and so what is right from it makes no difference as far as reality is concerned. (Stalnaker 1976: 69).

The metaphor of ‘possible worlds’ engenders a conception of them as legitimate standpoints or points of view, from which things seem to be a certain way, a different way. But, as Stalnaker aptly points out, what seems to be the case from a point of view that itself has not reality, offers us no reason whatsoever to think that those things are really the case. The modal case

is not a case of perspectival variance because it fails to meet the condition of there being no reason to privilege one perspective over the other, we have ample reason to privilege the one set of conditions over the other: the one perspective is real and the other not.

Consider a second example. One can imagine someone objecting along the following lines: ‘Say that we are on the phone, and I claim that the beach is close by and you disagree that it is not close by but far-away. Are we going to say that the beach is really close by and really far-away, but not the one insofar as it is the other? Those seem very dubious facts.’ The worry is, roughly, that given some linguistically indexical description, the fragmentalist implausibly takes the resulting sentences whose truth value varies across contexts to correspond to genuine facts that co-obtain with those facts that constitute the relevant context.

The matter is here more subtle than in the modal case. The objection can be applied to facts such as the following: *a*’s being here, *a*’s being heavy, *a*’s being tall, *a*’s being me, *t*’s being now, etc. It seems to me that, at least in most cases, we find such properties dubious because they do not seem to be qualitative ways in which anything is genuinely observed to be, but rather seem to be mere ways of conveniently stating certain other facts, which *are* observed to obtain. Now it’s in no way part of the *uniform application thesis* that we are to accept, for any predicate that gives rise to indexical statements, that these attribute genuine properties to things. We merely have the thesis that *if* any of these are genuine properties that things are observed to have, and things are observed to have incompatible properties under different (but equally good) sets of conditions, *then* we have a reason to think that things really are those ways. To put it crudely, first we must figure if these are genuine ways things could be regardless of perspectival variance. There is only a question of whether or not the fragmentalist framework should be applied to a pair of incompatible facts when we have good reasons to think that they are genuine ways things are observed to be.

Given some predicate, ‘is *F*’, we need to figure out if things described as *F* in different circumstances share a qualitative similarity, underwritten by a single property that is always expressed by the predicate, or whether the use

of the predicate attributes different properties in the different circumstances and the things described as *F* in different circumstances are qualitatively dissimilar. This can be a difficult question of course. When I observe a tree and describe it as close by instead of far away, and you are somewhere else and describe different things as close by, is there some qualitative similarity between the things that we observe as being close by? I believe there isn't but I am not at all confident on the matter, given that it ultimately relies on phenomenological reflections on the contents of what I exactly experience when I experience a bunch of things as close by. Do I experience them just as close *to me* or also as close simpliciter in some sense? If the latter, then the uniform application thesis will indeed imply that they genuinely obtain and are a source of fragmentation, if the former, then not. But the important point to note is that, either way the objection is no good. Either the facts are dubious for an independent reason (i.e. dubious because they clearly seem to be predicates whose attribution in different cases do not state incompatible facts) in which case we have an independent reason to think that they are not a source of perspectival variance, *or* there turns out to be no good reason to think the descriptions do not state incompatible facts, so that there is nothing dubious about those facts, in which case there is no force to the supposed objection. Either way, the objection has no force against the uniform application thesis as such.⁶

However, there might well be cases that we commonly think of as elliptic descriptions but really aren't when we come to see the possibility of a fragmented world. A likely example of this is the kind of perspectival variance that arise from differences in orientations. Let me return to the case we discussed in Chapter 1: imagine that we face each other, and that Sophie and Nora stand between us. We observe things differently: you observe Sophie as being on the left of Nora, I observe Sophie as being on the right of Nora.

⁶Note that, in the case of time, we did not argue for the fact that *some time is present* and *past*. This was precisely because I think that this is a misguided way of thinking of temporal matters (reifying times, and resorting to dubious properties as *being present* and *being past*). I similarly did not argue for the obtaining of *ML's being me* and *TN's being me* in the case of consciousness for the (independent) reason that these are just dubious facts.

Sophie and Nora are observed to have incompatible two-place relations from our different perspectives. One might object ‘are we now really going to say that Sophie is both on the left and right of Nora?’

Perhaps surprisingly, there are quite good reasons to think that this is indeed the right thing to say. I think that the two-place relations of *being on the left of...* and *being on the right of...* are genuine relations in which things stand, and that, all things that are observed to be on the left of something else, are qualitatively similar. When I turn around, the content of what I observe *changes*, indeed, the content of my earlier experience is incompatible with the content of my later experience in terms of these two relations.

There is a version of this case that conflates it with the indexical cases just discussed. Consider the description of this orientational case by Adrian Moore:

Dorothy sees a box on her left and says, ‘The box is there on the left.’ Again her representation is both true and perspectival. It is from her spatial point of view. Here the temptation’s grip is so slight as to be virtually negligible. It is the temptation to think that her representation is true in virtue of a feature of reality that is there from where she is situated but not from other positions, namely the box’s being on the left. [...] After a moment’s reflection it has no grip at all. Dorothy’s representation is true simply in virtue of the relevant spatial relation between her and the box. The representation is made true by the non-perspectival fact that the box is on the left of her. (A. Moore 1997: 44-55).

This is in fact a different case, namely a case in which things are described as on the left simpliciter. It’s indeed implausible to think that it is a genuine fact that *the box is on the left*. When we are facing the same way, and we both look at the box, and you say ‘it is on the left’ and I say ‘it is on the left’ we state different facts: you state that it is on the left of you and I state that it is on the left of me. Our descriptions, though using the same predicate, do not express a qualitative similarity, and this predicate ‘is on the left’ is an indexical. But not so in the case of the two-place predicate ‘is

on the left of...’, which - I want to claim - does captures a genuine relation that makes for an experienced similarity amongst things.

I think we find a similar situation with such notions as being tall, or being heavier; these are indexical predicates that express relations of *being taller than* and *being heavier than*, and these later relations are involved in perspectival facts, as we have seen in our discussion of special relativity. Whether something is a genuine case of perspectival variance depends on what we think are the genuine properties that things appear to have.

Note also that everyone is committed to there being proper and improper arities of the predicates that express genuine facts. This is not a commitment unique to the fragmentalist framework, the latter simply mobilises that commitment because it’s concerned with which facts are incompatible facts. To see this more clearly, consider a thought experiment. Take all the predicates in terms of which we describe objects. And now divide the universe in half, call one half A and the other half B, and add an argument place to all our predicates, which is intended to be occupied by the two halves of the universe. The things in region A are now described as being straight-at-A, existing-at-A, green-at A, taller-than-*b*-at-A, etc. and everything in region B is described in terms of the new predicates with region B in the relevant argument places. And imagine that instead of using our old predicates, we only use these predicates instead. We now think that our usual descriptions were all elliptic for our $n+1$ -aric predicates. On what basis do we reject such a proposal as misguided or inadequate? It seems clear to me that we reject such a description of the world because it fails to capture any qualitative similarities across objects that are in different regions. Only the objects that are in the same region and are described by the same relativized predicate are described as qualitatively similar. But, intuitively, the green objects from region A appear to be qualitatively similar to the green objects from region B, and this means that there is pressure on us to describe these objects as sharing some property.

