The dynamical foundations of physical geometry

Tushar Vasudev Menon

Balliol College, University of Oxford

Submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy

September, 2018

To Dimpy and Suresh Menon, and Rachel Fraser-Menon

Acknowledgements

I am grateful to the British Society for the Philosophy of Science for the doctoral scholarship that allowed me to write this thesis, and to the John Templeton Foundation's grant through the University of Illinois at Chicago for a six-month predoctoral fellowship in 2017-18 during which time most of this thesis was written.

I am most deeply indebted to David Wallace, and profoundly grateful to my other two supervisors, Simon Saunders and Adam Caulton. In the course of working towards finishing this DPhil, the following people have been extraordinarily supportive and intellectually influential: Dave Baker, Harvey Brown, Jeremy Butterfield, Richard Dawid, Nick Huggett, Dennis Lehmkuhl and (especially) James Read.

Two other philosophers of physics at Oxford deserve special mention: Niels Martens and my academic twin-sister Carina Prunkl. My deepest thanks to other philosophers for discussions, support and friendship: David Albert, Gordon Belot, Rhys Borchert, Ben Brast-McKie, Karamvir Chadha, David Tena Cucala, Erik Curiel, Radin Dardashti, Neil Dewar, Juliusz Doboszewski, John Dougherty, Patrick Dürr, Josh Eisenthal, Jamee Elder, Sam Fletcher, James Fraser, Sean Gryb, Josh Hunt, Rasmus Jaksland, Eleanor Knox, James Ladyman, Jessica Laimann, Vincent Lam, Martin Lesourd, Niels Linnemann, Annina Loets, Owen Maroney, Keizo Matsubara, Tom Moller-Nielsen, David Papineau, Tom Pashby, Brian Pitts, Oliver Pooley, Sebastian Murgueitio Ramirez, Alex Roberts, Katie Robertson, Laura Ruetsche, the Cambridge Simplex, Athamos Stradis, Nic Teh, Trevor Teitel, Patricia Contreras Tejada, Karim Thébault, Chris Timpson, Niranjana Warrier, Porter Williams, and Chris Wüthrich.

My unending gratitude to my friends: Nipun Nair, Archana Achal, Bhaskar Bhushan, Alma Causevic, Lia Costiner, Daniel Gildenlöw, Hrishikesh Halase, Dan Houghton, Stefanie Lenk, Anubhav Mohanty, Luke Machin, Shonali Mani, Léo Margarit, Dan Mash, Basil Mason, Nishita Nair, Venu Nair, Manasa Ramachandra, Savita Raman, Joel Rozario, Nigel Somera, Tulsi Srinivasan and Priya Urs.

Finally, my family, without whom none of this: my parents, Dimpy and Suresh Menon, Steve, Hilary, and Jack Fraser, and my wife, Rachel Fraser-Menon.

Abstract

This thesis is in two related parts. The first part presents arguments against the claim that the metric of general relativity is invariably spatiotemporal. This is done by discussing plausible criteria for spatiotemporality and demonstrating violations of those criteria. This leads to a question about the so-called *chronometric significance* of the metric, a central component of Brown and Pooley's *dynamical approach* to spacetime theories. The dynamical approach is scrutinised and, in the second part of the thesis, expressed, in opposition to the *geometrical approach* in a mathemtically novel way. This re-formulation of the position is then used to counter an ostensibly robust argument against the dynamical approach. This thesis is composed of five chapters.

Publication information

Three of the five chapters in this thesis are based either in part or wholly on papers that are forthcoming in peer-reviewed journals or collections; one is based on a paper that is currently under review. The publication information related to each of these chapters is as follows:

Chapter 1: Tushar Menon, Niels Linnemann and James Read, "Clocks and Chronogeometry: Rotating spacetimes and the relativistic null hypothesis", British Journal for the Philosophy of Science, 2018 (forthcoming)

Chapter 2: Tushar Menon and James Read, "On the limitations of inertial frame spacetime functionalism", 2018 (under review)

Chapter 3: Tushar Menon, "Taking up superspace: the spacetime structure of supersymmetric field theory", *Philosophy Beyond Spacetime: The Philosophical Foundations of Spacetime*, eds. Nick Huggett, Baptiste Le Bihan and Christian Wüthrich, Oxford University Press (forthcoming)

Chapter 5: Tushar Menon, "Algebraic fields and the dynamical approach to physical geometry", Philosophy of Science (forthcoming)

Permission to include material based upon jointly-authored work has been obtained from the Faculty of Philosophy, University of Oxford, as well as from all relevant co-authors. I am the lead author on all of the above works, and the author of the entirety of this thesis. Due acknowledgement has been given over the course of this thesis to the contributions of my collaborators.

This thesis is approximately 63,000 words long.

Contents

Introduction

Space and time are central to theorising in both philosophy and physics. This thesis is a systematic treatment of the latter.

It is crucial to realise that, even within physics, the term 'spacetime' means different things to different people. Here is one account: it is the thing whose structure is measured by rods and clocks. This, I take it, is the sort of thing Knox, for example, has in mind when motivating her project of spacetime functionalism:

> [O]ur empirical access to spacetime structure is mediated by a number of distinctly spatiotemporal phenomena—the behaviour of rods, clocks and light rays are some examples. [68, fn. 3]

This is distinct from, but not incompatible with, the *codificatory* conception of spacetime as expressed by Dewar [33]:

> [T]he characteristic feature of something's being *spacetime* structure—as opposed to material structure—is its pervasiveness, the fact that all dynamical processes take note of it. [33, p. 163]

Spacetime, on this latter view, is a useful construct insofar as it codifies certain significant 'pervasive' facts of which dynamical fields 'take note'. Of course, one needs to be very careful with claims like this—consider, for example, in quantum mechanics, the requirement that the probability current associated with a wavefunction is conserved by all dynamical evolution, irrespective of which system is under consideration. This is encoded in the structure of the Hilbert space of states and the requirement of unitarity on the operators, something of which the dynamical processes definitely ‘take note’. So why do we not consider a Hilbert space to be spatiotemporal? This is not a rhetorical question. I will have much more to say about this in this thesis.

In trying to assess the conceptual status of spacetime, I have found myself drawn not to the age-old (and perhaps outdated) debate between substantivalists (who assert that spacetime is a fundamental constituent of the ontology of the universe) and relationalists (who assert that spacetime’s existence is non-fundamental) but rather, the more recent debate between proponents of the so-called *dynamical* and *geometrical* approaches to spacetime theories.[1] The dynamical approach to spacetime theories is a position due to Harvey Brown and Oliver Pooley [18, 19], that asserts the priority of facts about matter fields (specifically symmetry-related facts) over facts about geometry (to the extent that the two can be distinguished) in accounting for the behaviour of those very matter fields. This might sound trivial—facts about matter field dynamics explain why matter fields behave the way they do; if this is a *minority* position amongst philosophers, then little wonder that Weinberg ([126, Ch. VII]) is wary of philosophy.

But accounting for (at least some)aspects of the behaviour of matter fields, in a sense to be made precise, by appeal to the geometric structure of spacetime

[1] It might be argued that the distinction between these two debates is not always clean—some proposals about, for example, emergent spacetime might be naturally understood as advocating a position in the substantivalist–relationalist debate. While this is true (consider, for example, Pooley’s claim that the dynamical approach in SR is a form of relationalism [92, §6]), my intention is to highlight that my focus is *primarily* on the dynamical–geometrical debate.

is the concensus, not only amongst philosophers, but perhaps even more overwhelmingly, amonst physicists. This powerful and pervasive intuition about certain forms of 'dynamics-neutral' structure[2] motivates a 'geometrical' account of the behaviour of matter fields. This is an account that has driven foundational work in physics for the last century, and find amongst its proponents high-profile physicists like Charles Misner, Kip Thorne and John Wheeler [79], as well as philosophers like Michael Friedman [45], Tim Maudlin [77] and James Weatherall [123]. According to proponents of this so-called *geometrical* approach, (at least some) dynamics-neutral facts about the behaviour of physical fields are accounted for by the existence of some ontologically independent geometrical structure.

This thesis naturally divides itself into two parts. The first part, titled 'on the epistemology of spacetime geometry', is primarily concerned with the behaviour of rods and clocks, the meaning of spatiotemporality, and our access to physical geometry. My novel analysis of the behaviour of clocks in general relativity in the first chapter prompts some serious questions about the nature of spatiotemporality, first as part of a general discussion of spacetime functionalism, then in the specific context of supersymmetric field theory. These questions are addressed in turn in chapters 2-3.

The second part of the thesis is intended, not so much as a defence of the dynamical position, as an exploration of its viability (of course, the positive claim that the approach is viable surely counts towards its defence!). Although not metaphysically radical, the dynamical approach does entail a somewhat radical semantic thesis—not all models of general relativity describe 'spacetimes', at

[2] I prefer 'dynamics-neutral' to 'dynamics-independent', since the real claim is not about what happens in the absence of dynamical fields, rather, about what happens irrespective of *which* matter field is under consideration.

least in the sense that the latter term is used by a majority of commentators on the subject. The thread that unifies the thesis is the *dynamical–geometrical* debate: it is to this debate that I now turn my attention.

The dynamical–geometrical debate

In special relativity, it is assumed that it is possible (even if in some idealised limit) to construct *rods* and *clocks*—devices that might be made of different materials, that survey the Minkowski metric, η_{ab}, i.e. measure intervals of proper time along a trajectory as given by that metric.[3] Although this universality feature is sometimes taken to abductively indicate the existence of an ontologically independent physical Minkowski metric field,[4] the intelligibility of the measurement process does not *require* a primitive Minkowski metric field. The process of measurement, seen as just a special case of dynamical evolution of matter fields, guarantees that rods and clocks measure intervals corresponding to the Minkowski spacetime interval only if all of their dynamical equations are Poincaré invariant (though, as we will discover in chapter 2, there are some subtleties here). A necessary condition for any measuring device to survey a metric field is that the (local) dynamical symmetries of its equations of motion are a subset of the (local) isometries of that metric field. On some versions of the geometrical account, for example, Maudlin's, this is arguably also a sufficient condition, given, in addition, a suitably detailed story of how rods and clocks supply information about the metric in some reference frame (or even

[3] In this thesis, lower case Latin letters early in the alphabet (a, b, c...) label *abstract* indices [86, 88], lower case Greek letters label coordinate indices, lower case Latin letters from the middle of the alphabet (i, j, k...) label spatial indices in a coordinate basis. Unless explicitly stated to the contrary, the Einstein summation convention applies (wherever applicable) and the constants are expressed in natural units, i.e. $c = G = \hbar = 1$.
[4] In [77], for example.

come into existence in the first place), and in the approximation that arbitrary accelerations do not destroy these stable bodies.

The dynamical approach

Here is an example of a particular type of locution that often comes up in discussions of special relativity:

Claim: Laws governing matter fields are Poincaré invariant because spacetime is Minkowskian (i.e. has a Minkowski geometry).[5]

There are (at least) two ways to read this sentence, *depending on an antecedent understanding of the referent of the term 'spacetime'*:

(1) An ontologically independent entity (spacetime) with a contingent property (Minkowski geometry) explains some feature of a distinct ontologically independent entity (matter fields).[6] On this reading, if spacetime had not been Minkowskian then it doesn't follow that all matter laws would not have been Poincaré invariant. Since spacetime is the kind of thing that could have had a different property, this is a *modally de re* claim: spacetime, in our world, is Minkowskian (for the sake of argument), but it might have been non-Minkowskian. In that non-Minkowskian possible world, we might still have some Poincaré invariant dynamical fields. Spacetime geometry is not necessarily the geometry determined by the shared symmetry group of all matter fields.

(2) Minkowski spacetime is ontologically dependent on the matter fields. To say that spacetime is Minkowskian is just to say that the geometry associated

[5] Consider, for example, Janssen [63, p. 26]: 'I argue that the space-time symmetries are the explanans and that the [Poincaré] invariance of the various laws is the explanandum'

[6] How this explanation might proceed is the subject of much debate. Balashov and Janssen, for example, propose a so-called *common origin inference*, presented in more detail in [63].

with Poincaré invariant matter is Minkowskian. Moreover, spacetime geometry *is* the geometry associated with matter. This is a *modally de dicto* claim: spacetime, being necessarily reflective of (the subset of shared) matter field symmetries, could not have been anything other than Minkowskian, given that all laws (by hypothesis, at least) are Poincaré invariant.

Brown and Pooley's view can be seen as advancing a definition of spacetime on which the *de dicto* reading follows, while, for example, Earman's principles implicitly require a definition of spacetime on which the *de re* reading follows. It is easy to see why Brown and Pooley's claim is, indeed, contentful, *pace* arguments by Acuña [1] and Myrvold [80]—it provides the definition of spacetime on which their claims of analyticity are built, while acknowledging that the other definition of spacetime is meaningful, though incorrect.[7]

A metric has *chronometric significance* if it the proper time that it determines along some trajectory is matched by the proper time measured by some (appropriately chosen, trustworthy) measuring device.To use Brown's terminology, such a matter field 'surveys the metric'. In special relativity, the chronometric significance[8] of the Minkowski metric motivates a reductive move made by Brown and Pooley, who take the Minkowski metric to be a 'glorious non-entity'. By this, they do not mean that there is no Minkowski metric structure in special relativity—quite the contrary—rather, they mean that a static Minkowski metric field, conceived of as an assignment of properties to regions of spacetime, is not fundamental.

[7] This highlights a tension in Brown and Read's [22, §3.1] decision to endorse Acuña and Myrvold's claims regarding analyticity, whilst maintaining that the arrow of explanation is asymmetric.

[8] As pointed out by Synge [109] and Penrose [86], the process of measuring 'spatial distances' in spacetime can, in theory, be carried out entirely by matter configurations that qualify as clocks. For this reason, Synge and Penrose use the term 'chronometry' rather than 'chronogeometry'. In this thesis, unless otherwise stated, I use both terms interchangeably.

It is clear that the Minkowski metric geometry is a constraint, in some sense, on the behaviour of the sorts of field to which we believe that special relativity applies—although exactly what kind of a constraint this is remains contested. Take Maudlin's interpretation:

> If we accept that in a vacuum there is no physical structure, except for the structure of space-time itself, then the behaviour of light in a vacuum implies that *the geometry of spacetime alone determines the trajectory of the light rays*. That is, given any point in the space-time, p, the structure of space-time ought to fix where light emitted from that p (in any possible direction) will go. [77, p. 68]

Balashov and Janssen [9, 63] argue that the constraints imposed by Minkowski structure are kinematical, in the sense that they are independent of the contingencies of dynamical laws:

> Special relativity is completely agnostic about what inhabits or ... *carries* Minkowski space-time. All the theory has to say about systems inhabiting/carrying Minkowski spacetime is that their spatiotemporal behaviour must be in accordance with the rules it encodes. Special relativity thus imposes the *kinematical constraint* that all dynamical laws must be Lorentz invariant. [63, p. 28]

This echoes Dewar's remarks about the pervasiveness of spacetime. On one reading, Balashov and Janssen can be thought of as advocating not the reification of a geometric object, η_{ab}, but instead its interpretation as some sort of geometrically-encoded meta-law that constrains the form of all dynamical laws.[9] Indeed, we do seem to have good evidence that this is the case—all

[9] Janssen indicates that this is a plausible reading of his position [63, p. 28, fn. 7]: 'Borrowing

observed fundamental laws of physics so far have, in fact, turned out to be locally Lorentz invariant.

For Brown and Pooley, Minkowski geometry is more like a demarcation constraint; all it is to be a special relativistic theory is to be a theory of matter fields whose equations are invariant under Poincaré transformations—there is no restriction on the form that future, undiscovered laws can take. To quote Brown,

> The appropriate structure [of spacetime in special relativity] is Minkowski geometry *precisely because* the laws of physics of the non-gravitational interactions are Lorentz covariant.[18, p. 133]

From the dynamical perspective, the Minkowski metric structure is ontologically reducible to the dynamical symmetry structure of matter fields, so it follows that the matter surveys this metric (at least to the extent that we assert that these dynamical fields also have the property that they can assume certain stable configurations that give rise to periodic processes). If, on the other hand, the metric were ontologically primitive, and not reducible to facts about matter dynamics, then, *prima facie*, it would be possible for that metric field to have different isometries from the symmetries of the matter dynamics.[10] The proponent

the language of Lange (2007) [[72]], one could say that the requirement of Lorentz invariance expresses a meta-law. A meta-law cannot be derived from mere ordinary laws (although there might be a deeper explanation of it). Even though it would serve my purposes, I hesitate to avail myself of Lange's machinery, as I fear that its operation commits one to the questionable notion of nomic necessity.' Frisch coherently occupies an intermediate position. Like Janssen, Frisch agrees that the Lorentz-invariance of all dynamical laws is indicative of a meta-lawlike constraint, but unlike Janssen, contends that 'once we have found a common constraint on physical theories, we are not required to offer a further explanation of this constraint ... [T]he realization that all our dynamical laws are Lorentz-invariant does not automatically generate a demand for a further explanation of that principle, as Janssen suggests' [47, p. 182].

[10] This possibility is what leads to methodological principles like the one suggested by Earman in [36, Ch. 2]—that the isometries of the metric of spacetime (or, more generally, symmetries of objects that play sufficiently similar roles to metrics, in classical spacetimes, like h^{ab} and t_a) be made to match the symmetries of the dynamical laws of matter fields.

of the geometrical view would then need to postulate a further principle that accounts for this coincidence while the proponent of the dynamical approach would not.

The dispute in general relativity

The gravitational interaction is pervasive, in a manner reminiscent of the pervasiveness of spacetime structure. It does not merely affect all matter fields, it also affects those matter fields *in the same way*. Thus Carroll infers:

> The idea is simply that something so universal as gravitation could be most easily described as a fundamental feature of the background on which the matter fields propagate, as opposed to as a conventional force. [30, p. 151]

Vizgin:

> The basic feature of general relativity that distinguished it sharply from all other physical theories...was the inherent idea of the geometrization of a physical interaction (the gravitational interaction). [115, p. xii-xiii]

And Misner, Thorne and Wheeler:

> "I weigh all that's here" is the motto of spacetime curvature. No physical entity escapes this surveillance. [79, p. 475]

These remarks all attest to a common intuition—that the metric field of general relativity 'sees' all other matter, and affects all matter in *the same* way.But the dynamical approach is often characterised as entailing an ontological reduction of the metric in special relativity. So construed, it is unclear how it could have implications in the context of general relativity. If we are to trace the consequences of the dispute between the dynamical and the geometrical approaches outside the context of special relativity, then both positions need to be re-characterised. Brown's suggestion is to recast the debate as one over the chronometricity of the metric.[11]

Brown's strategy in general relativity is what motivates my interest in construing the debate as being, at its heart, about the origin of chronometricity of the metric. He makes use of a principle that guarantees that general relativity, in an appropriate limit, can be approximated to arbitrary accuracy by special relativity. This argument for the chronometricity of the primitive metric field g_{ab} maintains that, insofar as, in a suitably small region, g_{ab} approximately shares certain features relevant to chronometricity with the *non-primitive* Minkowski metric, g_{ab} has chronometric significance. The principle that Brown claims is needed is the *strong equivalence principle (SEP)*:[12]

> There exist in the neighbourhood of each event preferred coordinates, called locally inertial at that event. For each fundamental non-gravitational interaction, to the extent that tidal gravitational effects can be ignored, the laws governing the interaction find their

[11] A radical Machian, for example, might want to resist this inference to the existence of g_{ab} and instead deny that vacuum solutions correspond to physically possible universes, in the same way as Gödel universes, for example, are ruled out by some as candidates for physically possible universes (Maudlin argues for this on grounds of causal determinism in [77, Ch. 6]; on similar grounds, Earman [37] and Joshi [64] argue for the restriction to globally hyperbolic spacetimes. For a discussion of these arguments, see [76]). Such a view would, potentially, allow for a reductive story to be told in general relativity as well. While this move is available, I shall not discuss it further, and instead focus on the Brown-style proponent of the dynamical approach who does not make such a move.

[12] Two brief comments are in order here: (i) the SEP only guarantees g_{ab}'s chronometric significance locally. The origin of the global chronometric significance of a particular g_{ab} is still mysterious. (ii) Although the intuitions behind the postulation of an equivalence principle are relatively straightforward—how does special relativity approximate general relativity in small enough regions?—there is little concensus on how to state such a principle precisely. For analysis of different conceptions of equivalence principles in relativity, see [51, 73, 82].

> simplest form in these coordinates. This is their special relativistic form, independent of space-time location. [18, p. 169]

A good deal turns on what counts as 'their special relativistic form'. But the overall structure of Brown's argument is as stated above—the autonomous g_{ab} field in general relativity gains chronometric significance through the SEP (along with other facts about dynamical fields). This leaves two unexplained facts about the world, first, that all the non-gravitational fields are governed by dynamical equations that are invariant under the same (non-trivial) symmetry group (following [97], refer to this as **MR 1**) and second that the local isometries of the autonomous dynamical field of general relativity coincide with those of the Minkowski metric tensor (referred to as **MR 2** [97]).

We can now state precisely the two main facets of the dynamical approach, the denial of both of which are constitutive of the geometrical approach:

(1) Fixed fields (i.e. non-dynamical, and identical across all *kinematically possible models*[13]), such as the Minkowski metric field η_{ab} of special relativity, are to be ontologically reduced to the symmetries of the dynamical equations governing matter fields (the dynamical view is, therefore, a modern form of relationalism—cf. [92, §6.3.2]).

(2) Ontologically autonomous metric fields, such as g_{ab} in general relativity, do not have their chronogeometic significance—i.e., are not surveyed by physical measurement apparatuses—of necessity (i.e., in all solutions of any theory in which they appear).

The disagreement between the two positions in special relativity is relatively clear: the geometrical approach—at least theversion that I understand

[13] For a definition of this term, see chapter 5.

Weatherall, Friedman and Maudlin to endorse—entails a claim about what one might call irreducible 'spacetime facts'.[14] In particular, these are facts pertaining to some ontologically irreducible entity, represented mathematically by the Minkowski metric tensor field. The dynamical approach denies that this entity is ontologically irreducible. But it is a mistake to construe this as their fundamental disagreement. The ontological dispute is just a consequence of the central dispute, which is about the origins of chronometricity, and this becomes clear in dynamical geometry theories like in general relativity. The disagreement there is not merely a generalisation of the ontological disagreement in special relativity, since both parties agree that the g_{ab} field is ontologically primitive.