So any of the mentioned reasons to be suspicious of positing the apparent facts have turned out be *ipso facto* reasons to doubt whether the relevant case is a case of perspectival variance. This does not mean that the *uniform*

application thesis is not a thesis that brings with it substantive commitments. Indeed, in the final example (of Sophie being on the left of Nora), we saw that the thesis can lead to cases that will be found counter-intuitive. The question is whether there is any reason to be suspicious of these facts other than that they are incompatible. My bet is that there isn't.

If that is right we should note a further theoretical virtue of fragmentalism, namely the potential uniformity we might reach in the treatment of cases of perspectival variance in contrast to some of the other conceptual strategies we encountered. Imagine that we adopt the multiplication strategy in the case of colours, and take - between our perspectives on the iridescent surface of a soap bubble - your sense datum to be red all over and my sense datum to be green all over. Each sense datum is exclusive to each experiential perspective. Imagine we apply this also to the case of orientations: there is an orientational token Sophie that is on the left of an orientational token Nora, and a different orientational token Sophie that is on the right of an orientational token Nora. Each orientation offers a perspective on its own exclusive sets of objects, with different properties. This is clearly quite an implausible view. Similarly, imagine we adopt a modalization strategy in the case of time, and say that the tree *was* straight and *is* bent. And now say we try to apply this again to the orientational case: it's other-orientationally the case that Sophie is on the left of Nora. Again we arrive at quite an unnatural view. Indeed, I do not see how any of the antirealist strategies could result in a plausible view of the world when applied uniformly to all cases of perspectival variance. To the extent we think that phenomena of the same kind should be given accounts of the same kind, the potential viability of the uniform application thesis might offer a surprising systematicity to our conception of the world.

7.3 Objectivity in a fragmented world

Nagel rightly observes that '[t]here is a tendency to seek an objective account of everything before admitting its reality' (Nagel 1979b:196). As Dummett

explains in a little more detail:

[R]eality must be something of which there exists in principle a complete description. I can make drawings of a rock from various angles, but if I am asked to say what the real shape of the rock is, I can give a description of it as in three-dimensional space which is independent of the angle from which it is looked at. The description of what is really there, *as it really is*, must be independent of any particular point of view. (Dummett 1960: 503).

The description that is only true in certain circumstances is not a description that is objectively true, and what is not objectively true is not a description of what is really there. Objectivity is assumed to be a necessary condition for a description of reality: any real fact is an objective fact, or contrapositively, when some description fails to be objective it is ruled out from being a description of what is really the case. The fragmentalist conception of the world might seem to sever this assumed link between reality and objectivity. Given that a fact may only obtain insofar as certain other facts obtain, it's not guaranteed that an observation in a certain circumstance will co-obtain with the fact we aim to observe, and hence it's not guaranteed the latter will be observed under those conditions. And in that case, the description of what is really there, as it really is, is not independent of any particular point of view, contrary to what Dummett and many others assume.

One might think that this a subversive idea. Our best guide to what the world is like, and what facts obtain is science, and scientific methodology is premised on the importance of objectivity. Indeed we think science offers the best picture of reality *because* it offers the most objective picture of reality. Does the fragmentalist conception of reality, and its break of the relation between reality and objectivity, undermine scientific methodology?

Before we can assess whether this is the case, and if so whether it is really worrying or not, we first need to be careful about the way we are understanding objectivity. Objectivity has been understood in many ways, from meaning something like being fair or bias-free, to being intersubjectively agreed upon, to being mind-independent or value-free (when applied

to facts), to simply being accurate or correct (when applied to representations).⁷ Objectivity in most of these senses is obviously left untouched by the fragmentalist framework.

The sense of objectivity in which it is potentially undermined is the sense of something objective being multiply accessible. Nozick formulates this criterion as follows:

First, an objective fact is accessible from different angles. Access to it can be repeated by the same sense (sight, touch, etc.) at different times; it can be repeated by different senses of the same observer, and also by different observers. (Nozick 2001: 75-76).

A truth or fact is objective in the sense of being accessible when it can be observed (1) at different times, (2) through different sense modalities, and (3) by different observers. If this is a criterion of objectivity, then indeed, the possibility of fragmented worlds shows that a fact can fail to be objective in this sense and yet be a real fact. Objectivity in this sense is not a necessary condition for a fact to be real.

But why should we think that descriptions that fail to be objective in this sense cannot be descriptions of what is really the case?

Consider first the idea that an objective fact should be observable at different times. Surely it should be possible that a real fact can obtain, and then stop to obtain, leaving no trace whatsoever? Although such a fact is not accessible at a later time this surely does not mean that it is not the case at the earlier time. There is no reason to think that the transience of an apparent fact is a bar to the reality of a fact.

The second dimension of accessibility mentioned by Nozick is the idea that an objective fact must be observable through different sense modalities (sight, touch, etc.) But say I note that my shampoo smells of coconut. It is plausible to think that I cannot see or touch the smell; that I can *only* smell it. Why should I be less convinced that my shampoo is not really as I smell it to be. More generally, certain facts may only be observable by

⁷See e.g. Gaukroger (2012) and Megill (1994).

particular measuring devices and not others due to the nature of the relevant facts and what the measuring devices are sensitive to. I do not see why the range of devices, including the different sense modalities of organisms, that can observe an apparent fact has any absolute bearing on the reality of the observed fact.

And a third dimension of accessibility mentioned by Nozick is the idea that different observers can observe the apparent fact. If the relevant observers are different in the sense of having different sense modalities, or in making their observations at different times, then this criterion collapses into the ones just discussed. So assume that we have two observers who are exactly alike in all these respects, i.e. that possess the same sense modalities and make their observations at the same time. If some fact is in principle only observable to the one observer but not the other, does this suffice for thinking the fact must be unreal?

Here the matter is more controversial than in the other cases of observability. In the discussion of consciousness, we saw that the fragmentalist framework makes room for subjective or private facts.⁸ However given the many other kinds of privateness that may arise, such as a temporal privateness of a fact, I do not see why any subjectively private fact should be banned from obtaining. Why should we assume this?

To be sure, what is undermined is solely the inference from the non-accessibility of an apparent fact to the non-obtaining of that fact. Accessibility here refers to the observation of the relevant fact and not the possibility of entertaining the thought that the relevant fact obtains. It's for this rea-

⁸Nozick remarks that quantum mechanics may also predicts facts that are only accessible to one observer, noting the following passage by David Albert:

There are going to [be] combinations of things that any particular observer can (in principle) simultaneously know about the world, and that nobody else can simultaneously know about; what those things are will depend on the identity of the observer in question, and some of those things will invariably be things that are about the observer itself. (Albert 1992: 186).

Nozick draws from this the conclusion that such facts are not 'intersubjectively available' (2001: 93). But I do not see how Albert's passage suggests anything more than the claim that these are just facts that no other observer can know *simultaneously*. If Nozick is right however, quantum mechanics may predict further types of facts that are only ever observable to one observer.

son important to distinguish the perspectival facts from facts specified using indexicals. The idea is not that there are, for example, propositions that only I could express using the sentence ‘I am making a mess’, and that such facts obtain. The private access in question is not one of expressibility but of verification: an observer may, because of certain contingencies, be in the position to verify and observe only a limited subsets of the facts. The issue is not one of ineffability.⁹

The fragmentalist position also does not imply that private accessibility should not raise any epistemic warning flags. Needless to say, when someone claims a certain fact obtains but also that he or she has sole access to it, we should be suspicious, as it may be a way of immunizing the relevant facts from falsification. Fragmentalism changes nothing in those respects. It’s tempting to adopt a conception of the world that enables us to make short work of such claims on the grounds that there cannot be such private facts; but such an assumption seems quite dogmatic.

Fragmentalism does however change how we think of scientific progress, in particular, how we think of the way in which scientific findings typically alter our conception of the world. There is a widespread image of scientific progress as cleansing our conception of the world from those properties that are parochial, and merely express how the world impresses upon observers of a certain constitution (like ourselves), instead of reflecting how the world is in itself. Bernard Williams is well known for this picture of science. Consider this (long) quote:

The suggestion is that there are possible descriptions of the world using concepts which are not peculiarly ours, and not peculiarly relative to our experience. Such a description would be that which would be arrived at, as C. S. Peirce put it, if scientific enquiry continued long enough; it is the content of that ‘final opinion’ which

⁹Perry notes that ‘if one finds plausible reasons elsewhere for believing in a universe that has, in addition to our common world, myriads of private perspectives, the idea of propositions of limited accessability will fit right in’ (Perry 1979:16). That may be right, but I do not see any independent reason to believe that, if facts may not obtain insofar as a certain subject experience and thinks about the world, the relevant fact should also be inexpressible to that subject.

Peirce believed that enquiry would inevitably converge upon, a ‘final opinion [...] independent not indeed of thought in general, but of all that is arbitrary and individual in thought’ (Peirce 1966: 82). [...] The picture, that offered by natural science, would explain the phenomena: it would explain them, moreover, even as they present themselves in other, more local, representations. It is this consideration that gives the content to the idea, essential to the traditional distinction, that the scientific picture presents the reality of which the secondary qualities, as perceived, are appearances. (B. Williams 1979/2005: 229).

Scientific enquiry converges on a view that leaves the facts that are only accessible by particular observers in particular circumstances ‘behind’: the scientific view explains those facts away as being the ways things appear to observers of the relevant type and situated in the relevant sort of circumstances.

Given the incompatibility of the facts we find across perspectives, there may seem to be a logical inevitability behind this progression. The progression proceeds in the following steps. We start with a certain set S of facts that we think obtain given the set of conditions under which we are able to observe the world, C . We think of new type of experiments, or develop new tools, which expands the set of conditions under we can observe the world to C^+ , and when we do observe the world under those condition, we find facts incompatible with some of the facts in S . We remove that type of fact from the set of facts we originally thought to obtain, arriving at a contracted set S^- and add to these certain new (types of) facts S^* to arrive at a new collection of facts $S^- \cup S^*$. For example, as we learn about relativity, we remove facts concerning definite length, shape, mass, etc. and add facts concerning new quantities, such as spacetime interval. The original set of facts shrinks in the process because those facts conflict with the apparent facts observed under the new conditions; whatever role scientific explanation plays, the mere incompatibility of the facts suffices for them be removed from the facts we think obtain in really. And, in the process, some of the original

conditions of observation are now deemed to be perspectival in the pejorative sense of leading to observations that are misleading with regard to the genuine content of reality.

Scientific enquiry is thought of a little differently when we consider it in the context of the fragmentalist conception of the world. As science proceeds, we increase the range of conditions under which we observe the world and uncover further facts. But even when some of the further facts clash with the original set of facts, this does not mean those facts are not amongst those that obtain in reality. The process is thus simplified, as the conditions of observation extends from C to $C \cup C^+$, the facts extend with it from S to $S \cup S^+$. Scientific enquiry does not typically narrow the set of facts we think are really the case, it typically broadens it. And the reason for this is that incompatible facts can be allowed to obtain across structural fragmentations of the world.

The conditions under which incompatible facts can be observed are not thought of as a source of misleading appearances, but simply as a limitation that explains a lack of access to a certain range of facts. Whenever facts are retracted from our conception of the world, this is simply because we think those facts were thought to obtain on the basis of errors, or that there is no good reason to believe they obtain, or because they are observed under conditions that prove to be a source of illusion. The clash with facts that are observed to obtain from other perspectives is no longer a sufficient reason for revisions to our conception of the world. There is no direct conflict between a 'manifest image' and 'scientific image', the manifest facts are just a subset of the wider set of facts countenanced in a scientific view of the world, with various explanatory relations amongst them. Scientific methodology remains exactly as it is, only the pressure to draw various sorts of metaphysical conclusions from empirical findings is affected by the possibility of fragmentation.

7.4 Summary and concluding remarks

When confronted with fragmentalism, many respond with a sense of unease; ‘this cannot be right - what about objectivity?, what about these other cases?, what about all the less radical alternatives?’ I hope the above discussion has taken away some of these worries, and that it has shown some further general reasons to think that fragmentalism might be the right story. The more standard responses to cases of perspectival variance have antirealist implications for an increasing range of properties (and objects), whereas fragmentalist allows a realist treatment. The uniform application of fragmentalism to cases of perspectival variance does not lead to an obviously implausible view - contrary to the more standard responses. And fragmentalism only severs the link between our conception of reality and multiple accessibility. But it’s in any case an implausibly strong assumption that a fact obtains if and only if it’s entirely multiply accessible.

7.5 Appendix: metaphysical relativism

One may doubt whether the arguments that I have offered in favour of fragmentalism do not just as well speak in favour of some form of metaphysical relativism, a view that might be easier to come to terms with. I don’t have the space to discuss metaphysical relativism in the full detail it deserves. Let me however mention a potentially conclusive objection raised by Adrian Moore (1997: Ch. 3), showing that metaphysical relativism is unstable.¹⁰

Let metaphysical relativism be the view that the apparent facts in cases of perspectival variance do really obtain, only their obtaining is relative to certain entities, so that we avoid saying that two incompatible facts obtain simpliciter. It’s really the case relative to you *that Sophie is on the left of Nora*, and it’s really the case relative to me *that Sophie is on the right of*

¹⁰I have put the argument in my own terms. Lewis (1986:204, 2002) raises a related worry for endurantists who hold that intrinsic properties are had relative to times. Fine (2005b:§11) raises a related worry for what he calls ‘external relativism’ about tensed facts.

Nora.¹¹

Thus far we have a claim about facts, about the *the fact that Sophie is on the left of Nora*, which is said to obtain relative to some object in the world, namely you. What we aim to capture however is the way *Sophie* and *Nora* appear to you from your perspective. How does the relative obtaining of the fact bear on the involved objects? When a fact obtains simpliciter, the story is straightforward: for *the fact that Sophie sits* to obtain is plausibly just for Sophie to sit (or at least implies that this is so). As Moore notes, the obtaining of perspectival facts similarly needs to consist in the world's being a certain way (1997: 46). The question is how the relative obtaining makes things be. What could it be for *the fact that Sophie is on the left of Nora* to obtain 'relative to you'? If we say that it's for Sophie to be on the left of Nora relative to you, then we have not captured the way Sophie appears to you. Again: if it appeared to you that Sophie is on the left of Nora relative to you, then it appeared to me that Sophie is on the right of Nora relative to me, and that isn't incompatible with the way it appeared to you. But Sophie and Nora *do* appear in incompatible ways.