The proponent of the dynamical approach has a cogent story to tell about special relativity, because (and this is a contingent fact about the world), whatever rods and clocks are built out of, their dynamical evolution ensures that they survey the Minkowski metric field. This does assume the fact that all fields are governed by laws that have the same dynamical symmetry group. Brown accepts that this is unexplained, and refers to this fact as *the big principle* [18, §8.4.1] (this is the special relativistic component of **MR 1**). This is not a huge problem, contends Brown, since all explanation has to cease at some point:

> In the dynamical approach to length contraction and time dilation ... the Lorentz covariance of all the fundamental laws of physics is an unexplained brute fact. This, in and of itself, does not count against the approach: all explanation must stop somewhere. [18, p. 143]

The dispute in general relativity is not over the ontological status of the g_{ab} field but over the origin of its chronometricity. The focus on chronometricity is the core of what Butterfield calls 'Brown's moral':

[14] This particular characterisation was suggested to me by Oliver Pooley.

> [T]he rough idea is that we should not simply postulate that a quantity in a physical theory has (chrono)-geometric significance. The point here is not just that it would be wrong to infer from a quantity's being *called* a metric that it mathematically represents (what the theory predicts about) the readings of rods and-or clocks... [W]e should [also] not infer from the fact that in the theoretical context, the quantity is mathematically appropriate for representing such behaviour, that it does. [27, p. 296]

Butterfield draws attention to two distinct points. The first is, in effect, a famous observation, originally attributed to Putnam, that (to use Lewis' more evocative prose) 'there is no semantic glue to stick our words onto their referents, and so reference is very much up for grabs' [74, p. 221]. More importantly, he highlights the fact that, as Brown argues, not all metrics are physically relevant—in fact, picking out g_{ab} as opposed to any other symmetric, rank-2 tensor as the dynamically significant metric, requires a further step. We will discover, in chapter 1, that this is, in fact, highly non-trivial.

Three roles for the metric

The metric tensor, g_{ab}, is undoubtedly the main character in general relativity—it deserves top billing. But what does it do to deserve this prominence? It depends on whom you ask. With this in mind, it is useful to make the following preliminary distinction between three roles that g_{ab} is taken to play in general relativity:[15]

A: It plays a role in articulating the dynamics of matter fields.

[15] This will be discussed again in chapter 4, where these roles are attributed to the Levi-Civita connection.

B: It plays a role in articulating the dynamics of matter fields and, in addition, ensures that symmetries of the matter fields are also symmetries of the metric *and vice versa*.

C: It plays a role in articulating the dynamics of matter fields and, in addition, ensures that symmetries of the matter fields are also symmetries of the metric and vice versa and, further, ensures that dynamical fields survey this field globally.

Consider **A**. This is a highly permissive description of a role that a mathematical object can play. It does not impose any specific requirements on the object, so in addition to metrics and connections, even topologies fall under this category. It is important, bearing in mind Brown's moral, that the mere existence of these objects in the dynamics does not mean that they play all (or any) of the roles that they could theoretically play; there is a gap here. So, for example, a connection might define a covariant derivative that appears in an equation without it playing the further role of identifying inertial trajectories (this is what happens in teleparallel gravity, for example; for a detailed discussion, see [67]).

We therefore need something stronger. Consider **B**: this requires a local symmetry coincidence of the metric and the dynamical laws. It is interesting to note (as is done in detail in chapters 1–3) that this requirement does not guarantee **C**. If one accepts **C**, then g_{ab}, in addition to playing a role in articulating the dynamical equations, is the metric with respect to whose judgements of proper time *all* ideal clocks agree.

This thesis

Chapter 1 argues against the view that the metric field of general relativity is always surveyed by ideal measuring devices. This assumption, equivalent to what is known as the 'clock hypothesis', is central to a majority of foundational work in general and special relativity. Chapter 1 demonstrates that this hypothesis is violated by light clocks in a certain class of solutions of general relativity. It discusses the operational meaning of the metric field in general relativity, and raises serious epistemological concerns regarding the status of spacetime in general relativity. This chapter is based on work done in collaboration with Niels Linnemann and James Read.

Chapter 2 introduces and elaborates on a distinction between *theoretical* and *operational* spacetime, and discusses two functionalist proposals, one each from Knox and Baker. It provides a novel analysis of the detailed inner mechanics of Knox's programme, shows that the proposal fails in certain cases, and offers a diagnosis of these failures. This chapter bridges the gap between general relativity discussed in chapter 1, and supersymmetry discussed in chapter 3. It is based on joint work with James Read.

Chapter 3 provides a brief introduction to superspace, (i) evaluates how arguments from the first two chapters might be brought to bear on assessments of the spatiotemporality of superspace, and (ii) how considerations from supersymmetry (SUSY) might mandate a change in our spatiotemporal conceptions. Philosophical intuitions are heavily influenced by, and tied to, particular frameworks of physical theories. Having a mathematically cogent, physically plausible extension to our standard notion of spacetime, thus allows us to identify the implicit and illicit assumptions present in our philosophical stances on a number of foundational issues across the discipline. It allows us to sidestep the aptly-named charge of modal provincialism, identified by Brian Pitts as 'the assumption that the whole realm of interesting theories that space-time philosophy should consider is basically like the familiar theories.'[90, p. 4]

Chapter 4 ventures a statement of the dynamical–geometrical debate in mathematical rather than metaphysical terms, using the machinery of fibre bundles, and some ideas about general relativity and Yang–Mills theories from David Wallace [119] and James Weatherall [122]. It ends with an argument against a claim Weatherall makes regarding the interpretation of gauge symmetries.

The final chapter argues against a claim made by John Norton [83], that the proponent of the dynamical approach is illicitly committed to spatiotemporal presumptions in 'constructing' spacetime from facts about dynamical symmetries.

Part I

The epistemology of spacetime geometry

CHAPTER 1

The clock hypothesis in Gödel 'spacetime'

1.1 Introduction

General relativity is our best current theory of gravitation. It is also the paradigm spacetime theory. The gravitational interaction is standardly modelled as a manifestation of spacetime curvature, in physics textbooks, popular science presentations and even many philosophical treatments.[1] One upshot of this way of thinking about general relativity is that some of the exact details of the dynamical laws of particular fields can be ignored when attempting to answer certain kinds of questions, for example, 'what is the path that a light ray would traverse in such-and-such a spacetime?' The answer, according to this view, is always 'along null geodesics of the metric, g_{ab}.'

When pushed, proponents of this view justify appeal to the behaviour of Maxwell fields (whose waves are constitutive of light) in a certain limit, the so-called *geometrical-optical* limit. With this in mind, in a recent paper, Fletcher

[1] Examples, respectively, include Padmanabhan: '[t]he gravitational field generated by matter manifests itself as the curvature of spacetime' [85], Penrose: 'a concept of 'curvature' for space-time can be used to describe the action of gravitational fields' [87, p. 207] and Torretti: 'the possible worldlines of the simplest (non-spinning, monopole) freely falling particles in general relativity are determined by the projective structure of spacetime' [111, p. 191].

proved a remarkable theorem, which he interprets as demonstrating that “for any timelike curve in any spacetime, there is a light clock that measures the curve’s length as accurately and regularly as one wishes” [44, p. 1370]. I take ‘measurement of a curve’s length’ to mean that a measuring device records intervals of proper time along the curve as given by the metric field of that spacetime (in the introduction, I referred to this as the ‘chronometric significance of the metric’); I thus take Fletcher to be advancing a proof of the *clock hypothesis* in general relativity—that is, the foundational principle in general relativity that there exist ideal clocks which can measure the proper time along their worldlines, regardless of whether or not those worldlines be geodesic.

Fletcher takes a kinematical stance on the nature of light in relativistic spacetime theories, insofar as he takes it to be a defining characteristic of light rays that they always traverse null geodesics of the metric field. This turns out to be a crucial premise in the construction of his central theorem. However, one might ask: is this strictly true, of physical light rays, constructed of Maxwell fields? If the answer to this question is ‘no’, then there remains room to doubt the practical import of Fletcher’s result.

In this chapter, I discuss a recent result by Asenjo and Hojman [7], which demonstrates that, at least in minimally coupled Maxwell theory in curved spacetime, the answer to the above question is indeed negative: in certain spacetimes, light rays, to the extent that they can be defined, do *not* traverse null geodesics; rather, their velocity is spacetime-dependent (that is, is dependent upon the position of the light ray under consideration). In light of this dynamical (rather than kinematical) underpinning of the behaviour of light rays, it can be argued that it is *not* the case that physical light clocks (even if idealised) can be used, in general, to measure a spacetime curve’s length arbitrarily accurately.

The structure of this chapter is as follows. In §1.2, I present Fletcher's theorem; in §1.3 I present the central result from Asenjo and Hojman, the rest of the chapter traces in detail the implications of the Asenjo-Hojman result for the clock hypothesis.

To be more specific, in §1.4 I discuss the import of these results for the clock hypothesis. Here, the central point is that the results of Asenjo and Hojman shake our confidence that certain clocks (in particular, light clocks) do indeed satisfy the clock hypothesis—*pace* Fletcher.

This done, I discuss the operational meaning, or 'chronometric significance', of the metric field in general relativity—that is, the fact that the metric field is surveyed by rods and clocks built from matter fields, in the sense of the latter reading off proper time intervals given by the metric field. Since the results of Asenjo and Hojman imply that the clock hypothesis need not be satisfied by what are traditionally understood to be 'good' clocks (*viz.*, light clocks), a broader problem arises for the operational meaning of the metric field in general relativity: it is much harder for the metric field to be surveyed by matter fields, and thereby to acquire chronometric significance, than has hitherto been appreciated. This has consequences in particular for advocates of the 'geometrical approach' to spacetime theories,[2] according to whom the metric field (in some sense) *compels* configurations of all matter fields to survey its structure—the results presented in this chapter provide evidence that, even in general relativity, such is not the case. Moreover, if, following Knox [70], one takes a field's having chronometric significance to be a necessary condition for its qualifying as *spatiotemporal*, then these observations raise broader concerns regarding the status of spacetime in general relativity, some of which will be

[2] For works advocating this view, see e.g. [45, 77].

explored further in the next chapter.

1.2 Fletcher's Theorem

Fletcher situates his result in the context of Maudlin's [77, ch. 5] argument regarding the *clock hypothesis* in special relativity:

> Maudlin … has recently argued that, given some additional assumptions, one can prove that the quantity an inertially moving light clock measures in Minkowski spacetime is the proper time along its worldline … The present paper generalizes [Maudlin's] result, indicating a direction in which one can extend Maudlin's argument to light clocks undergoing arbitrary acceleration in arbitrary spacetimes. [44, p. 1370]

Anticipating the upcoming objection to Fletcher's interpretation of his theorem, I use in this chapter the term 'null clock', rather than the more tendentious term 'light clock', to refer to a Langevin clock in which the oscillating material traverses null geodesics.[3]

The structure of this section is as follows. It begins in §1.2.1 by introducing the clock hypothesis (cf. §3.2.1) and examining Maudlin's argument for its approximate validity in special relativity; this relies on a particular idiosyncrasy of Minkowski spacetime—*viz.*, its globally flat metric and affine structure. I discuss, briefly, the representation of a null clock and the manner in which it can be taken to read off intervals of proper time along its worldline. I then consider Maudlin's conditions for an ideal clock to measure such intervals even

[3] By 'Langevin clock', I mean (following [22, §2]) a clock consisting of two mirrors and an oscillating medium. The light clock is the paradigm example of a Langevin clock.

when subjected to certain forces. In §1.2.2, I introduce Fletcher's notation and recast Maudlin's argument in his terms. In §1.2.3, I remove the restrictions that Minkowski geometry imposes, and present Fletcher's result in its original context—general relativity.

1.2.1 Maudlin on the clock hypothesis in special relativity

Maudlin states the clock hypothesis[4] as follows:

> The amount of time that an accurate clock shows to have elapsed between two events is proportional to the [i]nterval along the clock's trajectory between the two events ... [77, p. 75]

A significant point to note about this hypothesis is that it makes a claim regarding clocks *in all states of motion*—not just inertial. Maudlin argues for the validity of the clock hypothesis in special relativity for a particular class of clocks—*viz.*, null clocks. Let me review his construction.

Consider a smooth, paracompact, Hausdorff, Lorentzian metric manifold $\langle M, g_{ab}\rangle$.[5] In this subsection, we set g_{ab} to be η_{ab}, the flat Minkowski metric field of special relativity. Define two timelike (with respect to η_{ab}) curves $\gamma : I \to M$ and $\overset{\alpha}{\gamma} : J \to M$, where I and J are some open intervals on the real line, $\mathbb{R}$.

The Minkowski spacetime representation of a null clock consists of two material 'mirror' wordlines, represented by $\gamma[I]$ and $\overset{\alpha}{\gamma}[J]$, and a massless particle

[4] As Maudlin himself points out, 'hypothesis' is something of a misnomer for this statement. It is more accurately seen as leading to a definition of a clock, as a physical body which measures the proper time along its worldline. For a contrary attitude, see [97, §III.C], in which it is argued that no physical clock can be regarded as satisfying exactly the clock hypothesis—for all clocks will break eventually, when subject to sufficiently great accelerations. For these authors, what is of relevance is not exact satisfaction of the clock hypothesis, but rather approximate satisfaction, to some requisite degree of accuracy—this is what they dub the 'clock condition'. I return to these matters in §1.4.

[5] Here I use abstract indices, in order to align with Fletcher. In the following sections of this chapter, when performing physical calculations, I switch to coordinate indices. In addition, unless otherwise stated, in this chapter, I work in natural units in which $c = G = 1$.

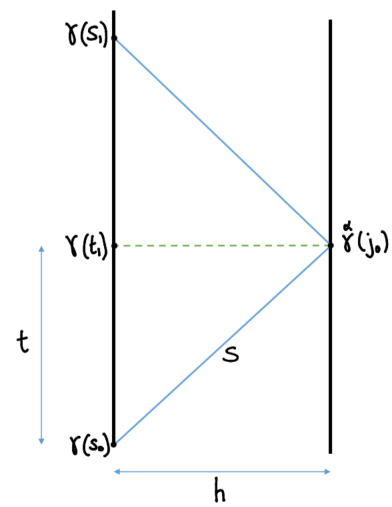

Figure 1.1: A null clock in an inertial frame of special relativity.

bouncing between $\gamma[I]$ and $\overset{\alpha}{\gamma}[J]$, the trajectory of which is represented by a series of null geodesics.[6] Call the trajectory of such a particle between two successive bounces a 'null ray'. In a Lorentz coordinate frame,[7] a configuration in which the null clock is at rest can be represented as in figure 1.1.

Introducing some useful notation, consider the closed interval $[I^1] = [s_0, s_1]$ of the domain of γ; we can safely assume that the intervals of parameter values are intervals of proper time along a worldline. The 'halfway point' on the image of γ (assessed with respect to η_{ab}) corresponds to $t_1 \in [I^1]$. Let h be the spatial distance, in a particular frame, between a point on $\gamma[I]$ and a point on $\overset{\alpha}{\gamma}[J]$, and let t be the proper time between s_0 and t_1. By construction, this makes t the proper time between t_1 and s_1 as well. Let S be the spatiotemporal distance between $\gamma(s_0)$ and $\overset{\alpha}{\gamma}(j_0)$.

[6] Both Maudlin's and Fletcher's discussions are situated wholly within the context of classical general relativity. I therefore use the word 'massless particle' not to represent a quantum of some quantised field, but rather to represent the geometrical-optical limit of a classical field defined on a spacetime manifold. For more on the geometrical-optical limit, see [54, ch. 10], [79, pp. 570-583], [89, ch. 5-6], and the extensive discussion below.
[7] That is, an inertial frame of special relativity, in which the laws of physics are understood to take their simplest form.

Minkowski geometry tells us that

$$S^2 = -t^2 + h^2. \tag{1.1}$$

By construction, the spatiotemporal distance between $\gamma(s_0)$ and $\gamma(s_1)$ is zero, because the spatiotemporal distance between any two points on a null ray is always zero. This means that

$$t^2 = h^2. \tag{1.2}$$

So, if we shoot a massless particle from $\gamma(s_0)$, the proper time elapsed on $\gamma[I^1 \subsetneq I]$ between $\gamma(s_0)$ and $\gamma(s_1)$ is just $2h$. Clock 'ticks' correspond to points on $\gamma[I]$ where the oscillating particle 'bounces off' $\gamma[I]$ and travels towards $\overset{\alpha}{\gamma}[J]$. Label these as $\gamma(s_2)$, $\gamma(s_3)...\gamma(s_n)$, where $n = \overset{\alpha}{n} \in \mathbb{N}$. The number of such points is referred to as the 'bounce number' for a segment of the trajectory, and denoted by $\overset{\alpha}{n}$. If we extend our interest to a larger segment $\gamma[I' \subset I]$ (where $[I' \supset I^1]$) of the image of $\gamma[I]$, then we will discover more points that correspond to clock ticks. The proper time elapsed on this larger segment of γ will then be

$$|I'| = |s_0 - s_n| = 2\overset{\alpha}{n}h. \tag{1.3}$$

What if we were to physically push the null clock, so that it changes its trajectory to one with a non-zero constant velocity, with respect to our original Lorentz coordinates? Call such an action a 'Lorentz push'—a Lorentz push thus implements an active Lorentz transformation. On intuitions imported from classical spacetimes like Galilean spacetime, it might seem reasonable to assume that the spatial distance between the mirrors be preserved after the push, as measured in the original coordinate system. This hypothetical situation

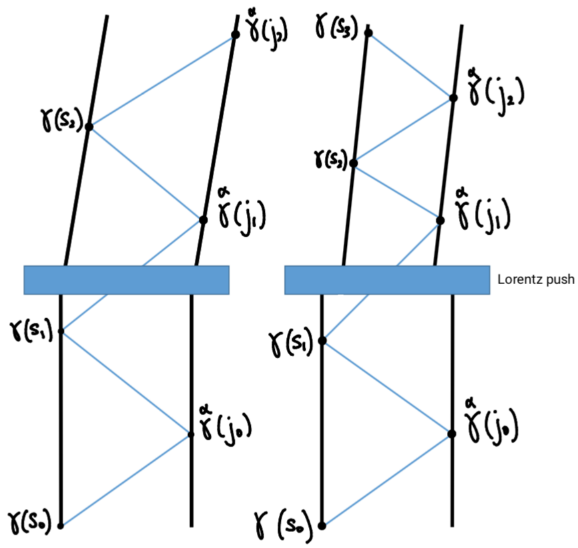

Figure 1.2: Two configurations for possible null clocks, before and after a Lorentz push. (The region over which the Lorentz push is implemented is shaded blue.) The clock on the left violates the relativity principle, for it ticks at different rates before and after the Lorentz push; the clock on the right satisfies the relativity principle, for it ticks at the same rate before and after the Lorentz push.

is represented on the left of figure 1.2. If such were the case, since our oscillating material traverses null rays, the post-Lorentz push clock would tick at a slower rate. This particular clock, then, would violate the relativity principle, for it would tick at different rates in different inertial frames.

The way out of this conundrum is well-known: as Maudlin states [77, pp. 112-113] (essentially following the moral of Bell's rocket experiment [14]), a Lorentz transformation will bring the two mirrors together. (How exactly Maudlin argues to this conclusion is elaborated below.) Thus, the post-Lorentz push state of the clock under consideration will appear as on the right hand side of figure 1.2. As a result, the null clock *will* continue to tick at the same rate, and so *does* satisfy the relativity principle. (In all considerations up to this point, I ignore the behaviour of the clock during the period of its acceleration over the course of the Lorentz push; we return to these matters shortly.)

In order to demonstrate the 'approximate validity of the clock hypothesis', Maudlin's argument proceeds in two phases: (A) demonstrate that inertially moving clocks measure intervals of proper time along their worldlines as given by the metric field, even after a Lorentz push; (B) demonstrate that clocks in arbitrary states of motion approximately measure intervals of proper time on their worldlines as given by the metric field.

To achieve (A), Maudlin constructs a null clock by attaching the two mirrors to opposite ends of a rigid rod. A rod is said to be 'rigid' just in case, after a Lorentz push, it assumes a configuration which, when expressed in the coordinates of the boosted frame (i.e., its rest frame after the Lorentz push), is identical to its configuration in the coordinates of the original frame, before the Lorentz push.[8] The state of a rigid rod in its rest frame is its 'equilibrium state'. The result of this construction is that the Lorentz push will cause the null clock to contract. The consequence, of course, is that such clocks satisfy (A), as given above.

Now to (B). If we restrict our attention to segments of the null clock's worldline in which the rigid rod is in an equilibrium state, then Lorentz pushes do not affect the clock—in the sense that our story regarding the behaviour of the clock is restricted to a story regarding its behaviour in its equilibrium states, in which (as we have already seen) the clock *does* read off intervals of proper time along its worldline. In the limit that arbitrary motion can be approximated as being composed of segments of inertial motion separated by arbitrarily short Lorentz pushes, Maudlin claims to achieve (B): a demonstration of the satisfaction of the clock hypothesis by null clocks in arbitrary states of motion.[9]

Fletcher [44, pp. 1370-1371] correctly and reasonably points to the restricted

[8] Thus, a rigid rod would *not* be distorted by a Lorentz push in the manner of the null clock on the left of figure 1.2, discussed above.

[9] Note that this now includes the blue regions in figure (1.2), in which the clock is accelerated in order to implement the Lorentz push.

scope of Maudlin's argument: it relies on the notion of an equilibrium state of a rigid rod, and, moreover, is only approximate, in its application to null clocks in arbitrary states of motion. Motivated by these concerns, Fletcher generalises Maudlin's argument to account (he claims) for the satisfaction of the clock hypothesis by null clocks in arbitrary states of motion in both special and general relativity.

1.2.2 Fletcher's result in special relativity

Let me now tell the first part of Maudlin's story from Fletcher's more general perspective, using his notation. Consider the tangent space T_pM at the point $p = \gamma(t_i) \in M$. Since Minkowski spacetime is a globally flat metric manifold, with a global smooth atlas, for an element $\rho^a \in T_pM$, the 'exponential map', $\exp_p(d\rho^a)$, where $d \in \mathbb{R}$, defines a smooth curve.[10]

If a normal open neighbourhood U_0 of T_pM can be exponentiated to recover an open neighbourhood U of $p \in M$, then U is known as a 'simply convex neighbourhood' [44, p. 1372].

For a timelike curve such as $\gamma[I]$, one can, at each point, identify an orthogonal unit spatial vector ρ^a, thus defining a vector field on $\gamma[I]$. Each vector can be exponentiated to map a point on $\gamma[I]$ to a point on another timelike curve, $\overset{\alpha}{\gamma}[J]$. If $\gamma[I]$ is an inertial trajectory, then multiplying each unit spatial vector by the same constant, $\overset{\alpha}{d}$, one can map $\gamma[I]$ to another timelike inertial trajectory. Each choice of α therefore generates a different parallel curve.

Call the exponential factor, $\overset{\alpha}{d}$, the 'scalar radius' of the curve $\overset{\alpha}{\gamma}$. In Minkowski

[10] The exponential map at $p \in M$, $\exp_p : U_p \to M$, is defined on a subset U_p of the tangent space T_pM as follows. First, $0 \in U_p$ and $\exp_p 0 = p$. Then any nonzero $\alpha^a \in U_p$ if and only if there is a geodesic $\gamma : [0, 1] \to M$ with tangent vector α^a at p such that $\gamma(0) = p$. Finally, for such nonzero $\alpha^a \in U_p$, $\exp_p \alpha^a = \gamma(1)$, which is well-defined since the geodesic γ corresponding to α^a is unique [44, pp. 1371-1372].

spacetime, this can be taken to be equal to the spatial distance between the points $\gamma(t_1)$ and $\exp_p(\overset{\alpha}{d}\rho^a) = \overset{\alpha}{\gamma}(j_0)$.[11,12] So, (1.3) can be rewritten as

$$|I'| = |s_0 - s_n| = 2\overset{\alpha}{n}\overset{\alpha}{d}. \tag{1.4}$$

1.2.3 Fletcher's theorem in general relativity

Maudlin's argument applies only to Minkowski spacetime. But this can be generalised to arbitrary Lorentzian manifolds. To do so, drop the requirement that the metric field be flat (indeed, also drop the requirement that it have constant curvature). As a result, global simple convexity is lost—i.e. the property that M is itself a simply convex neighbourhood. However, since the manifold is Lorentzian, we still have 'local simple convexity',[13] $\forall p \exists U_0 \subset T_pM, d \in \mathbb{R}$: $\exp_p(dU_0) = U_n \subset M$, and $\bigcup_n U_n = M$. If M is not flat, then $U \subsetneq M$ is an open neighbourhood of $p \in M$. In general, it is now only possible to recover local patches of the manifold by exponentiating element of the tangent space at a point.

For a particular family of so-called 'convergent companion curves', $\{\overset{\alpha}{\gamma}\}_{\alpha\in\mathbb{N}}$, Fletcher's theorem makes two assertions, call them **accuracy** and **regularity**:

Accuracy: $\lim_{a\to\infty} 2\overset{\alpha}{n}\overset{\alpha}{d} = |I'|$.

Regularity: $\limsup_{\alpha\to\infty} \left\{ |(\overset{\alpha}{s}_i - \overset{\alpha}{s}_{i-1}) - (\overset{\alpha}{s}_j - \overset{\alpha}{s}_{j-1})| : 1 \leq i, j \leq \overset{\alpha}{n} \right\} = 0$.

[11] This is because the tangent space at a point in Minkowski spacetime is isomorphic (as both a vector space and a manifold) to Minkowski spacetime. So there is a canonical one-one correspondence between vectors in T_pM and straight lines in M.

[12] We should stress, as Fletcher does, that although in this case one can think of the scalar radius as being equal to the distance between the mirrors, this is not a privileged measure. "One could very well pick some other spacelike vector field on γ and some other scalar parameter to trace out the same companion curve, and this new pair would bear a systematic functional relationship to ρ^a and $[\overset{\alpha}{d}]$. The constraints on $[\overset{\alpha}{d}]$, as determined by the theorem, would then fix constraints on this new scalar parameter" [44, p. 1381].

[13] See e.g. [84, p. 131] for a proof that all Lorentz manifolds are locally simply convex.

Accuracy is the statement that the times elapsed between 'ticks' measured on $\gamma[I]$ are proportional to proper time intervals on the worldline. **Regularity** asserts that the proper time interval measured between any two arbitrary pairs of successive ticks is the same. If we restrict our interest to physical clocks, we see that it is **regularity**, not **accuracy**, that is significant. Maudlin expands:

> An ideal clock is some observable physical device by means of which numbers can be assigned to events on the device's worldline, such that the ratios of differences in the numbers are proportional to the ratios of [i]nterval lengths of segments of the world-line that have those events as endpoints. [77, p. 106]

Accordingly, note that the null character of light rays does no operationally significant work here—we would still be able to construct ideal clocks in which the oscillating material traverses *timelike* paths. What is important is that the timelike curve $\gamma[I]$ is defined with respect to the same metric as the one surveyed by the oscillating matter—this is what ensures **regularity**. This is guaranteed by Fletcher's assumption that light travels on null geodesics of g_{ab}, but would just as easily be achieved by any oscillating material which travels at a constant velocity between the mirrors.[14]

Looking at Fletcher's theorem more closely, we see that, in general, we can still interpret $\overset{\alpha}{d}$ as being equal to the distance between the mirrors—$\overset{\alpha}{d}$ is constrained to be non-zero [44, p. 1376]—but only for the specific configurations of the null clock discussed below. As we will see, this is an important constraint on Fletcher's theorem: the theorem is only valid for clock configurations where this identification can be made.

[14] Fletcher makes this observation at [44, p. 1382].

Given $\overset{\alpha}{d} > 0$ (i.e. is strictly greater than zero), the curve $\gamma[I]$ is *not* an element of $\{\overset{\alpha}{\gamma}\}_{\alpha\in\mathbb{N}}$. This ensures that a physical 'bounce' is always possible. The limit in the family of convergent curves ensures two things: first, that the two curves are sufficiently 'nearby' that $\overset{\alpha}{\gamma}$ can be arrived at by exponentiating the unit spacelike vector field, ρ^a, on $\gamma[I]$—in other words, it ensures that they are in the same non-disjoint union of simply convex neighbourhoods. Second, that one can describe light clocks that are sufficiently small that changes of spatial length between mirrors does not undermine the clock's accuracy and regularity. In Minkowski spacetime, global simple convexity ensures that these two curves can be arbitrarily far apart, and the null clock (in principle) still functions as an ideal clock.

Since spacetimes in general relativity can be arbitrarily curved, the bounce number, $\overset{\alpha}{n}$, for a given timelike curve $\gamma[I]$ will, in general, depend upon the path traversed by the massless particle between $\gamma[I]$ and (the image of) the selected companion curve, $\overset{\alpha}{\gamma}[J]$. In order to avoid this dependence, one needs to be careful about *which* companion curve one chooses. This is where local simple convexity comes in. In regions in which it is possible, for any timelike curve, $\gamma[I]$, to find a companion curve, $\overset{\alpha}{\gamma}[J]$, which can be reached by exponentiating the spacelike tangent field on $\gamma[I]$, Fletcher's theorem holds.[15] In such regions, the scalar radius $\overset{\alpha}{d}$ *does* approximate the spatial distance between exponential map-related points across the curves. The fact that every point on a Lorentzian metric has a local simply convex neighbourhood guarantees that Fletcher's theorem holds for all Lorentzian spacetimes.

The physical interpretation that Fletcher gives of his use of the local simple

[15] Technically, what is required is that for any segment, $\gamma[\overset{k}{I}]$, of the curve $\gamma[I]$, such that $\bigcup_k \gamma[\overset{k}{I}] = \gamma[I]$, there exists a $\overset{\alpha}{\gamma}[\overset{m}{J}]$ such that $\exp_p(\overset{\alpha}{r}\rho^a) \approx \overset{\alpha}{\gamma}[\overset{m}{J}]$, $\forall p \in \gamma[\overset{k}{i}]$ and $\bigcup_m \overset{\alpha}{\gamma}[\overset{m}{I}] = \overset{\alpha}{\gamma}[I]$.

convexity assumption is the following: it allows a light clock to expand and contract arbitrarily as it moves through the manifold, thereby accounting for possible gravitational tidal forces in generic general relativistic spacetimes (as well as due to acceleration and internal forces). As long as the clock is sufficiently small that both mirrors are always within the same simply convex neighbourhood, Fletcher argues, it will measure proper time along $\gamma[I]$, and thereby satisfy the clock hypothesis. One way of putting our disagreement with Fletcher is the following: we argue (see §1.3) that the light rays *themselves* in physical light clocks will manifest spacetime-dependence. This dependence is what spoils the ideality of a light clock.[16,17]

Let us return now to Maudlin's null clock. Recall that the rigid rod between the mirrors ensures that their spatial separation, as measured in the frame in which the clock is at rest just before the push, is less after the push than before. But, as mentioned earlier, Maudlin says nothing about what goes on during the Lorentz push and in the time that it takes for the rod to reach its equilibrium state post-push. Therefore, his clock can only be guaranteed to measure proper times on trajectories to within the level of accuracy that the rigidity of the rod alone guarantees. One might abstract away from the use of the rigid rod, and instead impose a 'clock constraint', which restricts us to clocks whose mirrors' spatial separation changes in accordance with what an idealised rigid rod would have imposed on their configuration. But then one is presented with the non-trivial problem of showing that such systems exist.

In effect, the simple convexity of the open neighbourhoods around points on

[16] It is important to note that we are conceding to Fletcher many things—e.g. that it is possible to find a physical medium which emulates the behaviour of the mirrors of his light clock. In my view, there is good reason to doubt whether any such medium can be found—though I set the matter aside in the remainder of this chapter.

[17] Note that Maudlin's sense of the ideality of a clock is different from that deployed in footnote 26.

timelike trajectories in Lorentzian manifolds is what allows Fletcher to impose the clock constraint, and with it to prove the clock hypothesis for null clocks. More precisely, since, *ex hypothesi*, light travels on null geodesics and massive particles on timelike geodesics of g_{ab}, Fletcher's theorem proves that, within a given simply convex neighbourhood of a point on the manifold, an abstract version of Maudlin's construction holds, and any timelike trajectory *can* be approximated as a series of inertial trajectories linked by Lorentz pushes.

1.3 Electromagnetism and the Geometrical-Optical Limit

As discussed, Fletcher's assumption that light propagates on null geodesics—call this the 'relativistic null hypothesis'—is central to his theorem, for it suffices to ensure **regularity**. It is often taken for granted that the relativistic null hypothesis is satisfied in the 'geometrical-optical limit' of general relativity, in which, roughly speaking, the length scale over which the wavelength of the wave under consideration changes is much shorter than the length scale over which curvature changes in the ambient spacetime (cf. [79, §22.5]). In this section, I reconsider the physicality of the geometrical-optical limit, and thereby whether is it indeed true that light rays, *qua* solutions to Maxwell's equations in curved spacetime, invariably traverse null geodesics.

It turns out that the answer is 'no'—though a full explication of this answer will require some setup. First, some remarks on the field of optics. Research in this area is concerned with aspects of the behaviour of electromagnetic waves that are determined by how those waves propagate in the domain in which they can be approximated as rays. In this approximation, information about interference and diffraction is discarded, but, in compensation, one is able to derive a plethora of theorems and results regarding the observed behaviour of light. The approximation here is that the wavelength of the light being studied is significantly smaller than (a) the distances over which its amplitude varies, and (b) the distances over which curvature effects are non-negligible. Given that the wavelength of light studied in the context of optics is of the order of 100-1000nm, and typical experiments are conducted over distances at least 10^4 times as large, this optical approximation is generally a reasonable one.

In what follows, I focus on the behaviour of light in the limit that (a) and (b) are always satisfied. Since, for light of any wavelength, one can always find a Lorentzian manifold whose curvature varies on a scale comparable to that wavelength, the only way of guaranteeing the generality of a statement (such as Fletcher's) made on the basis of such optical approximations, is by considering light in the limit that its wavelength tends to zero (equivalently, its frequency tends to infinity), and its amplitude is constant—*this* is what is meant, precisely, by the 'geometrical-optical limit'.

It is worth pausing here, briefly, to discuss the physical intuitions behind taking this limit. Recall from §1.2 that what the construction of a null clock is intended to capture is the periodic motion of a particle bouncing back and forth between two mirrors. In classical physics, although light is not described by a particle, it is common to refer to the 'path traversed by light'. The standard interpretation of this locution is in terms of some limit. To the extent that we can interpret a solution to Maxwell's equations as instantiating a wave packet that is sufficiently localised and dynamically robust (i.e., continues to exist as a wave packet on the time scales of interest), we can talk about a path traversed by light. Moreover, the wave packet can be approximated as a *plane-wave* solution—the divergence in behaviour between the two types of solution only manifests itself on large length and time scales.

By choosing a suitably short wavelength, then, this wave packet can be localised to well below the length scale of interest. The trajectory of the wave packet can then be thought of as an integral curve to some vector field—these integral curves are the light rays. Call each element of this vector field a 'wave vector', k^μ. By construction, the wave vector always points in the direction of propagation of the wave in spacetime. In the plane-wave approximation, the wave vector is perpendicular to the 'wave front'—the line or surface connecting all points on a wave that are in phase. Therefore, if we know the wave vector field and the metric, we can calculate the wave fronts, which then constitute the one-form, k_μ, which is orthogonal to the vector field.

We can now reverse this construction. If we begin with the wave fronts, we can therefrom construct the wave vector field, and with it the integral curves that represent the wave packet's trajectory. On this picture, it is clear that a coherent picture of light as traversing *any* kind of path presupposes a robust notion of a wave front. Therefore, the solutions to Maxwell's equations that describe light as traversing paths have to be such that the wave fronts of the initial configuration are preserved under the dynamical evolution of the field. In Minkowksi spacetime, it is easy to come by such solutions—all plane wave solutions, for example, have this feature.

A plane wave in Minkowski spacetime satisfies what I refer to in §1.3.1 as a 'standard wave equation'—these are equations whose solutions are (possibly infinite superpositions of) plane waves, whose wave fronts move at a constant velocity, thus defining a constant wave front one-form k_μ. As a result, the wave vector field is also constant, since the metric is constant. Therefore, the following equation, known as a 'dispersion relation', describes the propagation of light in Minkowski spacetime:

$$\eta_{\mu\nu} k^{\mu} k^{\nu} = 0. \tag{1.5}$$

The wave vector is determined by the Maxwell wave equation (discussed below), which itself depends on the background geometry, and from the dispersion relation it is clear that the wave vector k^{μ} is null—so light propagates on null geodesics.

When considering the propagation of light in curved spacetimes, we still need to hold onto the notion of wave front preservation. But, as will be shown explicitly in §1.3.2, the generalisation of the notion of a plane wave in Minkowski spacetime to arbitrarily curved spacetimes (call them 'generalised plane waves') loses this important feature. It is therefore crucial to discover whether there exists some approximation in which this feature is still valid—this is what motivates the use of the geometrical-optical limit. Note that, in a generically curved spacetime, neither the metric field nor the one-form k_{μ} corresponding to a solution of the curved-spacetime Maxwell equations is, in general, constant across the manifold. More importantly, these k_{μ} need not even carry a straightforward interpretation as wave fronts of a propagating wave packet.

In generically curved spacetimes, call an equation of the following form a 'generalised dispersion relation':

$$g_{\mu\nu}(x) k^{\mu}(x) k^{\nu}(x) = f(x), \tag{1.6}$$

where $f(x)$ is some function of spacetime.

In the limit that the wavelength of light tends to zero, it can be shown that the evolution of generalised plane waves does preserve (at least locally) wave fronts and wave vectors [79, pp. 574-575]. In this limit, we can speak meaningfully of the trajectory of light. In order for it to be the case that these light rays are *null*, a further condition needs to be met,

$$g_{\mu\nu}(x)k^{\mu}(x)k^{\nu}(x) = f(x) = 0. \tag{1.7}$$

Call such a dispersion relation a ‘null dispersion relation’. In a certain class of rotating spacetimes, an example of which is discussed in §1.3.3, if we consider an arbitrary solution to the vacuum Maxwell equations, which does not describe the propagation of light *rays*, and then take the geometrical optical limit, the result is *not* guaranteed to be a solution to those vacuum Maxwell equations. In other words, $f(x)$ is not guaranteed to vanish for the dispersion relation associated with exact generalised plane wave solutions that preserve wave fronts, even though it might do so in the dispersion relation arrived at from geometrical-optical limit of an approximate solution. Although Fletcher never explicitly mentions geometrical optics, his assumption of the validity of the relativistic null hypothesis, irrespective of the background spacetime in which the wave is embedded, is equivalent to this assumption.

The purpose of this section, therefore, is to argue against Fletcher’s use of this assumption, by demonstrating that, in a certain class of spacetimes, there is no solution of the Maxwell equations which gives rise to a null dispersion relation. I will show that this leads to a violation of the condition of **regularity**, and this is ultimately what undermines the purported generality of Fletcher’s theorem. In §1.3.1, I introduce the vacuum Maxwell equations in curved spacetime, which describe the propagation of light. We then discuss, in §1.3.2, how the geometrical-optical limit can—if physically appropriate—demonstrate that light rays traverse null geodesics. Finally, I turn, in §1.3.3, to the vacuum Maxwell equations in certain rotating spacetimes. Such solutions deliver a startling verdict on the behaviour of light rays—that they violate the relativistic null hypothesis. I should stress that, in this section, my concern is with certain spacetimes violating the relativistic null hypothesis due to the misalignment between approximate and exact solutions that arise from *rotation*, rather than curvature effects.

1.3.1 Maxwell's equations in curved spacetime

Let us distinguish between two types of 'wave equation'—a 'standard wave equation' and a 'Maxwell wave equation'. A standard wave equation for, say, a scalar field ψ takes the form

$$\frac{\partial^2\psi}{\partial x^2} - \frac{1}{c^2}\frac{\partial^2\psi}{\partial t^2} = 0, \tag{1.8}$$

where c is the wave propagation speed. The most general solution to this kind of equation is a superposition of 'plane wave solutions', i.e. solutions of the form

$$\psi\left(x,\, t\right) = \psi_0 \exp\left\{i\left(\vec{k}\cdot\vec{x} \pm \omega\cdot t\right)\right\}, \tag{1.9}$$

where ψ_0 is a constant. The associated four-dimensional wave vector is $k^\mu = (\omega, \vec{k})$, and the four-dimensional wave front is $k_\mu = \eta_{\mu\nu}k^\nu$. From this, we see that a general formula for a four-dimensional wave vector is

$$k^\mu = \nabla^\mu\theta, \tag{1.10}$$

where θ is the phase of the wave and ∇^μ is the Minkowski metric-compatible derivative operator. This definition of the wave vector applies even to position-dependent phases. For any given plane wave, therefore, it makes sense to talk about a trajectory—it is just the integral curve associated with k^μ. Since k^μ is a constant, these integral curves are a family of straight lines in Minkowski spacetime corresponding to trajectories of rays travelling at velocity c.

The Maxwell equations in curved spacetime are given by[18]

$$\nabla_\mu F^{\mu\nu} = 0, \tag{1.11}$$

$$\nabla_\mu F^{*\mu\nu} = 0, \tag{1.12}$$

where ∇_μ is the derivative operator compatible with a generic Lorentzian metric field $g_{\mu\nu}$ (with respect to which index contraction is performed),

$$F_{\mu\nu} := \nabla_\mu A_\nu - \nabla_\nu A_\mu = \partial_\mu A_\nu - \partial_\nu A_\mu \tag{1.13}$$

is the antisymmetric Faraday tensor defined in terms of the electromagnetic four-potential A_μ,[19] and $F^{*\mu\nu} := \frac{1}{2}\epsilon^{\mu\nu\alpha\beta}F^{\mu\nu}$ the dual of $F^{\mu\nu}$.[20]

As we are dealing with the source-free Maxwell equations, (1.12) amounts to

[18] Throughout this section, I switch to (Greek) coordinate indices, since it is convenient to perform calculations in a particular coordinate basis.

[19] (1.13) holds in any coordinate basis, since connection components in this equation vanish by the symmetry of the connection. For details, see e.g. [35, p. 39].

[20] Of course, in general relativity, the dynamical equations for the $g_{\mu\nu}$ field are the Einstein field equations, which take the form

$$G_{\mu\nu} := R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = 8\pi T_{\mu\nu},$$

where $G_{\mu\nu}$ is the Einstein tensor, $R_{\mu\nu}$ is the Ricci tensor, R is the Ricci scalar, and $T_{\mu\nu}$ is the stress-energy tensor associated with the matter fields that serve as sources for the gravitational field. The contribution of the electromagnetic field to the energy-momentum tensor is given by

$$T^{\mu\nu}_{\text{Maxwell}} := F^{\mu\lambda}F^{\nu}{}_{\lambda} - \frac{1}{4}g^{\mu\nu}F_{\lambda\rho}F^{\lambda\rho}.$$

These equations are called the 'Einstein-Maxwell equations'. In this chapter, I largely drop consideration of the Einstein equations, given the irrelevance of back-reaction effects of the electromagnetic field on the metric field.

nothing but a Bianchi identity. So, the only equations that are not identically satisfied are (1.11); using (1.13), we have (cf. [79, p. 569])

$$\nabla_\mu \nabla^\mu A_\nu + R_\nu{}^\mu A_\mu = 0. \tag{1.14}$$

This is the second type of wave equation discussed in this chapter— a Maxwell wave equation.

Consider a solution of the Maxwell wave equation in curved spacetime of the following form—call it a 'generalised plane wave':

$$A_z(t,x) = \xi(x) \exp\{iS(x) \pm i\omega \cdot t\}. \tag{1.15}$$

Associated with this solution are $k^\mu = (\omega, \partial_i S(x))$, and $k_\mu = g_{\mu\nu} k^\nu$, both of which are spacetime dependent. Moreover, given the spacetime dependence of k^i, in general this will not have an interpretation as a wave vector for a wave packet state, since such states do not, in general, exist in curved spacetimes. In order for wave packet states to be guaranteed to exist, it must be the case that the spacetime-dependence of the amplitude and phase are negligible. So we go to the geometrical-optical limit in which ξ, S and $g_{\mu\nu}$ are approximately constant. In this case light rays are solutions to a standard wave equation, so guaranteed to travel at constant velocity, although the claim that this velocity is c requires further argument.