If we say instead that for *the fact that Sophie is on the left of Nora* to obtain 'relative to you' is in some way or other for Sophie to be on the left of Nora, then the relativization is idle: by the same token, for *the fact that Sophie is on the right of Nora* to obtain relative to me would be for Sophie to be on the right of Nora, so that Sophie is both on the left and right of Nora after all. The whole point of the relativization was to avoid saying that incompatible facts obtain simpliciter.

Can a relativist refuse to explain how the relative obtaining of the fact bears on the way Sophie and Nora are? No: the aim is to admit the reality of what appears to you from your perspective, and what appears to you is that Sophie is on the right of Nora - so whatever story we offer, it needs to

¹¹Distinguish metaphysical from semantic relativism, according to which the truth of an asserted or believed content is relative to contexts (as in MacFarlane 2005). The type of question that is of concern to us is whether Sophie can be both on the left and on the right of Nora, and this a metaphysical question about what objects can be like, and not a question about the ways in which representations of the world are true or false. There may be relations between the two kinds of relativisms, but such relations are relations between views that should be sharply distinguished.

tell us how things stand with regards to *them*.

Although we need to consider possible responses and adaptations of the view before we can conclusively reject it, I believe that Moore's argument reveals a serious conceptual instability in metaphysical relativism.¹² Fragmentalism avoids this instability: that the fact that Sophie is on the left of Nora obtains really means that Sophie is on the left of Nora, and similarly for the fact that Sophie is on the left of Nora. This is the respective ways in which Sophie and Nora appear to us. Given that metaphysical relativism is unstable in ways that fragmentalism isn't, I do not think that the reasons in favour of fragmentalism are similarly reasons to adopt metaphysical relativism.

¹²For a defence of the argument against various responses, see A. Moore (1997: Ch. 3). See Spencer (forthcoming) for a detailed discussion and defence of metaphysical relativism.

Conclusion

Each one of us finds a tiny fraction of the world. In this tiny fraction, certain things manifest certain qualities. You find Sophie to be on the right of Nora and not on the left of her; you find the surface of the soap bubble to be greenish and not blueish, you find the tree in the dunes to be straight and not bent; you find yourself sitting still whereas the cars outside drive by; and you find that things come to manifest themselves in these ways because you, the physical organism reading these sentences, is situated in a certain way in the world. You know that each one of us claims to find such tiny fractions of world, and that these can conflict with the tiny fraction given to you. What do you do?

You might hold fast to our own tiny fraction, after all, it's right there. Others claim to find different and even incompatible fractions but perhaps they are wrong. You could deem your tiny fraction to be all there is and, in this way, become Wittgenstein's solipsist:

For what the solipsist means is quite correct; only it cannot be said, but makes itself manifest. The world is my world: this is manifest in the fact that the limits of language (of that language which alone I understand) mean the limits of my world. The world and life are one. - I am my world (microcosm). - There is no such thing as a subject that thinks or entertains ideas. If I wrote a book 'The world as I found it', I should have to report on my body and say which parts were subordinate to my will, and which were not, etc. This then would be a method of isolating the subject or rather of showing that in an important sense there

is no subject: that is to say, it alone could not be mentioned in the book. - The subject does not belong to the world but it is a limit of the world. (Wittgenstein 1921/1961: 5.62-5.632).

Could it be that the limits of the world coincide with the limits of the tiny fraction that is manifest to you, your world? That the only language you truly understand is the language that predicates properties that are manifest in this tiny fraction? The thought is implausible. Not only do you find the manifestation of a tiny bit of world, you have memories of the way things were at different times, when the tree in the dunes was still growing straight; you have seen the soap bubble from a different angle, appearing more blueish; you know you can walk around in such a way that Sophie comes to be on the left of Nora instead of on the right of her; you can imagine being me in a certain sense, and - more generally - you can imagine the world being manifested in a different way. You know that the manifestation of the world depends on the state of the physical organism that you are, and that I and anyone else are similar such physical organisms. I claim that different things manifest themselves in my experience? How could you explain my mistake? How could you make sense of the things I say and do, if you make sense of yourself in light of the ways things are manifest to you in your experience?

‘Yes’, says the current paradigm, ‘and so do not be misled by the way the world is manifested in your experience’. To see how things really are, you must consider what *explains* the way things appear in your experience of the world. This already requires that we stand back from the way things are manifest to us and consider those things, not as the disclosure of a bit of world, but as various kinds of experiences of that world. And so, following the current paradigm, you might become William’s absolutist:

In understanding, even sketchily, at a general and reflective level, why things appear variously coloured to various observers, we shall find that we have left behind any idea that, in some way which transcends those facts, they ‘really’ have one colour rather than another. In thinking of these explanations, we are in fact

using a conception in which colour does not figure at all as a quality of the things. (B. Williams 1979/2005: 227)

Sophie is just on the right of Nora *relative to you*, and this will explain in some way why Sophie *appears* to you to be just on the right of Nora. The soap bubble *appears* greenish in your visual experience *because* of your angle and the reflectance properties of the surface; the tree is only straight at the current time; and so we move away, further and further, from the way things appear to be when we interact with the world.

The current paradigm has a long and esteemed history. But there are serious issues, issues that are often swept under the carpet. A judicious use of words helps to hide the issues. Instead of ‘Sophie is on the right of Nora’, we say ‘Sophie is on the right of Nora relative to you’, which makes it sound as if the latter fact is much like the former, ‘just a little more relational’. But we cannot simply confer our grasp of the two-place notion to the concept of the three-place notion, as the three-place relation does not describe the way we experience Sophie and Nora to be related. The content of this predicate must be a different property; one which is not manifest within our experience. But many properties of things are subject to perspectival variance, properties that feature in many successful explanations of the world. Various phenomena, such as the passage of time and our first person perspective on the world, become unintelligible. Our conception comes to suffer from an explanatory impoverishment, and the near-sceptical unravelling of our grasp of what the world is like in itself.

Fragmentalism recommends that we neither cling to our own embedded perspective, nor abandon it in favour of some more objective view, but that we recognize a different way of accumulating the multiple perspectives in one conception of the world, a conception according to which the world is not one unified scene. The fraction that is given in your experience might well be the way it’s given to you, and yet it might also be the way it’s given in my embedded perspective on the world. We must stand back from our embedded perspectives *in some way*, if we are to see the world in its entirety, as William’s absolutist rightly insists on. Fragmentalism offers a different

way of combining the contents of our conflicting perspectives.

Because no other facts co-obtain with all the facts in a given fragment, it can seem as if a given fragment is exhaustive of the world, as if it's all that obtains. And, in some sense it is: the facts that constitute a fragment are all the facts insofar as these facts all obtain. Yet, in another sense they are not all the facts, they are not all the facts that obtain as such. Fragments can mislead us when we do not distinguish between the co-obtaining and obtaining of facts; it's this that leads Wittgenstein's solipsist to think that the facts that do not co-obtain with those he finds fail to obtain at all. The limits of your perspective on the world may be a limit in the world itself, as Wittgenstein's solipsist assumed, but that does not mean that no facts obtain beyond that limit.

The result is that a tiny fraction of the world is manifested in our experience the way it is in itself; even if it's merely a fraction. Sophie really is on the right of Nora and, insofar as this is so, she is not on the left of her; the surface of the soap bubble is greenish and, insofar as this is so, it's not blueish; and so on. We can appeal to our experience for a sense of what the world is like, why something cannot be straight insofar as it's round, but can be straight insofar as it's red, or blue, or any other colour.