1.3.2 The geometrical-optical limit

The statement that light, even in arbitrarily curved spacetimes, can be taken to traverse null geodesics in the geometrical-optical limit[21] relies on an insuf-

[21] Such a statement is certainly pervasive in the literature. See, for example, Wald: "[I]n this approximation (known as the *geometrical optics approximation*), light travels on null geodesics,

ficiently general assumption about the behaviour of solutions to Maxwell's equations in curved spacetime. On this limit, Misner *et al.* state that

> [t]he fundamental laws of geometric optics are: (1) light rays are null geodesics; (2) the polarization vector is perpendicular to the rays and is parallel-propagated along the rays; and (3) the amplitude is governed by an adiabatic invariant which ... states that the number of photons is conserved. [79, p. 571]

They are careful to stress, however, that these laws are *derived* from the basic assumption of geometrical-optics (*viz.*, (a) and (b)) mentioned above. The relativistic null hypothesis (that is, Misner *et al.*'s (1)), follows from the basic assumptions of the geometrical-optical limit, only when one further important criterion is met—that the solutions arrived at in this limit do, in fact, approximate exact solutions to arbitrary accuracy. In this section, beginning with the pedagogical setup of Misner *et al.*, we discover the conditions under which the relativistic null hypothesis is valid.

Recall that a light ray is, by definition, a curve that is perpendicular to a surface of constant phase (i.e., a wave front) [79, p. 573]. In what follows, I will describe in detail the behaviour of approximate solutions that allow us to recover the ray-like characterisation of light familiar from optics—perhaps the easiest way to think about the approximation at play is to see it as being motivated by a desire to make the curved spacetime model resemble the flat spacetime model by considering progressively smaller regions of spacetime.

a suggestion which can be confirmed by studies of the Green's function" [116, p. 71]. Or Malament: "[The behaviour of] light [is determined by] the behaviour of solutions to Maxwell's equations in a limiting regime ("the optical limit") where wavelengths are small ... [W]hen one passes to this limit, packets of electromagnetic waves are constrained to move along (images of) null geodesics" [75, p. 147].

Begin with the picture in flat spacetime—here the electromagnetic equations are just the standard Maxwell equations (1.11) and (1.12), defined with respect to the fixed Minkowski metric field $\eta_{\mu\nu}$.

Consider a vector potential A^μ, which is a solution to (1.14). As a solution to a wave equation, it can always be decomposed into an 'amplitude' piece α^μ, and a phase piece $\theta \propto \frac{l}{\lambda}$, where l is the distance propagated by the wave and λ is its wavelength Thus,

$$A^\mu = \mathrm{Re}\left(\alpha^\mu e^{i\theta}\right). \tag{1.16}$$

Since we are in Minkowski spacetime, the curvature of the manifold, by definition, does not vary with distance. Consider a wave-like solution the amplitude of which is constant across the manifold—such solutions are guaranteed to exist, since the Maxwell wave equation in Minkowski spacetime is a standard wave equation.

In curved spacetimes, the geometrical-optical approximation attempts to preserve the wave packet behaviour associated with certain solutions of the Maxwell equations. If we start with a specific solution to the Maxwell equations the amplitude of which does not change appreciably over the length scale of interest, and for which the ambient spacetime is approximately flat over the same length scale then, as the wavelength is decreased, these approximations apply to a larger class of spacetimes (i.e., solutions to the Einstein field equations). Eventually, in the limit that the wavelength tends to zero, and the amplitude variation becomes negligible, it seems reasonable to assume that such solutions will coincide with exact ray-like solutions of the wave equation in any spacetime.

Let us examine one of the above-discussed limits—the constant amplitude limit. A general solution to the Maxwell equations in curved spacetime will describe a wave whose phase is a function of spacetime. If we restrict attention to regions of spacetime over which the amplitude varies very slowly, then minor corrections can be made to the amplitude at every point:

$$\alpha^{\mu} = a^{\mu} + b^{\mu}\lambda + c^{\mu}\lambda^{2} + d^{\mu}\lambda^{3} + \dots . \tag{1.17}$$

Since the amplitude really depends on the choice of length scale L, the expansion is in powers of $\epsilon := \frac{\lambda}{L}$, so the vector potential in (1.16) can now be expanded as

$$A^{\mu} = \mathrm{Re}\left\{\left(a^{\mu} + \epsilon b^{\mu} + \epsilon^{2} c^{\mu} + \dots\right) e^{i\theta/\epsilon}\right\} . \tag{1.18}$$

Substituting our expansion into the source-free version of the Maxwell wave equations (1.14) and once again gathering terms that are to the order $(1/\epsilon^{2})$, we get our familiar dispersion relation, $k^{\mu}k_{\mu} = 0$:

The relativistic null hypothesis is important for Fletcher's result only to the extent that it guarantees that light travels on trajectories of constant velocity—the actual value of this velocity is largely irrelevant. It is possible, therefore, that light rays travel at constant velocities, as judged by the affine connection, even if those geodesics are not null. In the terminology of §1.2, this would lead to a violation of **accuracy** but not **regularity**. So Fletcher's theorem would remain intact. If the dispersion relation were non-trivially spacetime-dependent, on the other hand, then the integral curves corresponding to the vector field defined by the wave vectors would not be straight lines according to the connection. This is what spells trouble for **regularity**, and so also for Fletcher's theorem.

In the class of curved spacetimes considered in [7, §2] for example, *in the geometrical-optical limit*, light does propagate on null geodesics of the metric, determined by the coupled Maxwell and Einstein field equations. Consequently, in such spacetimes, both **accuracy** and **regularity** are preserved. The important assumption that allowed us to derive this result was that the perturbative expansion presented was treated as approximating the exact solution with arbitrary accuracy. As we will see in the following, this is not guaranteed to be the case.

1.3.3 Rotating spacetimes

Consider two possible means of arriving at solutions of the Maxwell wave equation (1.14) in curved spacetime for an arbitrarily small wavelength. The first is to solve it exactly by some technique. The second is to solve the wave equation for a relatively large wavelength, and then take the geometrical-optical limit. *These two techniques are not guaranteed to agree on the space of solutions.* More precisely, it is not guaranteed that a particular solution, in the geometrical-optical limit, remains a solution to the Maxwell wave equations in curved spacetime. An example from Asenjo and Hojman [7, §3], pertaining to the Gödel solution to the Einstein field equations [53], demonstrates clearly this fact.

The Gödel metric can be expressed using natural coordinates, (t, x, y, z) on $\mathbb{R}^4$ as:

$$\begin{aligned} g_{00} &= -1 = -g_{xx} = -g_{zz}, \\ g_{yy} &= -2 + 4\exp(\sqrt{2}x\Omega) - \exp(2\sqrt{2}x\Omega), \\ g_{0y} &= \sqrt{2}[1 - \exp(\sqrt{2}x\Omega)], \end{aligned} \tag{1.19}$$

where Ω is a constant related to the angular velocity of the rotating universe.[22] Consider now the z-component Maxwell wave equation (1.14) in curved space-

[22] For a discussion of the difficulties in classifying a spacetime as 'rotating', as well as further foundational details on Gödel spacetime, see [75, ch. 3].

time,

$$\partial_0^2 A_z + \frac{1}{\sqrt{-g}\, g^{00}} \partial_x \left(\sqrt{-g}\, \partial_x A_z\right) = 0. \tag{1.20}$$

For the Gödel metric, there is no choice of coordinates such that (1.20) takes the form of a standard wave equation. There is, however, a choice of coordinates such that this equation takes the form

$$\partial_0^2 A_z + \sigma(\zeta)\, \partial_\zeta^2 A_z = 0, \tag{1.21}$$

where $\zeta = -e^{-\sqrt{2}x\Omega}/\sqrt{2}\Omega$. This looks somewhat similar to a standard wave equation, but with a position-dependent function determining the 'velocity' of the wave.

Let us approach (1.21) with a generalised plane wave ansatz of the form

$$A_z(x, t) = \xi(x) \exp\left[i\omega t \pm i S(x)\right]. \tag{1.22}$$

In the geometrical-optical limit, this solution does have the form of a standard plane wave solution, for in that limit, $S(x) \propto x$ and $\xi(x) \approx$ const. However, outside that limit, the wave vector takes the form $k_0 = \omega$; $k_i = \pm\partial_i S(x)$, where $S(x)$ need not be a linear function of x. Therefore, $k^\mu k_\mu$ is not guaranteed to vanish. The exact form, derived from (1.22), using the formula for a wave vector for a generalised plane wave (1.10), is

$$k_\mu k^\mu = \frac{1}{\xi\sqrt{-g}} \partial_x \left(\sqrt{-g}\, \partial_x \xi\right). \tag{1.23}$$

(1.20) also allows us to derive the condition

$$\partial_x \left(\sqrt{-g}\, k_x \xi^2\right) = 0. \tag{1.24}$$

Recall that, for the notion of a trajectory to be meaningful, the geometrical-optical approximation is used, in order to approximate solutions as plane waves. So, what we are looking for is a solution to (1.20), (i) for which the geometrical-optical approximation holds, and (ii) which satisfies (1.23) and (1.24). This is demonstrably impossible—there are no exact solutions to (1.20) for which $k^{\mu}k_{\mu} = 0$.

For the exact solution to (1.20), one first uses (1.24) to obtain

$$\xi(x) = \frac{\xi_0}{(-g)^{\frac{1}{4}}\, k_x^{\frac{1}{2}}}, \tag{1.25}$$

which can be plugged back into the dispersion relation (1.23) (cf. [7, pp. 3-4]), yielding the dispersion relation

$$k^{\mu}k_{\mu} = \frac{\Omega^2}{2} - \frac{k_x''}{2k_x} + \frac{3k_x'^2}{4k_x^2}. \tag{1.26}$$

Of central importance to the argument of this chapter is the fact that the expression is spacetime-dependent. The presence of first and second spatial derivatives of the spatial wave vector k_x indicates a dependence on (up to) third spatial derivatives of the function $S(x)$. The phase velocity $v_p := \frac{\omega}{k_x}$, and the group velocity $v_g := \frac{\partial\omega}{\partial k_x}$, will thus both generally deviate from c.

Consider for instance (following [7, §3]) the case of very small spacetime length scales ($\Omega x \ll 1$), such that the expression for the wave vector becomes

$$\omega \approx \left(1 + \frac{1}{2}\Omega^2 x^2\right) k_x. \tag{1.27}$$

We therefore find a phase velocity

$$v_p = \frac{\omega}{k_x} = 1 + \frac{1}{2}\Omega^2 x^2, \tag{1.28}$$

and thus a group velocity

$$v_g = \frac{\partial \omega}{\partial k_x} \approx 1 + \frac{1}{2}\Omega^2 x^2, \tag{1.29}$$

for $0 \leq \Omega^2 x^2 \ll 1$.

Both the phase and the group velocity of waves thereby exceed the speed of light. However, this does not necessarily mean that electromagnetic waves in Gödel spacetime can be used for faster-than-light signalling—it is a common misconception that the group velocity can straightforwardly be associated with the speed of information propagation (see [31] for a pedagogical clarification). Rather, it is the 'front velocity'—the velocity with which sharp pulses modulated onto the wave can propagate—which denotes signalling speed.[23]

But propagation of a signal requires that there can exist a well-defined signal in the first place. Thus, the pertinent question which arises now is the following: Does it make sense to talk about the propagation of wave front in Gödel spacetime *at all*? And thus, does it make sense to speak of a front velocity in such a spacetime? This would make sense only if the right-hand side of the dispersion relation (1.26) varied only moderately with spacetime position. Only then would we be licensed to assume that there exists a well-defined package of information to be sent from one side of a clock to another, allowing us to realise a Langevin light clock in the first place.

We now face a dilemma. Either (a) no reliable disturbance (against noise)

[23] Even the claim that the front velocity denotes the signal speed in an actual experimental setting has not gone unchallenged—see e.g. [48].

can be propagated via light; or (b) if this is possible, then this signal will be propagated at a speed varying with spacetime location (as the front velocity varies with spacetime location). Both cases are problematic for Fletcher, for in scenario (a), one cannot construct a Langevin light clock *at all*, whereas in case (b), such a clock will violate **regularity**.

1.3.4 Aren't Gödel spacetimes unphysical?

An immediate response to the above argument suggests itself: if this result has only been shown to be applicable to Gödel spacetimes, and we have good reason to believe that our universe is not described by such a solution (on various conceptual grounds related to causality [37, ch. 6] and the ability to define a consistent quantum theory [103], for example), then why should this result bother us? In more 'physical' spacetimes (Schwarzschild, FLRW, and de Sitter, for example), light rays do travel on null geodesics according to Maxwell's equations.

The most straightforward response to this objection is to note that it is plausible that the results presented here generalise to more physical spacetimes—and indeed, Asenjo and Hojman discuss the case of Kerr spacetimes in [7, §4]. Even this notwithstanding, however, the above response overlooks the epistemological crisis to which the results of this chapter give rise. In showing that, in a consistent solution to the Einstein field equations, a model of what is generally thought to be an ideal clock *does not* survey the metric, we have lost any straightforward reasons for believing that we live in the sort of universe in which we can trust such a measuring device.[24] All empirical claims about the sort of

[24] At this point, I only refer to light clocks when discussing 'such measuring devices'. I do not rule out the possibility that other dynamical systems might exist which do satisfy the clock hypothesis. However, we must countenance the existence of possible worlds which consist only of material fields which consistently violate the clock hypothesis as light clocks do in Gödel

universe in which we live are made based upon measuring devices that respond to the behaviour of matter and the metric field. This result shows that we have no way of using light clocks to determine whether the metric field we claim to have measured is, in fact, the metric field of the particular solution to the Einstein equations under consideration. In other words, it might be the case that our solution to the Einstein equations does contain, for example, closed timelike curves in the g_{ab} field, but the metric surveyed by our measuring devices (whose dynamics depends on g_{ab}, but are such that they do not survey g_{ab}) does not. The putative unphysicality of g_{ab} does not imply that the geometry associated with measuring devices governed by laws expressed with respect to g_{ab} is unphysical.

We find ourselves in a situation in which the very reason that we have for believing that we live in a particular spacetime is that we assume that clocks that we deem ideal always survey the metric field of the Einstein field equations, i.e. we accept (ii). We cannot, therefore, be guaranteed by our own theory that there exists a reliable method of inference from the behaviour of light rays to the geometry of the g_{ab} field. It is fitting that the solution which demonstrates that there might be truths about a spacetime that cannot be determined from observations confined to that spacetime bears Gödel's name![25]

spacetimes. It is to these spacetimes that the discussion in this subsection applies. In addition, two further points are in order here: (a) it is plausible that our arguments regarding light clocks generalise to a wide class of clocks typically considered to be 'good' clocks (cf. §1.4.2); (b) once the epistemological crisis delineated in this section arises for light clocks, it plausibly generalises to other clocks. Both of these matters are discussed in more detail below, and in §1.4.

[25] Whether or not one finds this troubling may boil down to one's stance on what constitutes an adequate justification of this method of inference. This is analogous to debates in epistemology over the status of rules of inference such as induction and abduction—inference rules which are themselves *rule circular*. Adopting an externalist position would entail that the observations in this chapter are not necessarily problematic, since all the externalist requires to be the case for us to be justified in believing that light surveys the metric field is that light does, in fact, survey geodesics of the metric field, whether or not we can point to some calculation or model that justifies (internally) our belief in its doing so.

1.4 The Clock Hypothesis and chronometry

Consider a clock—realised as a particular configuration of matter fields—in a particular frame of reference. Call this clock 'ideal' just in case it can be used to read off intervals of proper time along its worldline, as given by the metric field. Now ask: under what further conditions does this clock read off intervals of proper time along its worldline in *all* frames, i.e. in a frame-independent manner? As already elaborated in §1.2, a clock that satisfies this condition is one that satisfies the clock hypothesis.[26]

We have seen that a necessary condition for a clock to satisfy the clock hypothesis is that it satisfy **regularity**—*viz.*, the condition that the proper time interval measured between any two arbitrary pairs of successive ticks is the same. However, the results of Asenjo and Hojman [7] presented in §1.3 demonstrate that this principle is *not* satisfied for light clocks in Gödel spacetimes, for in such cases the velocity of signal propagation is a function of spacetime coordinates. Thus, light clocks in such spacetimes do not satisfy the clock hypothesis.

What is the philosophical upshot of this work? Such results demonstrate that what are often considered the simplest, most reliable conceivable clocks may, in certain spacetimes, not accurately measure intervals as given by the metric field—that is, they may fail to be good clocks in these spacetimes. Such an observation gives rise to broader operational concerns: if such clocks, built using light rays, need not accurately survey the metric field, should we expect that

[26] A clock which reads off intervals along its worldline in all *inertial* frames may be called 'ideal' *tout court*. Note that ideality is a much weaker condition than satisfaction of the clock hypothesis. While it would be reasonable to claim that we do not need a clock satisfying the exact clock hypothesis in order to obtain operational access to the metric field, but only a clock which satisfies the 'clock condition' (cf. footnote 4), or (weaker still) an ideal clock (in this sense), I am not convinced that all of these apparatuses are immune from the epistemological concerns raised in §1.3.4, and below.

the situation be any different for other clocks, built from different matter fields? If not, there arise pressing concerns regarding how one is to gain operational access to the metric field *tout court*.

The purpose of this section is to explore some of these philosophical concerns in more detail.[27] In §1.4.1, I begin by framing the results of this chapter in terms of Synge's distinction between 'natural observations' and 'mathematical observations'. In §1.4.2, I present a heuristic argument to the effect that one should not expect generic (Langevin) clocks to accurately record intervals along their worldlines. In §1.4.3, I consider the operational ramifications of situations in which different clocks read off different intervals along the same worldline, and in which we have no epistemic access to which of these readings, if any, correspond to the interval along this worldline as given by the metric field.

1.4.1 Natural and mathematical observations

It is helpful to view the work of this chapter through the lens of Synge's distinction between *mathematical observations* (MOs) and *natural observations* (NOs) (cf. [109, pp. 103-107]). The distinction is roughly the following: while NOs are empirical observations,[28] MOs are mathematical facts, constructs, and laws. Now, as Synge writes,

> The peculiar fascination of theoretical physics lies in the art of forcing meaningful truth out of the meaningless equation NO = MO, which is a symbolic form of the assertion that natural phenomena obey exact mathematical laws. The true inequality NO ≠ MO should not be spoken above a whisper, because it is extremely dangerous. If

[27] Ideas for the arguments in §1.4.1 are owed primarily to James Read; those in §1.4.2 primarily to Niels Linnemann.

[28] Synge distinguishes at [109, p. 103] between *uncontrolled*, *controlled*, and *imagined* NOs; this more fine-grained distinction will not be important for the purposes of this section.

> believed, it would sever mathematics from physics, and reduce both to sterility through lack of mutual fecundation. It is whispered here only as an apology to those readers who expect to see the mathematics of relativity tied to the physics of relativity by a strong chain of clear thought. It cannot be done. We have to muddle through. And if this book is dishonest in confusing MO with NO, it is no more dishonest than all similar books are, and necessarily must be. This sad state of affairs is not peculiar to relativity; every branch of mathematical physics has in its cupboard the skeleton MO ≠ NO. [109, p. 104]

The work in the current chapter brings to the fore an unexpected case of MO ≠ NO.[29] To see this, consider a particular interval along a timelike path in a solution of general relativity, and consider the proper time along this path as given by the metric field. Since this is a theoretical construct, it is an MO; call it $\mathrm{MO_g}$. Now consider the time along this path as given by a theoretical model of a Langevin clock in which the oscillating matter is described by Maxwell fields. Again, this is a theoretical construct, and therefore a MO; call it $\mathrm{MO_F}$. In rotating spacetimes, we have seen in this chapter, in light of the work of Asenjo and Hojman, that $\mathrm{MO_g} \neq \mathrm{MO_F}$.

Now, to which of $\mathrm{MO_F}$ or $\mathrm{MO_g}$ are associated the NOs of physical light rays? Naïvely, one might think, to both: to $\mathrm{MO_F}$ as light is described by Maxwell fields, and to $\mathrm{MO_g}$ given the mainstream view in relativistic physics that light rays propagate on null geodesics. But clearly, if in doubt, $\mathrm{MO_F}$—being directly concerned with the physical nature of light—should be given priority, so let us say that $\mathrm{MO_F} = \mathrm{NO}$. And since (as it turns out) $\mathrm{MO_g} \neq \mathrm{MO_F}$, we have that $\mathrm{MO_g} \neq \mathrm{NO}$. This result is more radical than standard cases of the kind MO ≠

[29] Of course, not one about which we should merely whisper!

NO, precisely in virtue of its being radically unexpected. And to put Fletcher in these terms: he has provided a mathematical model to read off MO_g, but since (as discussed) we expect that $MO_g \neq NO$, the physicality of his model is questionable.

1.4.2 Clock registry discord

In this subsection, I argue that a generic Langevin clock should not be expected to record its worldline interval as given by the metric field in generic spacetimes. This result is significant, for it calls into question a basic assumption in relativity theory, again brought out in an illuminating discussion by Synge:

> It is necessary to expose here a certain physical assumption inherent in the structure of relativity. Let C [figure suppressed] be the world-line of a material particle, and B, A two events on it, with B before A. The particle carries two standard clocks consisting of atoms of different types, or two atoms of the same type but with the use of different energy levels. Each clock registers a definite number of ticks between B and A; let these number be denoted by n_1 and n_2. The physical assumption just referred to is the following *hypothesis of consistency*: *For two standard clocks, the ratio* $n_1 : n_2$ *is a natural constant, independent of the world-line on which the observations are made and of the events on that world-line.* [109, p. 106] (Emphasis in original.)

To bring out the sense in which this chapter casts doubt upon this hypothesis of consistency, let us reconsider the circumstances under which a given clock does indeed read off its worldline interval as given by the metric field. The worldline length of a timelike path $\gamma[I]$ is given by

$$\Delta s = \int_\gamma g_{\mu\nu} dX^\mu dX^\nu. \tag{1.30}$$

Dividing up the curve into equidistant segments $\{\gamma_i\}$ with respect to an arbitrary curve parameter $\lambda \in I$ gives

$$\Delta s = \sum_i \int_{\gamma_i} g_{\mu\nu} dX^\mu dX^\nu. \tag{1.31}$$

Denote the beginning of each segment γ_i by the point p_i. Choosing small enough segments, and taking each point p_i as the origin for normal coordinates, the metric around the origin is given by (cf. [114, p. 22])

$$g_{\mu\nu} = \eta_{\mu\nu} - \frac{1}{3} R_{\mu\lambda\nu\rho} q^\lambda q^\rho + \ldots, \tag{1.32}$$

where q^ρ denotes the components of the vector to the point q considered in normal coordinates (the metric is Minkowskian at the origin). The worldline interval of the path γ thus splits into a curvature-free and a curvature-dependent part, i.e.

$$\Delta s = \int_\gamma \eta_{\mu\nu} dX^\mu dX^\nu - \sum_i \int_{\gamma_i} \frac{1}{3} R_{\mu\lambda\nu\rho} q_i^\lambda q_i^\rho dX^\mu dX^\nu + \ldots. \tag{1.33}$$

Remember now that a Langevin clock is realised through a back-and-forth signalling process in an oscillating medium. Such a physical clock can effectively read off the worldline interval just in case, for each segment γ_i, it can be seen to evolve as if it were situated in local Minkowski spacetime, while also being sensitive to curvature in just such a way as to register the higher-order terms on the right hand side of (1.33). However, different matter fields are governed by different dynamical equations, which in turn may feature different curvature couplings.[30] Thus, even for Langevin clocks, it is (on the above heuristic grounds) to be regarded as implausible that clocks built from two or more different matter fields should correctly record—or even agree upon—the full interval along a given worldline, as given by the metric field (1.33). But in that case, we have no guarantee—or even good reason to think—that the hypothesis of consistency will hold.

1.4.3 Chronometry

Consider cases in which **regularity** is lost—such as those discussed in §1.3—and in which Langevin clocks built from Maxwell fields accordingly do not correctly read off intervals as given by the metric field. Since such is the case for Maxwell fields, it is *prima facie* plausible that Langevin clocks built from other matter fields also do not correctly read off intervals as given by the metric field in such cases; moreover, in light of the reflections presented in §1.4.2, there exists no *a priori* reason to expect that such clocks will agree on the intervals read off along a particular section of a given worldline. The central question I discuss in this subsection is the following: if different clocks all read off different intervals along the same worldline, then how does one get access the 'true' geometry of the metric field—that is, how is the metric field afforded its operational *meaning*?

Focusing upon Langevin clocks for simplicity, there exist two scenarios worthy of consideration here: (A) the dispersion relation of the oscillating medium manifests constant spacetime dependence, and (B) the dispersion relation of the oscillating medium manifests variable spacetime dependence. In scenario (A), the matter fields constituting the oscillating media in the Langevin clocks under

[30] Not to mention different sensitivities to rotation—cf. §1.3.

consideration possess dispersion relations of the form $k^\mu k_\mu = \text{const}$. In this case, since this dependence is constant across spacetime, all clocks satisfy **regularity**, as discussed in §§1.2-1.3. Therefore, clock ticks remain proportional, and hence a universal notion of the time along a given worldline may be recorded. As mentioned above, this scenario is compatible with Fletcher's theorem, for all Fletcher requires is that the signal in one's clock travel at a *constant* velocity—he explicitly acknowledges that this may differ from c. In scenario (B), the matter fields constituting the oscillating media in the Langevin clocks under consideration possess dispersion relations of the form $k^\mu k_\mu \neq \text{const}$. In this latter scenario, genuine operational concerns do arise—for in this case, the variable spacetime dependence in the dispersion relations of the oscillating media means that the ticks of the clocks need no longer be proportional to one another; hence, **regularity** is lost. In this case, one cannot take the ratios of such intervals and infer universally the proper time along a given worldline; indeed, there appears to be no way, using these matter fields alone, to gain epistemic access to the intervals of proper time along this worldline as given by the metric field.

What are the consequences of the above results for our notion of spacetime more generally? From a purely terminological point of view, we call a structure 'spatiotemporal' just if, at least to a satisfactory degree, it relates to what we think is (or rather would be) measured by rods and clocks (if present).[31] Now, the metric field of general relativity is usually considered to be spatiotemporal in this sense as it is expected that rods and clocks—if present—would at least

[31] I take this to be the essence of Knox's 'spacetime functionalism'—cf. [70]; more on this in chapter 2. Of course, one may deny this view, and maintain (e.g.) that the metric field is inherently spatiotemporal. Though I do not find this view plausible, those readers who do embrace it (or another view diverging from that articulated above) are asked to consider the reflections in the body of this section in conditional form. Note also that the modal qualifications in the above mean that I am not committed in this chapter to an extreme form of material operationalism.

to a satisfactory degree measure distances and times as given by the metric field. The above findings, however, make a compelling case that Langevin clocks built from different matter fields may not agree on the interval along a given worldline. If, however, all clocks differ significantly in their operational temporal readings, and the term 'spacetime' is associated with the behaviour of rods and clocks, then it becomes questionable whether the metric field of general relativity deserves the title of 'spacetime' at all—for this field ceases to play the operational role of *codifying* the behaviour of rods and clocks.

So, whether the metric field of general relativity can be conceived of as spatiotemporal is contingent not only upon the types of matter fields at play, but also upon the nature of the solution to Einstein's field equations under consideration: with respect to a flat metric, for example, matter fields may be used to realise a clock which measures the worldline interval linked to the metric—for in this case, scenario (A) obtains. With respect to other solutions, however, scenario (B) may obtain, and there may exist no straightforward means of procuring epistemic access to intervals of proper time as given by the metric field.

The current line of thought serves as a further argument (developing upon [18, 96]) against the so-called 'geometrical approach' to spacetime theories, according to which rods and clocks in general relativity *invariably* survey the metric field $g_{\mu\nu}$ (cf. [96, §5]). Whereas e.g. [18, §9.5.2] presents various spacetime theories (such as the Jacobson-Mattingly theory [62], and Bekenstein's Tensor-Vector-Scalar theory (T*e*V*e*S) [13]) in order to argue that, locally, metric and dynamical symmetries need not coincide, and so the metric field need not have chronometric significance, the results presented in this chapter go further, for they demonstrate that, *even in the presence of such symmetry coin-*

cidence, the metric field need not necessarily have chronometric significance. This constitutes further grist to the mill of the argument that the metric field does not possess its chronometric significance necessarily (as on the geometric approach), but "earns its spurs" (to use Brown's phrase—cf. [18, p. 151]) via considerations of the dynamical behaviour of matter fields.

1.5 Conclusion

Both Maudlin and Fletcher argue for the satisfaction of the clock hypothesis in purely kinematical settings—settings in which light rays necessarily propagate along null geodesics. In this chapter I have called into question the extent to which such arguments remain sound, once it is recognised that light is a dynamical entity, which in certain spacetimes need not propagate at c. Particular trouble in this regard is found in the classes of rotating spacetimes considered by Asenjo and Hojman, in which **regularity** is lost. Consequently, light clocks cannot be regarded as ideal clocks in generic spacetime models.

These results lead to broad operational concerns: in certain spacetime models, it is not necessarily the case that we may have any operational access to the metric field. Such results also raise difficulties for the 'geometrical approach', for they provide further evidence that the metric field need not be surveyed by matter fields. In scenarios in which the metric field is *not* surveyed by rods and clocks built from matter fields, it is questionable whether this entity deserves the appellation 'spacetime' at all—thus, even in classical general relativity, spacetime may be significantly harder to come by than has hitherto been appreciated.

CHAPTER 2

Theoretical and operational spacetime

2.1 Introduction

"Spacetime is as spacetime does", or so goes the spacetime functionalist's slogan [71]. The spacetime functionalist provides a recipe for isolating the spatiotemporal component of a physical theory: first specify the functional role one associates with spacetime; second, identify the structures in the mathematics of one's theories which map onto this role. Clearly spacetime functionalism will come in many different flavours, with as many varieties as there are ways of filling out the first step of the recipe.

One of the best-known functional approaches to spacetime is due to Eleanor Knox, who states the following: [70, p. 5]

> I propose that the spacetime role is played by whatever defines a structure of local inertial frames.

Call this approach *inertial frame spacetime functionalism*. Knox motivates her brand of functionalism by appeal to Harvey Brown's view. In particular, she takes it that spacetime *just is* that structure which is surveyed by physical rods and clocks, built from matter fields. This is evident in passages such as the following (and her ensuing endorsement of the content thereof):

> Much of Brown's work is directly relevant to the question of defining a role for spacetime structure. In particular, two key themes emerge from his book [*Physical Relativity* [18]] that can serve as desiderata when seeking a concise way of expressing the spacetime role. First, Brown is concerned with the operational significance of the spacetime metric; the spacetime role had better ensure that the behaviour of rods, clocks, light rays and test particles appropriately (if not exactly) reflects the metric structure. Second, and related, he notes that ensuring such operational significance is a matter of dynamics. [70, p. 5]

The structure which "defines a structure of local inertial frames"—that I dub *theoretical spacetime*—need not coincide with that structure which is actually surveyed by physical rods and clocks built from matter fields—that I dub *operational spacetime*. Knox is committed to two separate claims: first, that the role of spacetime is to ensure that the behaviour of appropriate measuring devices reflects this structure; second, that this role is played by whatever defines a structure of local inertial frames. The main project of this chapter is to argue against the second of these claims. In addition, this chapter aims to clarify the connections between Knox's inertial frame spacetime functionalism, and one of this thesis' broader themes—the dynamical and geometrical approaches to spacetime.

The structure of this chapter is as follows. In §2.2, I present Knox's inertial frame spacetime functionalism. In §2.3, I draw the above-mentioned distinction between theoretical and operational spacetime. In §2.4, I present three cases in which theoretical spacetime comes apart from operational spacetime—while it turns out that Knox has the resources to argue that one of these three cases is unproblematic, the remaining two *do* seem to pose genuine problems for Knox's account. In §2.5, I consider means via which Knox might revise her spacetime functionalist approach, in order to overcome the gap between theoretical and operational spacetime. In §2.6, I compare Knox's spacetime functionalist approach with a different spacetime functionalist view, due to Baker [8]. In §2.7, I consider how inertial frame spacetime functionalism relates to the dynamical and geometrical approaches to spacetime theories.

2.2 Spacetime functionalism

There is a range of motivations that one might have for embracing a project that might be described as a spacetime functionalist one. On the one hand, one might be concerned with a project of giving a systematic account of the usage of the term 'spacetime'. On the other hand, one might believe antecedently in the existence of some substantival structure—spacetime—and seek a correct identification of this structure in one's theories, in order to identify some objective and ontologically autonomous entity in the world.

A third motivation for spacetime functionalism is the following: one might be interested in identifying some significant universal properties of dynamical systems that allows one to abstract away from the idiosyncrasies of particular systems, and to make generic claims about the behaviour of those systems. It is this latter motivation which I attribute to Knox—a reading made plausible by passages such as the following:

> [C]onsidering the inertial structure provides a shortcut that allows

> us to glean the empirical consequences of a theory without going into the messy details of our various measuring devices. [68, p. 347]

For Knox, 'spacetime' should be associated with this codificatory structure. The slogan of her own brand of spacetime functionalism—which in this chapter I dub *inertial frame spacetime functionalism*—is the following: "the spacetime role is played by whatever defines a structure of local inertial frames" [70, p. 9]. The idea is to functionally define 'spacetime' as any structure which itself picks out a structure of local inertial frames—the thought being that, in turn, it is this structure which can play the above-mentioned codificatory role.

Before looking at Knox's proposal in detail, it is worth discussing two distinct, antecedently understood notions that Knox takes to be co-extensional: first, the structure that defines a notion of privileged motion (i.e. inertial structure) and second the structure that is measured by rigid rods and regular clocks. That there are privileged forms of motion in spacetime is an intuition that goes at least as far back as Aristotle; in its modern guise, the privileged motion corresponds to that of *force-free* particles. While the details of what constitutes this form of motion are messy (and Knox herself gets her hands dirty with them), *that* there are such forms of motion is a respected position in the literature, even in the context of general relativity.[1]

Now, of course, if Knox's proposal is to have content, the meaning of an 'inertial frame' must be articulated. Knox gives the following (itself functional) characterisation of inertial frames:

> In Newtonian theories, and in special relativity, inertial frames have at least the following three features:

[1] Friedman [45] advocates a local notion of privileged motion in general relativity; a dissenting view can be found in Schrödinger [102, pp. 1-2].

> (1) Inertial frames are frames with respect to which force free bodies move with constant velocities.
>
> (2) The laws of physics take the same form (a particularly simple one) in all inertial frames.
>
> (3) All bodies and physical laws pick out the same equivalence class of inertial frames (universality). [68, p. 348]

Any structure which picks out a "structure of local inertial frames", i.e. a structure of local frames which satisfy these properties (initially identified as significant in the Newtonian/special relativistic context) qualifies, for Knox, as 'spacetime'.

It is illustrative to consider what Knox has in mind here in the particular, well-known case of general relativity. In order to do so, one further piece of machinery needs to be introduced: the foundational principle known as the *strong equivalence principle* (SEP). Here is how Knox puts the SEP:[2]

> To any required degree of approximation, given a sufficiently small region of spacetime, it is possible to find a reference frame with respect to whose associated coordinates the metric field takes Minkowskian form, and the connection and its derivatives do not appear in any of the fundamental field equations of matter. [68, p. 352]

In general relativity, precise satisfaction of the SEP guarantees that the metric field qualify as spatiotemporal, in Knox's sense. The reason is that, locally, the symmetries of the dynamical metric field coincide with those of the dynamical equations governing all of the matter fields; in any frame in which these dynamical equations take their simplest form, the metric field itself takes the form diag $(-1, 1, 1, 1)$. Thus, the metric field picks out a structure of local inertial

[2] Knox's version of the SEP draws on that presented by Brown at [18, p. 169].

frames for all of the matter fields, and so qualifies as spatiotemporal. This verdict seems correct, insofar as one thinks that it is the metric field g_{ab} of general relativity which codifies important aspects of the dynamics—e.g., intervals of distance and proper time as read off by physical rods and clocks built from matter fields.

At this point, however, one might ask the following question: how is it that *in general* such inertial frame structure captures—without recourse to the details—dynamical facts? It is clear that inertial frame structure does capture the common symmetries of the dynamical equations governing matter fields.[3] However, there certainly remains a conceptual gap to be bridged between such symmetries, and the full dynamics of matter, in particular the global structures of its trajectories. In this chapter, I argue that the *prima facie* success of Knox's programme is a consequence of contingent facts about the relationship between symmetries and full dynamics. When this relationship breaks down, we discover that Knox's prescription falters in its attempt to pick out an operationally significant structure.

It is worth distinguishing between two ways in which this operationally significant spacetime structure might fail to be picked out. It might be the case that inertial structure (in Knox's sense) can be identified, but there is no corresponding operational structure to be picked out, perhaps because, as a contingent fact about the initial conditions or the dynamical laws of the model, no stable matter exists. This is not a particularly serious problem, and can be sidestepped by recourse to some counterfactual claims about the relevant

[3] To the extent that (i) this common symmetry group is non-trivial and (ii) it is clear what the dynamical equations actually are. The latter is not always the case, as is evidenced by the the dialogue between Knox [69], Saunders [101] and Wallace [121] over the correct spacetime setting for Newtonian gravitation theory.

operational structure.[4]

A slightly more serious problem is that inertial structure, as Knox defines it, might fail to determine, via the purported 'shortcut', details of the behaviour of rods and clocks. §2.4, presents examples of the second kind of problem; the first kind is discussed in chapter 3.

2.3 Theoretical and operational spacetime

In this section, I argue that Knox's prescription correctly identifies a structure which I refer to as *theoretical spacetime*. But the promise of determining an operationally significant structure—i.e., one which is surveyed by physical measurement apparatuses such as rods and clocks—is achieved only via the identification of a (possibly) different structure—one which I refer to as *operational spacetime*. Knox's prescription does not invariably deliver operational spacetime.

I begin, in §2.3.1, by introducing the concept of theoretical spacetime. In §2.3.2, I discuss the distinct notion of operational spacetime. This sets up the discussion in §2.4, in which I present three cases in which theoretical spacetime comes apart from operational spacetime.

2.3.1 Theoretical spacetime

Consider two scalar fields, $\phi(x)$ and $\psi(x)$. A directional derivative V_x at a point is an element of the tangent vector space at that point. But often, there is more

[4] A similar problem afflicts the proponent of a decision-theoretic view on probability in the Everett interpretation of quantum mechanics as defended, in various forms, by Deutsch [32], Saunders [100] and Wallace [117, 118], for whom objective probabilities are given operational significance in terms of agents' betting preferences, but whose *definition* does not require the existence of these agents. A counterfactual dependence—if rational agents had existed, they would have bet in accordance with these rules—suffices.

than just linear (i.e. vector space) structure on a tangent space—on Lorentzian metric manifolds common in relativistic theories, the tangent spaces have an inner-product structure too. Consider a bimetric theory, i.e. a theory in which matter fields couple to one of two distinct metrics. V_x might be timelike with respect to the inner-product determined by one metric, but null with respect to the an inner-product determined by a different metric. So how might one determine, physically, what the appropriate metric is?

The answer—which comes in two parts—lies in the dynamical fields. First, the dynamical symmetries (defined as transformations that preserve the space of solutions) are taken to be co-extensional with the isometries of some metric (more generally, automorphisms of some geometrical structure, as in non-relativistic spacetimes; in this chapter I restrict my attention to Lorentzian spacetimes)—this fixes the transformation behaviour of the metric. Second, an invariant quantity, to be preserved by the symmetry transformations, is determined.

Recall from Einstein's first 1905 paper on special relativity [40] that the Lorentz transformations can, assuming spatial isotropy, spatial and temporal homogeneity and the Einstein-Poincaré synchrony convention, be derived from the relativity principle (RP) and the light postulate (LP).[5] The LP, together with the RP, contains two importantly distinct pieces of information: first, that the invariant quantity is a *speed*, and second, that the value of that speed is finite. The LP on its own states the independence of the speed of light from the speed of the source; it is only in conjunction with the RP that this means that the speed

[5] Einstein's statement of the relativity principle: '[t]he laws by which the states of physical systems undergo change are not affected, whether these changes of state be referred to the one or the other of two systems of co-ordinates in uniform translatory motion.' And the light postulate: 'Any ray of light moves in the "stationary" system of co-ordinates with the determined velocity c, whether the ray be emitted by a stationary or by a moving body' [40, p. 2].

of light is invariant in across all inertial frames. This invariant speed is encoded in the transformations between inertial frames that the RP determines.

In special relativity, the RP asserts that inertial frames are related by Lorentz transformations. But if the invariant speed, as determined by the RP and LP together, is treated as a variable, then there are actually infinitely many groups of transformations that relate inertial frames, each corresponding to a particular invariant speed, c. Thus, the transformation groups form a one-parameter family of (Lorentz) groups. Much of what follows turns on this fine-grained distinction, between different sorts of transformations that go under the same name. Consider a standard definition of a Lorentz transformation, as a four-dimensional transformation of the form: $x'^{\mu} = \Lambda^{\mu}{}_{\nu} x^{\nu}$, where $\Lambda^{\mu}{}_{\nu}$ is a tensor that satisfies the relation $\Lambda^{\alpha}{}_{\gamma}\Lambda^{\beta}{}_{\delta}\eta_{\alpha\beta} = \eta_{\gamma\delta}$. Here, the Minkowski metric tensor, $\eta_{\alpha\beta}$ can be expressed in coordinates in which it takes the form $diag(-1, 1, 1, 1)$. These are Cartesian coordinates in which the speed of light, $c = 1$. But there is nothing special about the choice of c; for a given situation, c could take any finite value. On setting up coordinates in which $c = 1$, the Lorentz transformation tensors would look the same. In those coordinates, a Lorentz boost in the x^1 direction gives us: $x^1 \mapsto \gamma(x^1 - vt) := \frac{x^1 - vt}{\sqrt{1-v^2}}$.

The shape of a cone in a Lorentzian tangent space depends on the choice of invariant speed. Suppose one constructs a model of vacuum general relativity, say one with globally vanishing Riemann curvature. This *still* does not uniquely determine the solution. This is because, if a model of vacuum general relativity is taken, as it standardly is, to consist of a manifold[6] with a Lorentzian metric on it, then absent a specification of the shape of the cone in the tangent space, there are uncountably many Lorentzian metrics compatible with this model—a

[6] Satisfying the standard conditions of paracompactness, Hausdorffness, smoothness, etc.

metric is Lorentzian just in case it has a signature $(-, +, +, +)$. Of course, when we speak of vacuum general relativity, we often take it for granted that the cones are such that the null direction is picked out by the speed of light (i.e. a wave in a Maxwell field), even in vacuum spacetimes. In such spacetimes, this is just a stipulation, albeit often a sensible one. The Minkowski metric is therefore not the unique globally flat Lorentzian metric; it is the unique globally flat Lorentzian metric whose null direction is the direction of propagation of electromagnetic waves in flat spacetime.

If the invariant speed is derived from the dynamics of a matter field, then the associated tangent space cone is a *matter field cone*.[7] Note that the process of discovering and articulating this invariant speed is highly non-trivial. A detailed analysis must eventually be given that imbues content to the speed, c. In this chapter, however, my concern is over a mathematical point—assuming that there is some physical, operational content to the claim that c is a dynamically relevant invariant speed, it is possible, and consistent with relativity theory, that it could have had a different value for different matter fields. A toy example of a superficially similar state of affairs—different tangent space cone structures for different matter fields—is discussed by Schuller and collaborators [52, 95, 103, 104], in the context of more general geometries like so-called *area metric* geometries.

Let us return to Knox's explicit characterisation of inertial frames introduced in §2.2. The first criterion—(1) "[i]nertial frames are frames with respect to which force free bodies move with constant velocities"—is of a different character to the other two, discussed below. It requires a specification of what is meant by 'force-free' and 'constant velocities'. This is exactly what I take to be delivered by

[7] The fact that all known matter field cones coincide can be seen as a version of what in [97, §5] is referred to as the 'first miracle of relativity'.

the second and third criteria, which effectively precisify the intuition captured by the first criterion. Recall that they state that (2) "[t]he laws of physics take the same form (a particularly simple one) in all inertial frames" and (3) "[a]ll bodies and physical laws pick out the same equivalence class of inertial frames."

In theories like special and general relativity the frames in which the laws take their simplest form are those in which the Levi-Civita connection coefficients vanish (these frames are local in general relativity and global in special relativity). If connection coefficients vanish in a frame, then, via the geodesic principle (and associated theorems—e.g. those found in [39, 49, 50]), these are frames in which force-free test bodies move in straight lines, i.e. at constant velocity. Thus, (2) and (3) pick out exactly those frames specified by (1).