The result is a background metaphysics that enables us to offer theories of the world in terms of concepts that seemed inapplicable in a unified world. The passage theory was one example of this. At least since McTaggart's work, the notion of a genuine passage of time has seemed incoherent, and so our choice seemed to be between the tensed conception of the A-theory, or the spatializing conception of the B-theory. Within a fragmented world, the concept of passage is itself applicable and this means that there is room for a new view, a passage theory of time.

The fragmentalist background theory also changes the dialectical landscape, freeing certain views from implications that made them seem implausible. We saw this in the case of the no-subject view of experience. The facts that constitute the appearance of your fraction do not co-obtain with the facts that constitute the appearance of my fraction. If we do not distinguish between the obtaining and co-obtaining of facts, then the claim that

my phenomenal facts do not co-obtain with yours collapses into the solipsist claim that my phenomenal facts do not obtain at all. With the distinction, there is no such collapse.

Besides advocating the particular fragmentalist theories, of time, change, subjectivity, secondary properties and frame-dependent properties, a more general aim has been to advocate fragmentalism as a general paradigm within which to formulate our theories. The background assumption of a unified and non-perspectival world has a crippling and indeed contortive effect. Without it, I believe that we may make some genuine progress in our philosophical understanding of things. A fragmented world - that should be our background assumption.

Bibliography

- Albert, D. (1992). *Quantum Mechanics and Experience*. Cambridge (MA): Harvard University Press.
- Alston, W. P. (1999). Back to the Theory of Appearing. *Philosophical Perspectives* 13, 181–203.
- Anscombe, G. E. M. (1975). The First Person. *Mind and Language*. S. Guttenplan (Ed.). Oxford: Clarendon Press, 45–65.
- Armstrong, D. M. (1969). Colour-Realism and the Argument from Microscopes. *Contemporary Philosophy in Australia*. R. Brown and C. D. Rollins (Eds.). London: Allen and Unwin, 119–131.
- (1973). *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- (1999). *The Mind-Body Problem*. Boulder: Westview Press.
- Averill, E. W. (1992). The Relational Nature of Color. *Philosophical Review* 101.3, 551–588.
- Bayne, T. and D. Chalmers (2003). What is the Unity of Consciousness? *The Unity of Consciousness: Binding, Integration, and Dissociation*. A. Cleeremans and C. Firth (Eds.). Oxford: Oxford University Press, 24–58.
- Berkeley, G. (1713/1979). *Three Dialogues between Hylas and Philonous*. R. M. Adams (Ed.). Indianapolis: Hackett Publishing Company.
- Block, N. (1990). Inverted Earth. *Philosophical Perspectives* 4, 53–79.
- (2003). Mental Paint. *Reflections and Replies: Essays on the Philosophy of Tyler Burge*. M. Hahn and B. Ramberg (Eds.). Cambridge (MA): MIT Press, 165–200.
- Boghossian, P. (2006). What is Relativism? *Truth and Realism*. P. Greenough and M. Lynch (Eds.). Oxford: Clarendon Press, 13–37.

- Bourne, C. (2006). *A Future for Presentism*. Oxford: Oxford University Press.
- Braithwaite, R. B. (1932). The Nature of Believing. *Proceedings of the Aristotelian Society* 33, 129–146.
- Broad, C. D. (1938/1976). *Examination of McTaggart's Philosophy*. London: Octagon Books.
- Burgess, J. P. (2002). Basic Tense Logic. *Handbook of Philosophical Logic, Vol. 7*. D. Gabbay and F. Guenther (Eds.). Dordrecht: Kluwer Academic Publishers.
- Burnyeat, M. (1979). Conflicting Appearances. *Proceedings of the British Academy* 65, 69–111.
- Byrne, A. (2001). Intentionalism Defended. *Philosophical Review* 110.2, 199–240.
- (2009). Experience and Content. *Philosophical Quarterly* 59.236, 429–451.
- Byrne, A. and D. R. Hilbert (2003). Color Realism and Color Science. *Behavioral and Brain Sciences* 26.1, 1–44.
- (2007). Color Primitivism. *Erkenntnis* 66.1, 73–105.
- Campbell, J. (1993). A Simple View of Color. *Reality, Representation and Projection*. J. Haldane and C. Wright (Eds.). Oxford: Oxford University Press, 257–268.
- Chalmers, D. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- (2010). Perception and the Fall from Eden. *The Character of Consciousness*. D. Chalmers (Ed.). New York: Oxford University Press, 381–454.
- Chisholm, R. (1976). *Person and Object*. Chicago: Open Court.
- Cohen, J. (2009). *The Red and the Real*. Oxford: Oxford University Press.
- Correia, F. and S. Rosenkranz (2012). Eternal Facts in an Ageing Universe. *Australasian Journal of Philosophy* 90.2, 307–320.
- Crane, T. (1988). The Waterfall Illusion. *Analysis* 48.3, 142–147.
- Crisp, T. (2003). Presentism. *Oxford Handbook of Metaphysics*. M. J. Loux and D. W. Zimmerman (Eds.). Oxford: Oxford University Press, 211–245.
- Cushing, J. T. (1994). *Quantum Mechanics: Historical Contingency and the Copenhagen Hegemony*. Chicago: University of Chicago Press.

- Dainton, B. (2000). *Stream of Consciousness: the Unity and Continuity of Conscious Experience*. London: Routledge.
- (2010). *Time and Space, Second Edition*. Durham: Acumen.
- Davidson, D. (2001). *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press.
- Davies, M. and I. L. Humberstone (1980). Two Notions of Necessity. *Philosophical Studies* 38.1, 1–30.
- Deng, N. (2013). Fine’s McTaggart, Temporal Passage, and the A versus B-debate. *Ratio* 26.1, 19–34.
- Dennett, D. C. (1991). Real Patterns. *Journal of Philosophy* 87.1, 27–51.
- Descartes, R. (1641/1970). Meditations on First Philosophy. *The Philosophical Works of Descartes, Vol. I*. E. S. Haldane and G. R. T. Ross (Eds.). Cambridge: Cambridge University Press.
- Dretske, F. (2003). Experience as Representation. *Philosophical Issues* 13.1, 67–82.
- Dummett, M. (1960). A Defense of McTaggart’s Proof of the Unreality of Time. *Philosophical Review* 69.4, 497–504.
- (2006). *Thought and Reality*. Oxford: Oxford University Press.
- Earman, J. (1974). An Attempt to Add a Little Direction to ‘the Problem of the Directionality of Time’. *Philosophy of Science* 41.1, 15–47.
- Ehring, D. (1997). Lewis, Temporary Intrinsics, and Momentary Tropes. *Analysis* 57.4, 254–258.
- Einstein, A. (1920/2001). *Relativity: the Special and the General Theory*. London: Routledge.
- Encyclopædia Britannica* (2015). Atomic Clock Time. Encyclopædia Britannica Online. URL: <http://www.britannica.com/EBchecked/topic/41796/atomic-time>.
- Evans, G. (1985). Does Tense Logic Rest on a Mistake? *Collected Papers*. G. Evans (Ed.). Oxford: Clarendon Press.
- Fine, K. (1977/2005). Prior on the Construction of Possible Worlds and Instants. *Modality and Tense: Philosophical Papers*. K. Fine (Ed.). Oxford: Oxford University Press.
- (2001). The Question of Realism. *Philosophers’ Imprint* 1.1, 1–30.