Each of the last two criteria effects a distinct universal quantification. Consider two dynamical fields, $\phi(x)$ and $\psi(x)$. If we assume that there is one frame in which the laws for $\phi(x)$ take a particularly simple form, then (2) tells us that any other frame in which these laws take the same form is an inertial frame. Similarly for $\psi(x)$—but the frames in which the laws for $\psi(x)$ take a simple form need not be the same as those in which the laws for $\phi(x)$ do so. It might, however, still turn out that the transformations between the inertial frames for $\psi(x)$ and the inertial frames for $\phi(x)$ are the same (for example, are both Lorentz transformations).

The distinct universal quantification effected by (3) asserts that, for all dynamical equations in $\phi(x)$ and $\psi(x)$, the inertial frames for $\phi(x)$ and $\psi(x)$ coincide. Thus, inertial frames, for Knox, are picked out by a geometrical structure whose automorphism group is picked out by the dynamical symmetries of matter fields whose cones antecedently coincide (that the cones antecedently coincide is specified by (3)). I therefore take Knox's prescription to deliver *cone-*

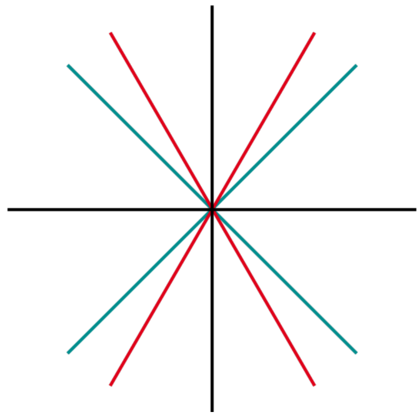

Figure 2.1: Lightcone structures associated with two different Lorentzian metrics, for which the invariant speeds do not coincide.

coincident theoretical spacetime, defined as follows:[8, 9]

> A *cone-coincident theoretical spacetime* is a geometrical structure picked out by the requirement that the intersection of the groups of local dynamical symmetries of all fields define the automorphism group of this structure, i.e. all matter field cones coincide.

Contrast this with *symmetry-coincident theoretical spacetime*:

> A *symmetry-coincident theoretical spacetime* is the geometrical structure picked out by the requirement that the groups of local dynamical symmetries of all fields define the automorphism group of this structure, up to a specification of invariant speed.

In the case of symmetry-coincident theoretical spacetime, although the symmetries of the geometrical structure under consideration coincide (modulo

[8] This picks out, as spatiotemporal symmetries, the so-called 'external dynamical symmetries'. Note the similarity to Earman's **SP1** and **SP2** [36, Ch. 2]. This is discussed in more detail in chapter 3.

[9] The 'geometrical structure referred to here includes projective as well as conformal structure—the improtant sources of disagreement in what follows are not due to disagreement on which timelike trajectories are inertial. Rather, it is over which trajectories are timelike to begin with.

a choice of parameter value, c) with the (antecedently-coincident) symmetries of the dynamical equations governing matter fields, the associated inertial frames need not coincide. Here is another way to put the issue. Consider two Lorentzian metrics η_{ab} and $\tilde{\eta}_{ab}$, related by

$$\tilde{\eta}_{ab} = \eta_{ab} - \xi_a \xi_b, \tag{2.1}$$

for a particular, suitable choice of 1-form ξ_a. The addition of the second term alters the shape of the cone structures associated with η_{ab} versus $\tilde{\eta}_{ab}$—thus, the situation is as in figure 2.1, in which (say) η_{ab} is associated with the outer (blue) cones, and $\tilde{\eta}_{ab}$ with the inner (red) cones. In turn, one can demonstrate that the isometries of η_{ab} correspond to Lorentz transformations with one particular invariant speed, and the isometries of $\tilde{\eta}_{ab}$ correspond to Lorentz transformations with a different invariant speed. It follows that the frames in which η_{ab} takes its diagonal form are different to the frames in which $\tilde{\eta}_{ab}$ takes its diagonal form—the former are related by the first class of Lorentz transformations, the latter by the second. In other words, the outer (blue) cone is one for which our coordinates are normalised, and $c = 1$; which makes the invariant speed corresponding to the inner (red) cone, $c' < 1$. The first set of matter fields will be Lorentz invariant under Lorentz transformations where the parameter is c (i.e. the γ parameter is of the form $\gamma = \frac{1}{\sqrt{1-\frac{v^2}{c^2}}}$); the second set will be invariant under Lorentz transformations where the parameter is c' (i.e. the γ parameter is of the form $\gamma = \frac{1}{\sqrt{1-\frac{v^2}{c'^2}}}$).

Now consider a situation in which the metric field ito which some fields are coupled is η_{ab}, but the metric field to which some other matter fields are coupled in their dynamical equations is $\tilde{\eta}_{ab}$. In such a case, η_{ab} would qualify as symmetry-coincident theoretical spacetime (for both it and the dynamical equations governing matter fields have Lorentz transformations with distinct invariant speed parameters as their symmetries), but not cone-coincident theoretical spacetime (for the specific Lorentz transformations which are symmetries of these objects are different—they feature different invariant speeds).

2.3.2 Operational spacetime

Anticipating that the metric field codifying the symmetries of the dynamical equations governing matter fields might nevertheless not be sufficient for that geometrical structure to be surveyed by *physical* measuring apparatuses built from matter fields, I introduce at this juncture a third species of spacetime, which I dub *operational spacetime*:

> An *operational spacetime* is a structure that is correctly surveyed by appropriate configurations of all dynamical fields.

Appropriate configurations of dynamical fields (in particular, rods and clocks) 'correctly survey' a metric just in case they (approximately) correctly read off intervals of distance, or of proper time along their worldlines, as given by that metric field.

It is clear that Knox's motivation in postulating her inertial frame functionalism is to provide a shortcut to identifying operational spacetime by identifying cone-coincident theoretical spacetime. When the metric structure is non-dynamical, and even in a large class of dynamical metric models, theoretical spacetime is indeed the same as operational spacetime, assuming stable matter configurations can be built. In the next section, however, I discuss four cases in which, even assuming that stable matter configurations exist, theoretical spacetime does not coincide with operational spacetime.

2.4 Theoretical/operational mismatches

In this section, I present three apparent problem cases for Knox's inertial frame spacetime functionalism, in which theoretical spacetime comes apart from operational spacetime. We will see that, while instructive, the first example is *not* genuinely problematic for Knox. By contrast, I take it that the subsequent two examples *are* genuine problem cases.[10]

2.4.1 T*e*V*e*S

Something analogous to the above occurs in Bekenstein's bimetric T*e*V*e*S ('Tensor-Vector-Scalar') theory, presented in [12, 13]. As discussed in [18, §9.5.2], in this theory the metric field which is surveyed by rods and clocks, the conformal structure of which is traced by light rays, and the geodesics of which correspond to the motion of free bodies, is not the 'fundamental' metric field g_{ab}, but rather a less 'fundamental' metric field $\tilde{g}_{ab}$, constructed from some of the other fields in the theory [18, p. 174]. In this case, both g_{ab} and $\tilde{g}_{ab}$ are Lorentzian metric fields; moreover, the matter fields in this theory obey (in the relevant regime) locally Poincaré invariant dynamical laws. Thus, one might think that, on Knox's inertial frame spacetime functionalism, *both* g_{ab} and $\tilde{g}_{ab}$ have local (Poincaré) symmetries coinciding with the local symmetries of the dynamical laws governing matter fields, so that *both* fields would qualify as spatiotemporal—an apparent problem.

To claim as much would be too fast—a defence of inertial frame spacetime functionalism can be mustered by drawing on the distinction between

[10] The idea of discussing T*e*V*e*S and minimally coupled massive scalar gravity as examples of the theoretical/operational spacetime mismatch, as well as many of the details, are owed to James Read.

symmetry-coincident theoretical spacetime, and cone-coincident theoretical spacetime. In T*e*V*e*S, it is the metric field $\tilde{g}_{ab}$, and not the metric field g_{ab}, which takes a diagonal form in the frames in which the dynamical equations governing matter fields take their simplest form (cf. [97, §6]). Thus, while g_{ab} qualifies as symmetry-coincident theoretical spacetime, it does not qualify as cone-coincident theoretical spacetime. If one takes "picking out a structure of local inertial frames" to require cone coincidence, then there is room for a Knoxian spacetime functionalist to make the claim that, in T*e*V*e*S, the metric field g_{ab} is *not* spatiotemporal—rather, only $\tilde{g}_{ab}$ is spatiotemporal. For Knox, this is the correct verdict, since it is $\tilde{g}_{ab}$ which is to be regarded as operationally significant, in T*e*V*e*S. Thus, once again, it is not obvious that this example should be taken to be a genuine problem case for Knox.

2.4.2 The Gödel solution

Although a valuable tool in the quest to identify spatiotemporal degrees of freedom in theories of quantum gravity, Knox's spacetime functionalism was developed initially with the goal of ascertaining the structure which plays the spacetime role in established physical theories, such as Newtonian gravitation theory and general relativity. It is, therefore, especially interesting to discover a failure of Knox's programme *within the context of general relativity itself*. In this subsection, after a brief recapitulation of the central argument from the last chapter, I demonstrate how the Gödel solution presents a problem for Knox's programme—one in which her criteria for cone-coincident theoretical spacetime are met, but in which the metric field nonetheless does not have operational significance. (So once again, in this case, the metric field qualifies as cone-coincident theoretical spacetime, but not operational spacetime.)

In order for the metric field to have operational significance, it must be possible for certain matter field agglomerations—clocks—to be able to read off intervals of proper time (as given by that metric field) along their trajectories. The degree of operational significance depends, among other things, on the degree to which these matter field agglomerations can maintain their structural integrity. A matter field agglomeration satisfies the *clock hypothesis* if it can be used to read off intervals of proper time along its worldline—*whatever that worldline may be* (see e.g. [21, 77] for some recent philosophical discussion of the clock hypothesis). Clearly this is an unphysical requirement; all physical clocks have a breaking point. But modulo such concerns, approximate satisfaction of the clock hypothesis suffices for approximate operational significance of the metric field under consideration, in terms of clock readings.

Light clocks—simple constructions consisting of two perfectly reflecting mirrors with a photon bouncing back-and-forth between them—are often regarded as being ideal clocks, i.e. are regarded as being candidates for satisfaction of the clock hypothesis. Therefore, a scenario in which such clocks can be constructed is usually taken to be one in which the metric field has operational significance. Fletcher [44] offers a mathematical precisification of this intuition, purporting to demonstrate how it is, in principle, possible to construct a light clock that ticks arbitrarily accurately and regularly in an arbitrary general relativity solution. If correct, Fletcher's argument would underwrite the success of Knox's programme in general relativity—for the metric field, which Knox identifies as spatiotemporal, could always be regarded as having operational significance, via the readings of light clocks. However, as demonstrated explicitly in the previous chapter, Fletcher's result in not universally valid. It fails in the Gödel and Kerr solutions to general relativity, for example. And this failure is not down to a failure of cone-coincidence.

Fletcher's proof represents a differential-geometric generalisation of a heuristic argument for the clock hypothesis in Minkowski spacetime due to Maudlin [77, pp. 106-114]. Maudlin's argument proceeds as follows. The path traversed by a photon (more precisely, the classical analogue of a photon; a wave-packet state of the electromagnetic field) between two mirrors is a null geodesic. This is easily demonstrated by solving Maxwell's equations. Consider a light clock at rest. The proper time elapsed along the worldline of one mirror can be calculated by using the generalisation of Pythagoras' theorem to pseudo-Riemannian flat space. Since the distance covered by a light ray is zero (as it propagates on null geodesics), the proper time elapsed between two successive bounces of the photon off the same mirror is just twice the spatial distance between the two mirrors.

A similar argument holds for a clock in a boosted frame. For an ideal clock, the description of the clock in the boosted frame, with respect boosted coordinates is the same as the original description with respect to the original coordinates; this is just the relativity principle. If the clock were to be a good clock, and one were to regard the boosted clock from the original unboosted frame, the mirrors would appear to have moved closer together. So we can ensure the clock is good by joining the two mirrors using a rod built out of matter governed by Lorentz-covariant laws. The natural contraction of the rod ensures the appropriate distance between the mirrors is maintained even after a boost. This takes care of the clock hypothesis for a boosted clock. Finally, to the extent that any boost can be modelled as being approximated by inertial motions separated by instantaneous impulses, the clock hypothesis is approximately satisfied by a light clock so constructed in Minkowski spacetime. Fletcher's result is similar in spirit; it just requires that the light clock be confined to an appropriately small region of the manifold. Relying on a particular feature of Lorentzian manifolds, Fletcher demonstrates that such a region always exists, and with it, demonstrates that Maudlin's result can be generalised.

Mathematically, all is well with the proof. However, like Maudlin, Fletcher relies on the assumption that light traverses null geodesics on its path between the mirrors. But this is not true in all solutions to Einstein's field equations. As Asenjo and Hojman demonstrate [7], in the Gödel and Kerr solutions, to the extent that one can have a persisting wave packet that represents light, this packet does not traverse null geodesics; indeed it does not even traverse a path of constant velocity, but instead its velocity manifests spacetime position dependence. This means that light clocks, if constructable, do not accurately read off intervals of the general relativity metric. The metric, therefore, cannot be afforded operational significance via light clocks. Since light clocks are one of the simplest kinds of clock which one could imagine, in such spacetimes, the metric field does not, it appears, qualify as operational spacetime—for if even light clocks do not read off intervals of proper time along their worldlines, there is no *a priori* reason to expect this to be true of more complicated clocks either.

To repeat: it is interesting to note that locally, the cone structure of the metric field g_{ab} and that of the dynamical equations governing the Maxwell fields do coincide. After all, the metric g_{ab} is the generalisation of the Minkowski metric, η_{ab}, which was derived from the *light* postulate—in other words from Maxwell's equations themselves. Thus, we have an unequivocal case of cone-coincident theoretical spacetime, but not operational spacetime.

2.5 Proposed revisions

We have seen that one putative problem case for inertial frame spacetime functionalism—presented in §§2.4.1—is easily dispensed with, because it fails to appreciate the distinction between symmetry-coincident and cone-coincident theoretical spacetime. If one takes it that it is the latter notion of theoretical spacetime in which Knox is interested when she proposes her functional characterisation of spacetime in terms of picking out a structure of local inertial frames, then Knox's approach *does* deliver the intuitively correct verdict on what counts as spatiotemporal in that case. On the other hand, other apparent problem cases for inertial frame spacetime functionalism—in particular, those presented in §§**??**-2.4.2—are effective, for in these cases cone-coincident theoretical spacetime does not qualify as operational spacetime—which, for Knox, is the intuitively correct notion of spatiotemporality.

These problem cases lead one to question how Knox's spacetime functionalist proposal could be modified in order to deliver the intuitively correct verdict on what counts as spatiotemporal in *all* cases. The natural revision to propose is the following:

> Spacetime is that structure which has chronometric significance.

Making this move would mean that the correct verdict on spatiotemporality would be delivered in all the cases presented in §2.4. However, such a move would have the profound *disadvantage* that it would deprive Knox's characterisation of spacetime of *easy applicability to new cases*. For example, consider the case of Gödel/Kerr spacetime presented in §2.4.2. When one considers Maxwell fields in such a spacetime, it is easy to note that the metric field of general relativity qualifies as cone-coincident theoretical spacetime—and so spacetime for Knox *tout court*. However, it is very difficult—and involves substantial, non-trivial calculations—to ascertain whether this structure does or does not qualify as operational spacetime. Thus, to define spacetime in terms of operational spacetime would, arguably, diminish the ready applicability of Knox's programme.[11] Indeed, Knox says as much herself, when she writes:

> We would ideally like a formulation that entails appropriate phenomenological behaviour without requiring us to model the behaviour of complex systems. [70, p. 5]

In light of this, one might instead wonder whether there is any way in which Knox's original proposal could be defended. Ideally, what Knox would need to argue here is that (cone-coincident) theoretical spacetime is in general a good—albeit defeasible, in light of the cases presented in §2.4—guide to operational spacetime. If such a link can be rendered explicit, then there is room for Knox to retain her original functional definition of spacetime, in terms of picking out a structure of local inertial frames. What the examples presented in §2.4 of this chapter demonstrate, however, is that such a link is not inevitable—and therefore, that there is a burden on Knox to spell out the connection between theoretical and operational spacetime, if her account is ultimately to be compelling.

2.6 Baker's proposal

When it comes to spacetime functionalism, Knox's approach is not the only game in town. One distinct but noteworthy spacetime functionalist approach is due to Baker [8]—not sharing Knox's above-mentioned operationalist leanings,

[11] It is also interesting to note that making this move would compel one to identify spacetime at the level of *solutions*, rather than *theories*.

Baker advances a very different functional conception of spacetime. He begins [8, pp. 11-12] by noting that many *different* factors seem to contribute to the spacetime concept:[12]

> I won't make any attempt to give an exhaustive list of candidates here, but the following are examples of criteria which are logically independent of Knox's inertial criteria and which seem to also count toward a structure's satisfying our spacetime concept:
>
> - The structure is non-dynamical, at least with respect to non-gravitational interactions.
> - The structure is (in some sense) located everywhere in all states of the theory.
> - The structure does not carry energy or momentum.
> - "Vacuum" solutions exist which describe the (putatively) spatiotemporal structure in the absence of other structures.
> - There are no other structures in the theory which can exist without the (putative) spacetime structure.
> - The structure grounds or explains a family of modal facts about which states are geometrically possible, where geometric possibility does not reduce to physical possibility [16, pp. 50-51].
> - It is a (higher-order) law of nature that the geometric symmetries of the structure are dynamical symmetries of the theory [63, 105].

[12] Citations in the following quotation have been modified for consistency with the present chapter; there is no change in content.

> - Forces propagate across the spatial distances defined by the metric characterizing the structure (so that long-range forces like electromagnetism fall off proportionately to the inverse square of this distance, and so on).
> - The structure is fundamental.
>
> Again, this is not meant to be an exhaustive list. Rather it is meant to illustrate that a vast number of different criteria could plausibly figure into our ascription of the name 'spacetime' to a given theoretical structure, depending on the details of the laws that define that structure. And indeed, Knox's own criterion,
>
> - The structure determines the difference between inertial and non-inertial frames of reference,
>
> belongs high on this list, perhaps even at the top. She has certainly shown that it's a very important criterion. My only disagreement is with her claim that it is the sole criterion.

On the basis of these apparently myriad factors which appear to contribute to the spacetime concept, Baker draws the natural conclusion that, ultimately, there is no unequivocal notion of spatiotemporality; rather, spacetime is a *cluster concept*: [8, p. 2]

> [O]ur spacetime concept has the structure of a cluster concept. Rather than possessing a single set of necessary and sufficient conditions, cluster concepts can be satisfied in a variety of different ways by different entities falling under them.

What to make of this proposal, especially as compared with Knox's own? On the one hand, Baker is likely *correct* that our pre-theoretic concept of spacetime (insofar as we have such a pre-theoretic concept!) cannot be analysed via one unequivocal set of necessary and sufficient conditions. (This, of course, is part of a broader lesson against conceptual analysis familiar from the latter half of the 20th century in all branches of analytic philosophy). On the other hand, in light of its heterogeneity, Baker's analysis lacks *practical applicability*. For example, consider the cases presented in §2.4—on Baker's cluster concept approach to spacetime, which structure in each of these theories is to be regarded as spatiotemporal? In light of the complexity of the analysis, it is difficult to give any definitive answer to this question. Thus, while Baker is *morally* right on the nature of spacetime, his analysis has limited practical value.

Such is not the case for Knox's proposal: Knox gives a simple, functional characterisation of spatiotemporality, which is readily applied to new spacetime theories (consider, for example, the novel work to which Knox puts her programme in [67, 69]). It might be the case that Knox's criterion does not *fully* capture our notion of spatiotemporality (including Knox's own—this is, of course, the central point of §2.5 above); nevertheless, the claim is that this account of spacetime can deliver intuitively correct verdicts on spatiotemporality, and so should feature as a (defeasible!) guide to spatiotemporality in new cases also. While, as argued above, Knox should be more explicit that inertial frame structure need not always constitute the *sine qua non* of spatiotemporality (whether (i) because one has Knox's operationalist leanings—in which case there is a gap between theoretical and operational spacetime, or (ii) because one embraces Baker's notion of spacetime as a cluster concept), her approach has the capacity to be put to novel interpretative work, in a manner in which

Baker's approach does not.

Thus, the final verdict on Knox's inertial frame spacetime functionalism versus Baker's cluster concept spacetime functionalism is the following: while Baker's analysis is likely closer to our overall conception of spatiotemporality, Knox's analysis has the virtue of readily applicability to new cases. Insofar as one takes inertial frame structure to be a guide to the other qualities which feature in the spacetime concept (perhaps those on Baker's list), one may continue to be justified in following Knox's approach. Of course, however, one should ideally make explicit the link between inertial frame structure, and those other factors featuring in the spacetime concept.

2.7 Dynamical and geometrical approaches

It is sometimes claimed (see e.g. [71, 80]) that Knox's inertial frame spacetime functionalism "extends" previous work on the dynamical approach. To briefly recapitulate, the dynamical approach, for the purposes of this chapter, can be taken to involve two claims:

(1) Fixed fields, such as the Minkowski metric field η_{ab} of special relativity, are to be ontologically reduced to the symmetries of the dynamical equations governing matter fields (the dynamical view is, therefore, a modern form of relationalism—cf. [92, §6.3.2]).

(2) Ontologically autonomous metric fields, such as g_{ab} in general relativity, do not have their chronometric significance—i.e., are not surveyed by physical measurement apparatuses—of necessity (i.e., in all solutions of any theory in which they appear).

Focussing on (2), advocates of one version of the opposing *geometrical approach* to spacetime theories would state that ontologically autonomous metric fields, such as g_{ab} in general relativity, *do* have their chronometric significance necessarily. However, in [96, 97], it was argued that this particular version of the geometrical approach is not viable—precisely because there exist problem cases for such a view, in which one has a metric field g_{ab} in one's theory, but that structure is not surveyed by physical rods and clocks (some of the examples presented in §2.4 of this chapter would count as problem cases of this kind for the geometrical view).

In any case, regardless of the particular view which one espouses in the dynamical/geometrical debate, one can ask: is it true that Knox's spacetime functionalism is an "extension" of the dynamical view? In my view, this claim is not correct; commitments in the dynamical–geometrical debate are *orthogonal* to whether one endorses Knox's spacetime functionalism. The reasons are the following: whether one thinks that a given metric field is ontologically reducible to dynamical symmetries, or does or does not have its chronometric significance necessarily, is distinct from the question of whether one should regard this object as being spatiotemporal, on Knox's functional analysis of spacetime. Suppose, for example, that one endorses the dynamical approach to spacetime theories—then on Knox's programme one will, in light of (2) above, deny that e.g. a Lorentzian metric field g_{ab} always qualifies as spatiotemporal—for whether this is so will depend upon particular facts about the matter sector of the theory under consideration. On the other hand, if one denies (2) (*à la* the strong version of the geometrical approach above), then one will think that a generic Lorentizan metric field g_{ab} always qualifies as spatiotemporal. Not only are both of these dynamical and geometrical views perfectly compatible with

Knox's spacetime functionalism, but, moreover, they would also be compatible with a different functional conception of spacetime—or, indeed, with certain *non*-functional approaches to spacetime.

One further remarks is in order. One way to understand the the contrast between the geometrical and the dynamical approach is that the former is willing to make certain 'riskier' assumptions about the chronometric status of a given field (e.g. g_{ab}) than the dynamical approach is willing to countenance. We have seen in §2.5 above, however, that an underlying assumption of Knox's approach is that theoretical spacetime is always a good guide to operational spacetime. In this sense, Knox too is (arguably) making an *a priori* assumption about the nature of certain fields which appear in our physical theories. In this very particular sense, one might argue that such assumptions place Knox closer to the geometrical rather than the dynamical view. This should be surprising, since, as noted above, several authors tie Knox's spacetime functionalism more closely to the *dynamical* approach, than to the geometrical approach.

2.8 Conclusion

The central aim of the present chapter has been to identify and diagnose problems for inertial frame spacetime functionalism. The diagnosis made it clear that more needs to be done than Knox might have initially anticipated in identifying the chronometric structure of dynamical fields—Knox's epistemic shortcut from universal symmetries to generic field behaviour is not universally valid.

Baker has a more permissive account than Knox when it comes to his particular spacetime functionalist approach; thus from the point of view of practical utility and applicability, Knox's approach is to be preferred. If one shares Knox's point of view regarding operational access to spacetime geometry, then one faces the urgent task of bridging the gap between theoretical and operational spacetime. To be clear: the intended lesson of this chapter is not that Knox's preferred brand of functionalism is a failure, but rather that it faces a cluster of challenges which have, thus far, gone unremarked upon.

CHAPTER 3

Spacetime in supersymmetric field theories

3.1 Introduction

Supersymmetry (SUSY) is a proposed symmetry between quantum fields of integer spin (bosons) and quantum fields of half-integer spin (fermions). SUSY is a bizarre and brilliant idea, and it deserves to be scrutinised by philosophers. In this chapter, we will merely dip our toes into its deep waters. The motivation for the project of looking at SUSY from a philosophical perspective is not as radical as it might appear at first. Philosophical intuitions are heavily influenced by, and tied to, particular frameworks of physical theories. Exploring the logical space of such possibilities has value for the light it sheds on the constraints imposed by our current theorising. As Friedman argues [46], physical theories are embedded in a package of background formal and conceptual assumptions, which can be hard to see explicitly, or to appear matters of necessity when they are seen.

Friedman's contention is that knowledge is stratified into three layers. The first consists of concepts and principles of 'empirical natural science properly so-called', by which Friedman has in mind the actual physical laws like those of general relativity or quantum mechanics which make directly testable predictions. The second is the set of constitutive principles required in order to be able to articulate these laws in the first place. These include logical and geometrical structures. The third level is a ‘meta-framework’ that motivates a choice of second-level structures. None of these levels is unrevisable, and the choice of superspace in this thesis reflects a decision to assess the utility of the spacetime concept in light of a revision to some of this second-level structure. SUSY places current spacetime and particle physics in a broader logical landscape, revealing hidden assumptions and contingencies. In this regard, its epistemic value is independent of whether or not SUSY is realised in nature and this is especially clear when seen in light of Friedman’s own interest in ‘a genuine expansion of our space of intellectual possibilities’.

Further pragmatic motivation for studying SUSY comes from the role that it plays in string theory, one of our prime candidates for a quantum theory of gravity. Any theory which purports to describe our world needs to have the conceptual resources to describe fermions. The only way, that we know of, to incorporate fermions into a string theory is through SUSY (for more on this, see e.g. [11, Ch. 4-5], [128, Ch. 14]). So any thesis regarding spacetime in a string theory that hopes to model our actual world must include reference to the role that SUSY plays. In this chapter, I choose to focus only on SUSY, divorcing it from the context of string theory, but it is useful to bear this motivation in mind throughout.

In this chapter, I attempt to identify the appropriate spacetime setting for a supersymmetric field theory. In chapter 2, I discussed a few motivations for engaging in the project of identifying structures might might deserve the name ‘spacetime’in. Whatever the motivations, I hope to make clear in this chapter that SUSY introduces a new dimension (several, actually) to the discussion, and places under the microscope our motivations for being interested in spacetime structure in the first place.

Constructing an appropriate spacetime for a supersymmetric theory is tantamount to taking a stance on the viability of a spatiotemporal interpretation of *superspace*[1]—an extension of four-dimensional Minkowski spacetime to include (at least) four new dimensions, coordinatised by mathematical objects known as *supernumbers*. These objects are, in one significant way, quite different from real or complex numbers—some of them have the property that their order of multiplication makes a difference; in mathematical terms, they are said to have *non-trivial commutation properties*.

This chapter argues for two theses: first, that some standard arguments that motivate a particular choice of spacetime structure in familiar spacetime theories also motivate the choice of superspace as the appropriate spacetime for SUSY field theories. And second, that the broader metaphysical utility of the concept of spacetime requires more than just the satisfaction of this universality condition. In short, superspace is spatiotemporal, but in an attenuated sense; in the terminology of chapter 2, superspace is a theoretical, but not operational, spacetime.

This chapter is structured as follows. In §3.2, I examine how dynamical symmetry considerations factor into arguments concerning spacetime structure. In §3.3, I briefly introduce SUSY by analogy with constructions familiar from standard spacetime theories. In §3.4, I discuss problems and considerations analogous to those brought up in §3.2, this time in the context of SUSY, and

[1] This is not to be confused with the term 'superspace', (or any associated terms like 'mini-superspace') as used in the context of geometrodynamics in e.g. [60, 79]. They use the term to refer to a mathematical space of spatial 3-metrics.

show that all of the arguments presented in that section mandate a choice of superspace, rather than Minkowski spacetime, as the appropriate spacetime setting for SUSY. Finally, in that section, I discuss, in more detail, motivations for being interested in spacetime structure, and whether the identification of superspace as the appropriate spacetime setting for a certain class of SUSY field theories is theoretically valuable.

3.2 Arriving at spacetime structure

Our only means of epistemic access to spacetime structure (whatever its associated metaphysics) is through the behaviour of dynamical fields.[2] As a contingent fact about a particular dynamical fields, it might be the case that certain configurations exhibit a behaviour that allows for the (possibly approximate) measurement of proper time according to some metric along intervals of a worldline. As a further contingent fact, it might be the case that a number of different matter fields can be brought into configurations which measure the same intervals of proper time along their worldlines. If it is the case that *all* of the matter fields to which we have access are such that intervals of proper time along their worldlines are seen to correspond to the same metric, then, on purely practical grounds, it makes sense to see this structure, as codifying some universal features of the dynamics of these fields. In particular, granting this 'universality' supposition, we can then make claims about certain types of behaviour (e.g. the measurement of length intervals, claims about signalling and causality) in a manner which is neutral on further idiosyncratic details of

[2] I use the term 'dynamical fields' to include test bodies, i.e. matter fields that do not themselves act as sources in the Einstein field equations. Thus deducing metric structure from, for example, conformal and projective structure à la Ehlers, Pirani and Schild [39] from the behaviour of test particles (and light rays) still qualifies as using dynamical fields.

the dynamics; this is what is sometimes taken to be characteristic of *kinematical* structure in dynamical theories. The methodological importance of spacetime symmetries is precisely in their ability to encode some of these kinematical facts. This much is uncontroversial. Disputes arise over further claims about what explains this feature of dynamical fields; I return to the *dynamical–geometrical* debate in the next two chapters.

It is important to be clear about the use to which the spacetime concept is to be put. What I have described above is a fairly minimal requirement, but one might be interested in even less—instead of encoding facts about the behaviour of certain contingent matter agglomerations (i.e. surveyors), we might think of spacetime structure as encoding nothing more than certain geometrical facts common to all dynamical fields. On this view, all there is to being a spacetime symmetry is to be a transformation such that when applied to *any* dynamical field, solutions are mapped to solutions. In chapter 2, this was referred to as *theoretical spacetime*, to contrast it with *operational spacetime* which deals with the behaviour of matter configurations like rods and clocks; call such matter configurations *surveyors*.

Theoretical spacetime: That geometrical structure whose automorphisms are the (common) symmetries of the dynamical fields.

Operational spacetime: That structure which is correctly surveyed by physical surveyors built from matter fields.

One might require even more than what I described above—that the appellation 'spacetime' refer to some structure that reproduces our first-personal experience of a world through which we 'move' in some sense; call this *phenomenal spacetime*. In this chapter, my interest is in the relationship between theoretical and operational spacetime; phenomenal spacetime will not be discussed. We will see that the theoretical spacetime concept is useful for metaphysical theorising in virtue not only of the universality condition, but also other contingent facts about certain types of dynamical fields. Equipped with this distinction, let us now review some standard proposals concerning spacetime structure and determine whether they refer to theoretical or operational spacetime.

3.2.1 Earman's principle

Earman presents his famous prescription—make your dynamical and spacetime symmetries (as defined below) coincide—henceforth referred to as *Earman's Principle*, in the second chapter of *World Enough and Spacetime* [36]. In order for this dictum to make sense, we therefore need (at least some) pre-theoretic idea about what spacetime and dynamical symmetries are. Note that, insofar as reference is made *only* to symmetries and no stipulation is made of the invariant quantity to be preserved under these symmetry transformations, this is a prescription for arriving at symmetry-coincident theoretical spacetime.

In what follows, I briefly discuss two proposals for determining which symmetries count as spacetime symmetries. In §3.2.1, I describe Earman's own proposal and its drawbacks. I propose my own criteria for identifying spacetime symmetries in §3.2.1.

Earman on absolute objects

On a standard way of understanding the structural set up of a spacetime theory, a theory is identified with (or, more loosely, associated with) a set of models.[3] Models can be thought of as set-theoretic entities whose mathematical struc-

[3] This view of scientific theories is known as the *semantic view*, and was introduced by Suppes [108].

ture is representative of the structure of the world. A model $\mathscr{M}$ of a spacetime theory is an ordered tuple of the form $\langle M, A^i, P^i \rangle$ where M is the set of independent variables of our theory (for a classical field theory, this is taken to be a 4-dimensional smooth manifold).

The A^is represent the *absolute (geometric) objects* which characterise the spacetime. An absolute object can be thought of as a geometric object which is the same across all models. Consider, for example, *Newtonian spacetime*, which privileges a certain inertial frame corresponding to a standard of absolute space. Such a frame would be identified using a privileged timelike vector field in this formalism.

The P^is are geometric objects which represent the dynamical elements of the theory—the matter fields and force fields which are subject to dynamical equations. For a classical field theory, these will be maps from M to some appropriate mathematical space in which the fields take their values. The models thus specified form the set of *kinematically possible models* (KPM). The dynamics is a matter of stipulating which of these maps are nomologically allowed. A theory is then simply the collection of models whose fields obey these dynamical constraints, *the dynamically possible models* (DPM).[4]

On this setup, it is easy to distinguish between dynamical and spacetime symmetries. The former are the transformations to the P^is under which the dynamical equations retain their form. These transformations are actually carried out by pushing forward the P^is along (a subset of)[5] diffeomorphisms on

[4] This construction glosses over an important technical point about whether the absolute objects remain unchanged across all kinematically possible models (KPMs) or merely across all dynamically possible models (DPMs). I take Anderson's [5] terminology of absolute objects to refer to the latter, while Pooley's [93] 'fixed fields' refer to the former. For the purpose of this chapter,the Andersonian conception of absolute objects suffices.

[5] There is a subtlety here which is discussed at length by Pooley [94, p. 128, p. 134 fn. 41]. When dealing with coordinate-free formulations, *everything* is, by construction, covariant under diffeomorphisms (a property known as general covariance, and distinct from diffeomorphism

the base manifold. The condition that the dynamical equations retain their form is the requirement that these diffeomorphisms map dynamically possible models onto dynamically possible models. A diffeomorphism which, when the A^is are pushed forward along it, leaves the A^is invariant, is known as a *spacetime symmetry*. In other words, spacetime symmetries leave invariant the absolute objects of a theory. One might then consider the base manifold together with its absolute objects to constitute the *background structure* of the theory. The spacetime symmetries are then the automorphisms of this structure. The method for picking out absolute objects comes from some independent considerations (usually extra-theoretic, metaphysical or methodological). At least, this is the case on the coordinate-independent view presented here. Friedman himself follows this route in arriving at the symmetry group of a theory [45, p. 42]:

> the symmetry group...of a theory is the largest subgroup of...the group of automorphisms...leaving certain of the geometrical objects of the theory—...the absolute objects of the theory—invariant.

The problem with this setup is that it assumes that we have some independent handle on, at the very least, the sorts of things that could qualify as absolute objects (or fixed fields). This prescription then allows us to narrow down the options for this structure. Earman's book is primarily concerned with the debate between *substantivalists* for whom spacetime is a matter-independent substance in which matter distributions are embedded, and *relationalists* for

invariance as characterised by Pooley in [93]). All this means is that if $\mathscr{M}$ is a model of a theory, then so is the tuple obtained by pushing forward *all* the objects defined on the manifold. But these two models need not be identical. If they are not, then the subgroup of diffeomorphisms which leaves each of them invariant will be different for each model. But they will be isomorphic as groups. If the models are of the form $\mathscr{M} = \langle M, A^i, P^i \rangle$, then the abstract group which is constrained to act only on the A^i to leave the models unchanged will be the spacetime symmetry group, while the group which acts on the P^i to map it onto another model will be the dynamical symmetry group.

whom spacetime is constituted by a set of relations amongst matter distributions, which are fundamental.

The examples presented in Earman's own treatment constitute a highly restricted class of spacetime theories. In particular, the fields and particles are assumed to have no degrees of freedom other than positions and momenta. Incorporating, say,$U(1)$–symmetric dynamical gauge fields would, according to Earman's prescription, require the spacetime symmetries to include that group. An obvious response to this worry is to somehow naturally restrict the class of dynamical symmetries to which Earman's principle is to apply to the so-called *external* dynamical symmetries.

Spacetime symmetries from external symmetries

Earman takes both parties to have agreed on the sort of thing that constitutes a spatiotemporal relation; their disagreement is over what grounds these relations. Earman himself does not offer any further clarification on this matter: spatiotemporal relations are things like relative distances, angles and time intervals. Implicit in this, therefore, is some universality constraint which would explain our interest in the kinematical structure of spacetime. In this subsection, I make that intuition precise.

A state of a system is specified by some collection of values of variables. Some of these variables are *independent*, which means we are free to choose their values, other *dependent*, which means that they are determined by certain functions on the independent variables which we had chosen. This distinction is one of convenience, and it is most common for the set of independent variables in a field theory to be identified with spacetime points. Since this is the very structure we are trying to get an independent handle on, we cannot presuppose it.

Consider the state of a complex-valued classical Lorentz-invariant field. It is specified by an uncountable number of tuples of complex numbers. We are interested in dividing this parameter space into 'internal' and 'external' parameter spaces. This is where our antecedent knowledge of the symmetry structure of the theory comes in. Given the dynamical symmetries of this theory (say, $U(1)$ and $SO(1,3)$), the parameter space splits naturally into a two-dimensional parameter space invariant under $U(1)$ rotations and a four-dimensional parameter space invariant under $SO(1,3)$ rotations. This split is the basis of the internal/external distinction. If we now consider a number of other Lorentz-invariant fields, their parameter spaces will also split. Further, they will all have the four-dimensional $SO(1,3)$-invariant parameter space in common. The symmetry group of this common parameter space is what I will refer to as *external dynamical symmetries*. Note that, in the case of a theory with just one field, although the natural split still exists, there is no principled reason to refer to one parameter space rather than the other as the 'external' space.

We can use Earman's prescription to structure spacetime appropriately. For example, in Newtonian theory, force-free particles execute straight-line trajectories in spacetime. The external dynamical symmetry group is the Galilean group (the internal symmetry group is the singleton $\{e\}$). Earman's prescription tells us that the spacetime symmetry group must, therefore, also be the Galilean group. In other words, the subset of the diffeomorphism group which induces push-forwards on the absolute objects of the theory is the Galilean group.

It is important to bear in mind that our identification of external dynamical symmetries with spacetime symmetries is *extensional*. We assume that absolute objects exist, and that they have some associated symmetry group. We then use

Earman's principle as a guide for how to use dynamical symmetries to ascertain this structure. The dynamical and spacetime symmetry transformations themselves act on distinct objects, and it is a methodological convention that they are co-extensional. And, using the terminology introduced in chapter 2, the resulting spacetime is a *theoretical* spacetime.

We can think of Earman's prescription as mandating that we use the most minimal structure that allows us to encode universal dynamical facts. As such, it can be seen as an Occamist norm on our theorising. But the question that now arises is why we ought to restrict our attention to external dynamical symmetries. In the following sections, I canvass three interpretations of spacetime structure that account for this restriction.

3.2.2 The geometrical approach

There appears to be a simple answer to the question posed above—we restrict to external symmetries just because they *are* the symmetries of the underlying spacetime geometry. In other words, we assume that spacetime, whatever it is, can be represented by a smooth manifold with some geometrical structure on it. This geometrical structure is privileged in the sense that it is pervasive—all dynamical fields adhere to its directives. This is the central assumption of the so-called *geometrical approach* to spacetime as advocated by e.g. Friedman [45], Belot [15] and Maudlin [77], detailed in the introduction to this thesis. The mechanism by which this is achieved is disputed;[6] indeed whether such a mechanism (in the quotidian physical sense of it being some dynamical process) is even required is up for debate. But, whatever the reason, on this view it is the

[6] Misner, Thorne and Wheeler, for example, rely on the geometrical-optical limit of Maxwell's equations to determine that light rays move on null geodesics of the metric of general relativity. Fletcher uses this assumption to argue that arbitrarily accurate light clocks can, in theory, be constructed. But, as we discovered in chapter 1, both of these claims can be disputed.

case that matter fields are compelled, by the geometry, to evolve in such a way as to survey it. Ultimately, absent a detailed dynamical account for why this is the case, this view is justified by appeal to some form of abductive reasoning.

On this approach, we are in a good epistemic situation. All we need to do is pick a suitably stable matter configuration, and use it to measure distances, angles and time intervals. If we corroborate these readings with other suitably chosen matter configurations of different materials, this eventually gives us a sufficiently accurate picture of the underlying geometry of spacetime, using Earman's prescription. External dynamical symmetries are thus, by definition, the symmetries of matter fields which match the symmetries of the underlying geometry (i.e. the spacetime symmetries) because they preserve the 'pervasive' geometrical structure. This is, of course, jut a necessary, not sufficient condition for matter to 'survey' this geometry, i.e. for configurations to read off intervals of the time parameter associated with worldlines embedded in that geometry; this is a specific case of the observation presented in chapter 2 that theoretical spacetime structure does not, in general, guarantee operational spacetime structure. However, in a large class of physically relevant spacetimes (e.g. Minkowski, Schwarzschild, FLRW) the standard matter that we use to construct clocks and rods is such that theoretical spacetime structure *does*, in fact, determine the metric structure that they read off.

Thus the geometrical approach and Earman's principle together form a package; indeed the latter requires the former in order to be meaningfully formulated. In §3.2.3, I drop the assumption that some independent element of reality (like a geometry) compels matter to behave in a specific way.

3.2.3 The dynamical approach

The dynamical approach, answers the above question of why we restrict ourselves to external symmetries in a different way. Rather than presupposing that an independent element of the world, represented by a geometry, is the ultimate explanans of the spatiotemporal (i.e. pervasive or universal) behaviour of matter fields, the proponent of the dynamical approach asserts that the dynamics itself is the ultimate explanans. Of course, this does lead to an unanswered question about why there is such a large symmetry coincidence across all known matter field theories—on the dynamical approach, this is just an unexplained brute fact about the world.[7]

For a Brown-style proponent of the dynamical approach, a spacetime symmetry is connected to the behaviour of chronometric matter fields, i.e. matter fields which can read off intervals of proper time as given by g_{ab}; therefore as a prescription for identifying operational spacetime, the dynamical approach only works with solutions to the Einstein field equations in which theoretical spacetime structure guarantees operational spacetime structure (as in the aforementioned examples of Minkowski, Schwarzschild and FLRW solutions). The dynamical approach is, however, perfectly consistent with theories in which this link does not exist; it still allows us to identify symmetry-coincident theoretical spacetime in all theories. In §3.2.5, we will see that this link is provided by the so-called *clock hypothesis*.

In summary, on the dynamical approach, (theoretical) spacetime structure is wholly determined by the behaviour of matter fields. Thus, a version of Earman's principle holds trivially—all external dynamical symmetries will be spacetime symmetries, because the latter can only be defined as external dynamical sym-

[7] For a deeper discussion of this fact, see [18, 97].

metries.

3.2.4 Spacetime functionalism

The dynamical approach technically does not commit its adherent to a particular view on *spacetime*; Brown himself is only interested in physical geometry; whether the structure thus identified deserves the appellation 'spacetime' is additional. Knox's spacetime functionalism (as presented in [68, 70] and chapter 2) is one way of executing this additional step (note, however, that spacetime functionalism is compatible with the geometrical approach as well). For Knox, spacetime is any structure in a physical theory that plays a role in characterising inertial structure. As we saw in chapter 2, she defines inertial frames as:

> (1) Inertial frames are frames [footnote suppressed] with respect to which force free bodies move with constant velocities.
>
> (2) The laws of physics take the same form (a particularly simple one) in all inertial frames.
>
> (3) All bodies and physical laws pick out the same equivalence class of inertial frames (universality). [68, p. 348]

Knox builds into her notion of inertial structure, through the second criterion, the transformation behaviour of matter field equations that the proponent of the geometrical view builds into its characterisation of the metric. At this point, then, it is clear that spacetime functionalism guarantees us symmetry-coincident theoretical spacetime. Of course, as was discussed at length in chapter 2, the third criterion is what ensures that spacetime functionalism determines *cone-coincident* theoretical spacetime. We also know, from Knox's own writings, that her inertial frame spacetime functionalism is motivated by a desire to capture operational spacetime as well. Consider, for example, the following passage:

> By defining a structure of local inertial frames in the way described by the strong equivalence principle, the metric succeeds in filling the desiderata set for spacetime [...]: the local coupling ensures that the local symmetries of the dynamics coincide with the local symmetries of the metric, and hence ensure that the metric governs the behaviour of rods and clocks which obey those dynamical laws. [70, p. 5]

This makes sense given the intended domain of Knox's programme—classical spacetime theories like Newtonian Gravitation, and special and (some solutions of) general relativity. In such scenarios—scenarios in which, as will be discussed in §3.2.5 the clock hypothesis is satisfied—cone-coincident theoretical spacetime structure is enough to determine operational spacetime structure.