- Fine, K. (2005a). Necessity and Non-Existence. *Modality and Tense: Philosophical Papers*. K. Fine (Ed.). Oxford: Oxford University Press.
- (2005b). Tense and Reality. *Modality and Tense: Philosophical Papers*. K. Fine (Ed.). Oxford: Oxford University Press.
- (2007). Response to Paul Horwich. *dialectica* 61.1, 17–23.
- (2012). A Guide to Ground. *Metaphysical Grounding: Understanding the Structure of Reality*. F. Correia and B. Schnieder (Eds.). Cambridge: Cambridge University Press.
- Galileo (1632/1967). *Dialogue Concerning the Two Chief World Systems*. S. Drake (Ed.). Berkeley: University of California Press.
- Gaukroger, S. (2012). *Objectivity: A Very Short Introduction*. Oxford: Oxford University Press.
- Gazzaniga, M. S. (1970). *The Bisected Brain*. New York: Appleton-Century-Crofts.
- (1985). *The Social Brain: Discovering the Networks of the Mind*. New York: Basic Books Inc.
- Gazzaniga, M. S., R. B. Ivry, and G. R. Mangun (2014). *Cognitive Neuroscience: Biology of the Mind, Fourth Edition*. New York: W. W. Norton.
- Gert, J. (2008). What Colors Could Not Be: an Argument for Color Primitivism. *The Journal of Philosophy* 105.3, 128–155.
- Gödel, K. (1949/1990). A Remark About the Relationship between Relativity Theory and Idealistic Philosophy. *Kurt Gödel, Collected Works*. S. Feferman (Ed.). Oxford: Oxford University Press.
- Godfrey-Smith, W. (1979). Special Relativity and the Present. *Philosophical Studies* 36.3, 233–244.
- Gray, J. (2005). Felix Klein’s Erlangen Programme, Comparative Considerations of Recent Geometrical Research. *Landmark Writings in Western Mathematics 1640-1940*. I. Grattan-Guinness (Ed.). Amsterdam: Elsevier, 544–553.
- Hardin, C. L. (1988). *Color for Philosophers: Unweaving the Rainbow (Expanded Edition)*. Indianapolis: Hackett Publishing Company.
- Hare, C. (2009). *On Myself, and Other, Less Important Subjects*. Princeton: Princeton University Press.

- Harman, G. (1990). The Intrinsic Quality of Experience. *Philosophical Perspectives* 4, 31–52.
- (2001). General Foundations versus Rational Insight. *Philosophy and Phenomenological Research* 63.3, 657–663.
- (2007). Explaining the Explanatory Gap. *American Philosophical Association Newsletter on Philosophy and Computers* 6.2, 2–3.
- Haslanger, S. (1989). Endurance and Temporary Intrinsics. *Analysis* 49.3, 119–125.
- (2003). Persistence Through Time. *The Oxford Handbook of Metaphysics*. M. J. Loux and D. W. Zimmerman (Eds.). Oxford: Oxford University Press, 315–354.
- Hellie, B. (2011). There it Is. *Philosophical Issues* 21.1, 110–164.
- (2013). Against Egalitarianism. *Analysis* 72.2, 304–320.
- Hinchliff, M. (2000). A Defense of Presentism in a Relativistic Setting. *Philosophy of Science (Proceedings)* 67.3, S575–S586.
- Hume, D (1739/1975). *A Treatise of Human Nature*. L. A. Selby-Bigge and P. H. Nidditch (Eds.). Oxford: Clarendon Press.
- Hurley, S. (1998). *Consciousness in Action*. Cambridge (MA): Harvard University Press.
- Jackson, F. (1982). Epiphenomenal Qualia. *Philosophical Quarterly* 32.127, 127–136.
- (1996). The Primary Quality View of Color. *Philosophical Perspectives* 10, 199–219.
- (2004). Representation and Experience. *Representation in Mind*. H. Clapin, P. Staines, and P. Slezak (Eds.). Amsterdam: Elsevier, 107–124.
- James, W. (1909/1967). A Pluralistic Universe. *Essays in Radical Empiricism: a Pluralistic Universe*. R. B. Perry (Ed.). Gloucester (MA): Peter Smith Publishers.
- Jaśkowski, S. (1948/1969). Propositional Calculus for Contradictory Deductive Systems. *Studia Logica* 24, 143–157.
- Johnston, M. (1987). Is there a Problem about Persistence? I. *Proceedings of the Aristotelian Society (Supplementary Volume)* 61.3, 107–135.
- (1992). How to Speak of the Colors. *Philosophical Studies* 68.3, 221–263.

- Johnston, M. (2007). Objective Mind and the Objectivity of Our Minds. *Philosophy and Phenomenological Research* 75.2, 233–268.
- Kalderon, M. E. (2007). Color Pluralism. *Philosophical Review* 116.4, 563–601.
- Kamp, H. (1971). Formal Properties of Now. *Theoria* 37.3, 227–273.
- Kant, I. (1781/1997). *Critique of Pure Reason*. P. Guyer and A. W. Wood (Eds.). Cambridge: Cambridge University Press.
- Kaplan, D. (1977). Demonstratives. *Themes from Kaplan*. J. Perry J. Almog and H. Wettstein (Eds.). Oxford: Oxford University Press, 481–561.
- Kennedy, M. (2009). Heirs of Nothing: the Implications of Transparency. *Philosophy and Phenomenological Research* 79.3, 574–604.
- Kim, J. (1993). Events as Property Exemplifications. *Supervenience and Mind*. J. Kim (Ed.). Cambridge: Cambridge University Press, 33–52.
- Kripke, S. A. (1982). *Wittgenstein on Rules and Private Language*. Cambridge (MA): Harvard University Press.
- Lakoff, G. and M. Johnson (1980). *Metaphors We Live By*. Chicago: University of Chicago Press.
- (1999). *Philosophy in the Flesh: the Embodied Mind and its Challenge to Western Thought*. New York: Basic Books.
- Langsam, H. (1997). The Theory of Appearing Defended. *Philosophical Studies* 87.1, 33–59.
- Le Poidevin, R. (2010). Time Without Change (in Three Steps). *American Philosophical Quarterly* 47.2, 171–180.
- Leibniz, W. G. (1686-87/1989). Correspondence with Arnauld, 1686-87 (Selections). *Philosophical Papers and Letters, Second Edition*. L. E. Loemker (Ed.). Dordrecht: Kluwer Academic Publishers.
- Levine, J. (2001). *Purple Haze*. Oxford: Oxford University Press.
- Lewis, D. (1979). Attitudes De Dicto and De Se. *The Philosophical Review* 88.4, 513–543.
- (1982). Logic for Equivocators. *Nous* 16.3, 131–141.
- (1986a). Events. *Philosophical Papers, Vol. II*. D. Lewis (Ed.). Oxford: Oxford University Press, 241–269.
- (1986b). *On the Plurality of Worlds*. Oxford: Blackwell.

- Lewis, D. (1999). Individuation by Acquaintance and by Stipulation. *Papers in Metaphysics and Epistemology*. D. Lewis (Ed.). Cambridge: Cambridge University Press.
- (2002). Tensing the Copula. *Mind* 111.441, 1–14.
- Lichtenberg, G. C. (1804/1990). *The Waste Books*. R. J. Hollingdale (Ed.). New York: The New York Review of Books.
- Lockwood, M. (1989). *Mind, Brain and the Quantum: the Compound 'I'*. Oxford: Wiley-Blackwell.
- Lowe, E. J. (1988). The Problems of Intrinsic Change: Rejoinder to Lewis. *Analysis* 48.2, 72–77.
- MacBride, F. (2001). Four New Ways to Change Your Shape. *Australasian Journal of Philosophy* 79.1, 81–89.
- MacFarlane, J. (2005). Making Sense of Relative Truth. *Proceedings of the Aristotelian Society* 205.1, 305–323.
- Mackie, J. L. (1976). *Problems from Locke*. Oxford: Oxford University Press.
- Marcus, R. B. (1990). Some Revisionary Proposals about Belief and Believing. *Philosophy and Phenomenological Research* 50, 133–153.
- Marks, C. (1980). *Commissurotomy, Consciousness and Unity of Mind*. Cambridge (MA): MIT Press.
- Martin, M. G. F. (2002). The Transparency of Experience. *Mind and Language* 4.4, 376–425.
- Maudlin, T. (2007). *The Metaphysics within Physics*. Oxford: Oxford University Press.
- (2011). *Quantum Non-Locality and Relativity: Metaphysical Intimations of Modern Physics, Third Edition*. Oxford: Blackwell.
- (2012). *Philosophy of Physics*. Princeton: Princeton University Press.
- McDowell, J. (1985). Values and Secondary Qualities. *Morality and Objectivity*. T. Honderich (Ed.). London: Routledge and Kegan Paul, 110–127.
- McGinn, C. (1983). *The Subjective View*. Oxford: Clarendon Press.
- (1996). Another Look at Color. *The Journal of Philosophy* 93.11, 537–555.

- McLaughlin, B. P. (2003). Colour, Consciousness and Colour Consciousness. *Consciousness: New Philosophical Perspectives*. Q. Smith and A. Jokić (Eds.). Oxford: Clarendon Press, 97–154.
- McTaggart, J. E. (1908). The Unreality of Time. *Mind* 17.68, 457–474.
- Megill, A. (1994). Four Senses of Objectivity. *Rethinking Objectivity*. A. Megill (Ed.). Durham (NC): Duke University Press.
- Mellor, D. H. (1981). *Real Time*. Cambridge: Cambridge University Press.
- (1986). Tense’s Tenseless Truth-Conditions. *Analysis* 46.4, 167–172.
- (1998). *Real Time II*. New York: Routledge.
- Merlo, G. (2013). Specialness and Egalitarianism. *Thought* 2.3, 1–10.
- (Dissertation). *The World as I Found It*.
- Meyer, U. (2013). *The Nature of Time*. Oxford: Clarendon Press.
- Minkowski, H. (1964/1908). Space and Time. *Problems of Space and Time*. J. J. C. Smart (Ed.). New York: Macmillan Company, 297–312.
- Monton, B. (2006). Presentism and Gravity. *The Ontology of Spacetime, Vol. 1*. D. Dieks (Ed.). Amsterdam: Elsevier.
- Moore, A. (1997). *Points of View*. Oxford: Oxford University Press.
- Moore, G. E. (1955). Wittgenstein’s Lectures in 1930–33. *Mind* 64.253, 1–27.
- Nagel, T. (1979a). Brain Dissections and the Unity of Consciousness. *Mortal Questions*. T. Nagel (Ed.). Cambridge: Cambridge University Press, 147–164.
- (1979b). Subjective and Objective. *Mortal Questions*. T. Nagel (Ed.). Cambridge: Cambridge University Press.
- (1979c). What is it Like to be a Bat? *Mortal Questions*. T. Nagel (Ed.). Cambridge: Cambridge University Press.
- (1986). *The View from Nowhere*. Oxford: Oxford University Press.
- Norton, J. D. (2010). Time Really Passes. *Humana Mente: Journal of Philosophical Studies* 13, 23–34.
- Nozick, R. (2001). *Invariances: the Structure of the Objective World*. Cambridge (MA): Harvard University Press.
- Oaklander, N. L. and Q. Smith, eds. (1995). *The New Theory of Time*. New Haven: Yale University Press.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.

- Park, D. (1972). The Myth of the Passage of Time. *The Study of Time*. J. T. Fraser, C. Haber, and G. H. Müller (Eds.). Berlin: Springer, 110–121.
- Peacocke, C. (1983). *Sense and Content*. Oxford: Oxford University Press.
- (1984). Colour Concepts and Colour Experience. *Synthese* 58.3, 365–381.
- Peirce, C. S. (1966). A Critical Review of Berkeley’s Idealism. *Charles S. Peirce, Selected Writings (Values in a World of Change)*. P. P. Wiener (Ed.). New York: Dover Publications.
- Percival, P. (2002). A Presentist’s Refutation of Mellor’s McTaggart. *Royal Institute of Philosophy Supplement* 50, 91–118.
- Perry, J. (1979). The Problem of the Essential Indexical. *Noûs* 13.1, 3–21.
- Price, H. (2011). The Flow of Time. *The Oxford Handbook of the Philosophy of Time*. C. Callender (Ed.). Oxford: Oxford University Press, 276–311.
- Priest, G. (2007). Paraconsistency and Dialetheism. *Handbook of the History of Logic*. D. Gabbay and J. Woods (Eds.). Amsterdam: North Holland, 129–204.
- (2008). Jaina Logic: A Contemporary Perspective. *History and Philosophy of Logic* 29.3, 263–278.
- Prior, A. N. (1960). The Runabout Inference Ticket. *Analysis* 21.2, 38–39.
- (1962/2003). Changes in Events and Changes in Things. *Papers on Time and Tense, New Edition*. P. Hasle et al. (Eds.). Oxford: Oxford University Press, 7–19.
- (1967). *Past, Present and Future*. Oxford: Oxford University Press.
- (1968/2003a). Egocentric Logic. *Papers on Time and Tense, New Edition*. P. Hasle et al. (Eds.). Oxford: Oxford University Press, 223–240.
- (1968/2003b). Tense Logic and the Logic of Earlier and Later. *Papers on Time and Tense, New Edition*. P. Hasle et al. (Eds.). Oxford: Oxford University Press, 117–138.
- (1970). The Notion of the Present. *Studium Generale* 23, 240–248.
- (1977/2003). Worlds, Times and Selves. *Papers on Time and Tense, New Edition*. P. Hasle et al. (Eds.). Oxford: Oxford University Press, 241–256.
- Putnam, H. (1967). Time and Geometry. *Journal of Philosophy* 64.8, 240–247.

- Putnam, H. (2005). A Philosopher Looks at Quantum Mechanics (Again). *British Journal of the Philosophy of Science* 56.4, 615–634.
- Quine, W. V. O. (1960). *Word and Object*. Cambridge (MA): MIT Press.
- Rescher, N. and R. Brandom (1980). *The Logic of Inconsistency: a Study in Non-Standard Possible-Worlds Semantics and Ontology*. Oxford: Basil Blackwell.
- Restall, G. (1997). Ways Things Can't Be. *Notre Dame Journal of Formal Logic* 38.4, 583–597.
- Rodriquez-Pereyra, G. (2003). What is Wrong with the Relational Theory of Change? *Real Metaphysics: Essays in Honour of D. H. Mellor*. H. Lillehammer and G. Rodriquez-Pereyra (Eds.). London: Routledge, 184–195.
- Rosen, G. (1994). Objectivity and Modern Idealism: What is the Question? *Philosophy in Mind: The Place of Philosophy in the Study of Mind*. M. Michael and J. O'Leary-Hawthorne (Eds.). Dordrecht: Kluwer Academic Publishers, 277–319.
- (2006). The Limits of Contingency. *Identity and Modality*. F. MacBride (Ed.). Oxford: Oxford University Press, 13–39.
- Sartori, L. (1996). *Understanding Relativity, a Simplified Approach to Einstein's Theories*. Berkeley: University of California Press.
- Saunders, S. (2002). How Relativity Contradicts Presentism. *Royal Institute of Philosophy Supplement* 50, 277–292.
- Savitt, S., ed. (1995). *Time's Arrow Today: Recent Physical and Philosophical Work on the Direction of Time*. Cambridge: Cambridge University Press.
- (1998). There's No Time Like the Present (in Minkowski Space-time). *Philosophy of Science (Proceedings)* 67, 563–574.
- (2002). On Absolute Becoming and the Myth of Passage. *Time, Reality & Experience (Royal Institute of Philosophy Supplement, Vol. 5)*. C. Callender (Ed.). Cambridge: Cambridge University Press, 153–168.
- Schaffer, J. (2009). On What Grounds What. *Metametaphysics*. D. Chalmers, D. Manley, and R. Wasserman (Eds.). Oxford: Oxford University Press.
- Schechter, E. (2013). The Unity of Consciousness: Subjects and Objectivity. *Philosophical Studies* 165.2, 671–692.

- Schechter, E. (2014). Partial Unity of Consciousness: A Preliminary Defense. *Sensory Integration and the Unity of Consciousness*. D. J. Bennett and C. S. Hill (Eds.). Cambridge (MA): MIT Press, 347–74.
- Schlick, M. (1936). Meaning and Verification. *Philosophical Review* 45.4, 339–69.
- Schwitzgebel, E. (2002). A Phenomenal, Dispositional Account of Belief. *Noûs* 36, 249–275.
- Shoemaker, S. (1969). Time Without Change. *Journal of Philosophy* 66.12, 363–381.
- (1982). The Inverted Spectrum. *Journal of Philosophy* 79.7, 357–381.
- (1994a). Phenomenal Character. *Noûs* 28.1, 21–38.
- (1994b). The First-Person Perspective. *Proceedings and Addresses of the American Philosophical Association* 68.2, 7–22.
- (2006). On the Way Things Appear. *Perceptual Experience*. T. S. Gendler and J. Hawthorne (Eds.). Oxford: Oxford University Press, 461–480.
- Sider, T. (2001). *Four Dimensionalism: an Ontology of Time and Persistence*. Oxford: Oxford University Press.
- Simons, P. (2006). The Logic of Location. *Synthese* 150.3, 443–458.
- Sklar, L. (1981/1985). Time, Reality and Relativity. *Philosophy and Space-time Physics*. L. Sklar (Ed.). Berkeley: University of California Press, 289–304.
- Smart, J. J. C. (1949). The River of Time. *Mind* 58.232, 483–494.
- Speaks, J. (2009). Transparency, intentionalism, and the nature of perceptual content. *Philosophy and Phenomenological Research* 79.3, 539–573.
- Spencer, J. (forthcoming). Relativity and Degrees of Relationality. *Philosophy and Phenomenological Research*.
- Sperry, R. (1968). Hemisphere Deconnection and Unity in Awareness. *American Psychologist* 23.10, 723–733.
- Stalnaker, R. (1976). Possible Worlds. *Noûs* 10.1, 65–75.
- (1984). *Inquiry*. Cambridge (MA): MIT Press.
- Stein, H. (1968). On Einstein-Minkowski Space-Time. *Journal of Philosophy* 65.1, 5–23.

- Strawson, P. F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. London: Routledge.
- Stroud, B. G. (2000). *The Quest for Reality, Subjectivity and the Metaphysics of Colour*. Oxford: Oxford University Press.
- Tallant, J. (2013). A Heterodox Presentism: Kit Fine's Theory. *Contributions to Philosophy*. R. Ciuni, K. Miller, and G. Torrenco (Eds.). Munich: Philosophia Verlag GmbH, 2–26.
- Taylor, R. (1960). Pure Becoming. *Australasian Journal of Philosophy* 38.2, 137–143.
- Tooley, M. (2000). *Time, Tense and Causation*. Oxford: Oxford University Press.
- Trevarthen, C. (1970). Analysis of Cerebral Activities that Generate and Regulate Consciousness in Commissurotomy Patients. *Hemisphere Function in the Human Brain*. S. J. Dimond and J. G. Beaumont (Eds.). New York: Halsted Press Book.
- Tumulka, R. (2006). A Relativistic Version of the Ghirardi-Rimini-Weber Model. *Journal of Statistical Physics* 125.4, 825–844.
- Turner, J. (2011). Ontological Nihilism. *Oxford Studies in Metaphysics, vol. 6*. D. Zimmerman and K. Bennett (Eds.). Oxford: Oxford University Press.
- Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge (MA): MIT Press.
- (2000). *Consciousness, Color, and Content*. Cambridge (MA): MIT Press.
- (2002). Visual Qualia and Visual Content Revisited. *Philosophy of Mind: Classical and Contemporary Readings*. D. J. Chalmers (Ed.). Oxford: Oxford University Press.
- (2003). *Consciousness and Persons: Unity and Identity*. Cambridge (MA): MIT Press.
- (2006). Nonconceptual Content, Richness, and Fineness of Grain. *Perceptual Experience*. T. S. Gendler and J. Hawthorne (Eds.). Oxford: Oxford University Press.
- (2014). Transparency, Qualia Realism and Representationalism. *Philosophical Studies* 170.1, 39–57.

- Varzi, A. (1997). Incoherence Without Contradiction. *Notre Dame Journal of Formal Logic* 38.4, 621–638.
- Warmbrod, K. (2004). Temporal Vacua. *Philosophical Quarterly* 54.215, 266–286.
- Watkins, M. (2005). Seeing Red: the Metaphysics of Colour Without the Physics. *Australasian Journal of Philosophy* 83.1, 33–52.
- Williams, B. (1979/2005). *Descartes: the Project of Pure Enquiry*. Oxford: Routledge.
- Williams, D. C. (1951). The Myth of Passage. *The Journal of Philosophy* 48.15, 457–472.
- Wittgenstein, L. (1921/1961). *Tractatus Logico-Philosophicus*. D. F. Pears and B. F. McGuinness (Eds.). London: Routledge and Kegan Paul.
- (1929). Some Remarks on Logical Form. *Proceedings of the Aristotelian Society (Supplementary Volume)* 9, 162–171.
- Yablo, S. (1995). Singling Out Properties. *Philosophical Perspectives* 9, 477–502.
- Zeh, H. D. (2007). *The Physical Basis of the Direction of Time, 5th Edition*. Berlin: Springer.