3.2.5 How to do things with spacetime

It seems intuitively obvious why one might be interested in operational spacetime; an account of the reasons that measuring devices behave in the way that they do tells us a great deal about the dynamical structure of the matter fields in the world—after all, measuring devices form part of the background structure that allow for the testing of hypotheses. Another reason (which I will not address further in this chapter) is to overcome the problem of *empirical incoherence*, a term first used in the context of quantum mechanics by Barrett [10], of quantum gravity theories. Briefly, the worry is that many (perhaps all) of our best guesses as theories of quantum gravity seem to indicate that operational spacetime disappears at sufficiently small scales due to quantum effects. All of our empirical science, however, seems to be predicated on the idea that scientific data about the world is recorded by spatiotemporally located devices. Therefore, if some particular theory of quantum gravity without spacetime is true, then the truth of that theory would undermine any empirically-motivated reasons we might have for believing in it. Huggett and Wüthrich's [59] response is to point out that there is no requirement that spacetime be fundamental; an emergent account of spacetime would suffice as long as such an account can be robustly provided. Indeed, this is arguably an arena in which spacetime functionalism is of significant value.

Our reasons for being interested in theoretical spacetime are less obvious, beyond its being a necessary component of an account of operational spacetime. Here, I offer one reason for our interest in theoretical spacetime for its own sake—its role in our metaphysical theorising about the world. Several debates in metaphysics turn on accounts of space and time—causation, causality, identity, modality, and the nature of time itself, to name a few; as I mentioned in the introduction to this thesis, space and time are central to most metaphysical theorising. But getting bogged down in the complicated dynamical details of matter fields in the actual world would preclude a great deal of this theorising. Which is why so-called 'kinematical' facts about space and time play such a vital role—they abstract away from conceptually irrelevant contingent characteristics of particular dynamical theories, and allow for general claims to be advanced about the world. Take causation and causality, for example. It is now taken for granted that faster-than-light communication is prohibited by relativity theory. But it is extremely uncommon (I know of no examples) for a metaphysical text to demonstrate this fact by, for example, solving Maxwell's equations in the geometrical-optical limit in all relevant spacetimes and demonstrating the physical significance of the light-cone structure of each tangent spaces of these spacetimes. Far more common (and sensible) is to begin by just assuming that this structure is, in fact, a feature of spacetime.

It will turn out that the three procedures outlined in this section—Earman's, Brown's and Knox's—all agree that superspace is spatiotemporal. However, the utility of this fact is dampened by the severance of the link between theoretical spacetime structure and operational spacetime structure. This link, in commuting spacetime theories, is provided by the *weak clock hypothesis*:

Weak clock hypothesis: There exist configurations of matter whose dynamics is such that they can be used, in inertial as well as suitably gently accelerated states of motion, to record intervals of proper time, as determined by some metric, along their worldlines.[8]

In special relativity, the clock hypothesis is defined with respect to the Minkowski metric; in general relativity, with respect to the metric g_{ab} of the Einstein field equations. The local symmetry coincidence given by theoretical spacetime structure, guarantees the Lorentzian signature of the metric, at least for matter fields described by the (classical approximation of the) standard model. Given local symmetry coincidence of the dynamical laws and the geometry, the minimal condition one needs to satisfy to be able to have epistemic access to operational spacetime *approximately*, is that certain matter configurations are *boostable*, i.e. they are just as good at reading off intervals of proper time along a metric before and after an active Lorentz boost, even if not *during* the boost. Call the hypothesis that such matter configurations exist the **approximate clock hypothesis**. This is a strictly weaker condition than satisfaction of the weak clock hypothesis. In what follows, I will not discuss the approximate clock hypothesis.

[8] This restriction to 'suitably gently accelerated states of motion' is important. To expect the clock hypothesis to hold for *any* state of motion is far too restrictive; every physical clock has some breaking point.

The weak clock hypothesis is a field-dependent statement—it makes reference to certain field configurations. However, it is standard to think of it as being field-neutral, i.e. the choice of the metric with respect to which the hypothesis is defined is chosen in such a way as to be independent of the choice of matter field.[9] Let us therefore refine the above definition to a *strong clock hypothesis* (this is what is often referred to simply as 'the clock hypothesis'. See e.g. [44, 77]):

Strong clock hypothesis: For each kind of matter, there exist appropriate configurations whose dynamics is such that they can be used to record, in an inertial or suitably gently accelerated state of motion, intervals of proper time along their worldlines, as determined by some metric. Further, this is the metric with respect to which clocks built out of other matter fields will record proper times along their worldlines.

Satisfaction of the strong clock hypothesis then guarantees that theoretical spacetime structure determines operational spacetime structure. The strong clock hypothesis is what ensures the universality of a class of spatiotemporally-oriented metaphysical theses. We will see, in §3.4, that the kinematical structure of superspace is not as friendly to such metaphysical theorising as Minkowski spacetime is, because of the failure of the strong clock hypothesis.

3.3 Supersymmetric field theory and superspace

In order to extend the arguments presented in the previous section, we will need a small number of new formal tools. The SUSY presented in this chapter is sufficient for this purpose; it does not constitute anything like an introduction

[9] For the proponent of the dynamical approach, the physical utility of this move is a brute unexplained fact; for the proponent of the geometrical approach, this is explained by the constraining power of the metric.

to SUSY for philosophers. In particular, I make a number of important technical restrictions—I consider only *superclassical* fields—these are fields that have the algebraic properties of bosonic and fermionic fields, but are themselves classical field-theoretic and not quantum field-theoretic objects[10]—that are symmetric under rigid (i.e. constant across the space of independent variables) and non-extended (i.e. there is only one irreducible spin representation of the SUSY generators) SUSY transformations. The physics presented in this section is based on Weinberg [124] and Wess and Bagger [127]; the mathematics on De Witt [34], Rogers [98], and Kuzenko and Buchbinder [23].

SUSY can be approached from a number of directions; in this chapter, our trajectory is mapped out by our interest in SUSY as a (putative) spacetime symmetry. I will therefore mostly ignore its quantum field theoretic roots, and present it as a generalisation of a spacetime symmetry (in the sense of §3.2). This section is structured as follows. In §3.3.1, I define, in the most general possible terms, SUSY transformations and then discuss how SUSY, a transformation between particle types, can be given an interpretation as a spacetime symmetry. In §3.3.2, I introduce the necessary mathematical concepts to set up superspace. Finally, I present, in §3.3.3, a simple supersymmetric field theory.

3.3.1 What is supersymmetry?

SUSY is a collective term for a set of transformations between *bosons* (which typically represent the carriers of force, such as photons, w-particles, or gravitons) and *fermions* (which typically represent the quanta of matter, such as electrons or quarks) that leave the Lagrangian density associated with a particular theory invariant (up to a possible surface term). This is the first symmetry of its

[10] A quantum field is standardly interpreted as a spacetime-indexed operator-valued distribution.

kind in physics—until now, all symmetry transformations have been such that bosons were taken to other bosons and fermions to other fermions, leaving a force-matter distinction invariant.

A natural expectation, on learning the definition of SUSY, is that an even number of applications of a SUSY transformation returns the original particle type (for example, a boson would switch to a fermion then back to a possibly distinct boson). This is indeed what happens, but the process has a surprising consequence—the transformed boson is displaced in spacetime compared to the original. It is this novel mixing of symmetry types that gives SUSY such philosophical promise.

One can learn everything about the symmetry structure of a physical theory by studying the groups associated with the symmetry transformations under consideration. So the study of rotations in Euclidean space, for example, reduces to the study of the *orthogonal group*. Physicists often use a trick to simplify the mathematics—they focus on infinitesimal versions of the transformations, and study the associated *algebras*. For most theories, this reproduces most of the interesting structural features of the theory at a fraction of the mathematical cost. We can, therefore, learn a good deal about SUSY (whose dynamical symmetry group is the *super-Poincaré* group) by studying its algebra, known as the *Super-Poincaré* algebra.[11]

Infinitesimal SUSY transformations are represented by an operator, Q_α: applying Q_α to a boson, ϕ, turns it into a fermion, ψ. The process of 'acting on an object' e.g. turning a boson into a fermion is enacted, mathematically by finding a *representation* of the algebra element, Q_α. The algebraic structure of

[11] Technically, the super-Poincaré algebra is a *superalgebra*, corresponding to the super-Poincaré group, which is a *super-Lie group*. Their role in SUSY theories is analogous (with certain technical caveats explored in [98]) to the role of a symmetry algebra and its corresponding group in ordinary field theory.

the set of such operators is encoded in what is known as the *anti-commutator*: for each pair of operators Q_α and Q_β, the sum of the results of acting with Q_α first and Q_β second, and the other way around. It encodes the information about the repeated application of the transformation to which we referred above. We find, on working through the mathematics, that the anti-commutator of two of these generators is, completely unexpectedly, proportional to the generator of spacetime translations.[12] We therefore conclude that there is something inherently spatiotemporal, at least at the infinitesimal level, about SUSY. In other words, a purely *internal* transformation (i.e. one which explicitly does not involve changing the position or orientation of an object in spacetime) between kinds of matter, has somehow led to a spacetime displacement.

It is natural to think of a symmetry transformation as being something that does not change the identity of the object on which it acts. The fact that the dynamics does not see this change is an indication that this is a good inference. A rotation of, say, a plate might change its configuration, but it does not change the fact that the rotated object is still a plate, albeit in a (possibly indistinguishable) different state. So how do we make sense of a *symmetry* transformation between a boson and a fermion, when it is clear that both of these objects are *distinct*.[13] This is a two-stage process.

The first stage is to realise that, if SUSY is, in fact, a symmetry of the actual world, then the current state of the world is one in which the symmetry is *broken*. We come across broken symmetries all the time, even in classical

[12] The explicit form of the super-Poincaré algbera relations is:

$$\{Q_a, Q_b^\dagger\} = (\sigma^\mu)_{ab} P_\mu$$

$$[Q_a, P_\mu] = 0$$

$$[Q_a, M_{\mu\nu}] = (\sigma_{\mu\nu})_a{}^b Q_b$$

[13] Even referring to bosons as objects is not uncontroversial. See [99].

physics. The laws of mechanics are translation-invariant, but that does not mean that if I suddenly move myself twenty feet to the right, that I won't notice a change. What has happened is that my state is no longer symmetric under all the transformations under which the laws are symmetric.

The second stage is to consider the regime of unbroken SUSY, and think of bosons and fermions as being components of a more general type of field. This is known as a *superfield*, and is the focus of §3.3.3. In order to define a superfield, we need to be able to talk about supernumbers and superspace.

3.3.2 Supernumbers and superspace

Supersymmetric laws are invariant under a transformation of a bosonic field to a fermionic field and vice versa. To simplify things a little, we work with so-called superclassical fields. Quantum operators, as elements of an algebra that is not necessarily commutative, are the sorts of things that can commute or anti-commute.

All classical fields, on the other hand, commute, so they are represented as maps from spacetime to some commuting space (usually this is also a manifold and vector space; standard examples are $n-$tuples of real or complex numbers). To construct a superclassical field, all we do is widen the scope of the target space of the maps to include spaces of anticommuting numbers (and tuples thereof). We do not interpret these as operator-valued distributions; instead we see that they take their values in spaces of (possibly tuples of) number-like objects known as *supernumbers*.

Supernumbers

In the introduction, we characterised supernumbers as generalisations of real and complex numbers in such a way as to have non-trivial commutation properties. An explicit construction of these objects, as elements of an infinite *Grassmann algebra*, can be found in [23]. For the purposes of this chapter, we can think of supernumbers as elements of a space of objects (of uncountable cardinality) which splits into two subspaces. The first consists of (a set of objects isomorphic as an algebra to the) real numbers, and other objects that commute with each other (and the real numbers), the second consists of objects that *anticommute*. If ξ_1 and ξ_2 are supernumbers, then if $\xi_1 \cdot \xi_2 = -\xi_2 \cdot \xi_1$, but $\xi_1 \neq 0 \neq \xi_2$, these supernumbers anticommute. In many respects that are significant to the physics, they behave like ordinary scalars—they can be added, subtracted, multiplied and divided. These spaces are, respectively, the space of real commuting supernumbers, denoted by $\mathbb{R}_c$, and the set of all real anticommuting supernumbers, $\mathbb{R}_a$. In what follows, $x^a \in \mathbb{R}_c$ and $\theta^a \in \mathbb{R}_a$. It is notationally convenient to define θ (and its complex conjugate, $\bar{\theta}$) as belonging to the *complex* anticommuting supernumbers—these are merely ordered pairs of real anticommuting numbers.

We now have the resources to define the classical analogues of fermionic fields as maps from some spacetime manifold (more on this in §3.3.3) to the space of anticommuting supernumbers (and tuples composed therefrom), and the classical analogue of bosonic fields as maps into the space of commuting supernumbers and tuples composed therefrom (this is a space that includes the real numbers as well).

Superspace

We would like to write down classical field equations which are symmetric under interchange of commuting and anticommuting fields, and impose this symmetry in the standard field-theoretic way, by considering representations of the algebra of infinitesimal generators. Since we are dealing with fields which transform into one another, and given that the two fields have very different algebraic properties, we need a new kind of algebra. It turns out that what we need are infinitesimal anticommuting generators, infinitesimal commuting generators, and a rule for how they interact. This is the information contained in the super-Poincaré algebra.

Generalising the spacetime manifold is the crucial step that allows us to construe superspace as being spatiotemporal; this generalised manifold is the titular superspace. It is standard to construct a manifold by beginning with a set of points, call it M, defining maps from these points into the real numbers (i.e. coordinatisations), and then specifying the structure on M by specifying the allowed coordinatisations of M. This is the Kleinian approach to geometry (for more details, see e.g. [3, 120]), and modulo limited success with trying to define this structure nominalistically (see [6, Ch. 8]), it is the only way in which we can construct a smooth manifold.

A supermanifold is composed by analogy with an ordinary manifold, using supernumber-valued coordinatisations.[14] It is the product space of real commuting supernumbers and real anticommuting supernumbers (under the DeWitt topology and supersmooth structure; see [98]). Formally, the definition

[14] There are a number of thorny technical issues that come up. For example, because of the nilpotent elements, there is no natural Hausdorff topology on the space of supernumbers in the way that there is on the real numbers. So the generalisation to a supermanifold loses some of the physically desirable qualities of ordinary manifolds. For more details on this construction, see [98, Chs. 2-6]

of a superspace, denoted by $\mathbb{R}^{p|q}$, is

$$\mathbb{R}^{p|q} := \mathbb{R}^p_c \times \mathbb{R}^q_a = \{z^M = (x^1, x^2, \ldots x^p, \theta^1, \theta^2, \ldots \theta^q),\ x^m \in \mathbb{R}^p_c,\ \theta^\mu \in \mathbb{R}^q_a\}.$$

A point in such a space now has $p+q$ coordinates, p of which are commuting supernumbers, q of which are anticommuting supernumbers. An infinitesimal translation in the superspace coordinates can be easily defined:[15] The transformations are

$$x'^a = x^a - i\epsilon\sigma^a\bar{\theta} + i\theta\sigma^a\bar{\epsilon}, \tag{3.1}$$

and

$$\theta'^\alpha = \theta^\alpha + \epsilon^\alpha, \tag{3.2}$$

where the α index runs from 1 to q, and σ^a are the Pauli spin matrices. A similar expression to (3.2) holds for the complex conjugate coordinate $\bar{\theta}$. With these transformations, along with the standard Lorentz transformations,[16] it is easy to find an expression for the analogue of a Minkowski interval for superspace, a function that is left invariant by these transformations. This two point function, defined below, is the super-Minkowski interval.

$$w^a(z_1, z_2) = (x_2 - x_1)^a + i\theta_\alpha\sigma^a(\bar{\theta}_\beta - \bar{\theta}_\alpha) - i(\theta_\alpha - \theta_\beta)\sigma^a\bar{\theta}_\alpha. \tag{3.3}$$

where z_1 and z_2 are coordinates in superspace. The super-Minkowski interval, arrived at here using coordinates, can be thought of as describing an absolute geometric object (in the sense of §3.2), call it the *super-Minkowski metric*.

[15] A discussion of this can be found in [23, p. 158]

[16] It is worth noting here that there is no 'super' analogue of a Lorentz transformation; the super-Poincaré group is the semi-direct product of the ordinary Lorentz group and the super-translation (super)group.

3.3.3 Superfields

A real classical scalar field $\phi(x)$ can be represented as:

$$\phi(x) : \mathbb{R}^4 \to \mathbb{R}_c. \tag{3.4}$$

This idea can be generalised to incorporate a set of both commuting as well as anticommuting (the classical analogue of fermionic) fields. Each element of a set of dynamical fields $\{\phi^i\}$ can be represented as

$$\phi^i(x) : \mathbb{R}^4 \to \mathbb{R}^{p|q}. \tag{3.5}$$

where p is the number of commuting fields components and q the number of anticommuting ones.[17]

Finally, one might generalise this procedure to produce an object, call it a *generalised superfield*, Φ, that transforms as a generalised rotation in superspace, much like a Lorentz transformation can be thought of as a rotation in spacetime:

$$\Phi^i : \mathbb{R}^{4|4} \to \mathbb{R}^{p|q}. \tag{3.6}$$

With the setup presented above, it is possible to combine the hitherto separate types of field—bosonic and fermionic, into a single field, Φ. At this stage, there is nothing radical about this move. Expressing fields in terms of combinations of other fields is a common move in QFT—it is the first step in constructing a Lagrangian for a field theory with a spontaneously broken symmetry, for example. Consider, now a mapping from this superspace (supermanifold) into

[17] We temporarily gloss over the fact that we need to do some more work in order to ensure that certain essential features of our fields, like smoothness, translate appropriately into this formalism—the short answer that they do. A precise definition of what it is for a function to be 'supersmooth', for example, can be found in [98].

the superspace of real commuting supernumbers

$$V : \mathbb{R}^{4|4} \to \mathbb{R}_c$$

On Taylor expanding this superfield in its anticommuting coordinates, we recover all the expected bosonic and fermionic fields (for simplicity, I have returned to using θ_1 and θ_2 instead of θ_α and $\bar{\theta}_\alpha$):

$$V(x, \theta_1, \theta_2) = A(x) + B(x)\theta_1 + C(x)\theta_2 + D(x)\theta_1\theta_2 \tag{3.7}$$

where $A(x)$ and $D(x)$ are commuting (bosonic) fields, $B(x)$ and $C(x)$ are anticommuting (fermionic) fields.

Notice how, by incorporating the supernumber-valued coordinates into the argument of the function V, we have been able to express both fermionic and bosonic fields (we do not, as yet, know that they are bosonic/fermionic—it is their transformation behaviour under the Poincaré group/algebra which determines that) as functions of commuting coordinates, weighted by the appropriate commuting or anticommuting factor. The mathematical advantage of this superspace setup is clear—the equations of motion associated with these fields are manifestly supersymmetric. In the next section, we discuss the physical significance of interpreting superspace as a theoretical spacetime.

3.4 On the spatiotemporality of superspace

We now have all the pieces in play to decide on the spacetime structure of a classical flat supersymmetric field theory. In §3.4.1 I revisit the three approaches discussed in §3.2. I then discuss the issue of chronometry, in §3.4.2 and see how the link between theoretical and operational spacetime is severed in superspace.

3.4.1 The three approaches to spacetime

Earman's Principle

Using the generalised mathematical notions introduced in §3.3, we can set up a supersymmetric field theory in a model-theoretic way that allows us to articulate Earman's Principle. The models of a flat superclassical field theory are tuples of the form $\mathscr{M} = \langle SM, A^i, P^i \rangle$, where SM is a supermanifold. The dynamical symmetries are elements of the super-Poincaré group, by construction. Since this is the largest common symmetry group, by the arguments presented in §3.2, this qualifies as the analogue of the external symmetry group as well. So Earman's Principle mandates that these be the symmetries of the absolute objects as well; the 'spacetime' symmetry group. This is why one the absolute object characteristic of this 'spacetime' is a generalised tensor superfield that defines the super-Minkowski interval on the supermanifold.

Contrast this with the alternative approach that maintains that the absolute object be only the Minkowski metric tensor, η_{ab}. The symmetries of this object are just the Poincaré transformations. But the dynamical symmetry group continues to be the super-Poincaré group. This leads to an unacceptable mismatch, and a violation of Earman's Principle. Therefore, the proponent of this principle has to accept that superspace (equipped with the relevant geometrical object that specifies the super-Minkowski interval) is a theoretical spacetime.

The Dynamical Approach

On the dynamical approach, we begin with the fields themselves, and ask what their (largest) common dynamical symmetry group is. There is no underlying commitment to an independently existing spacetime metric, to which we might only have epistemic access through the behaviour of matter fields; the spacetime metric is itself nothing more than a codification of the symmetry properties of matter fields. Therefore we are forced to conclude that, for a universally super-Poincaré-invariant theory, only superspace equipped with the super-Minkowski metric function could be spatiotemporal; the ordinary Minkowski metric simply does not encode all aspects of the universal behaviour of matter fields.

Spacetime Functionalism

Equipped with the generalisation of the Minkowski interval, the super-Minkowski interval defined in (3.7), it is possible to construct the superspace analogue of inertial frames, called *super-inertial* frames.[18] Recall that a crucial ingredient of Knox's definition of inertial frames was a universality requirement. This requirement can only be satisfied by taking the relevant relativity principle to be the one expressed by the statement that the largest common external symmetry group is the super-Poincaré group—this way we avoid making reference, in our definition of inertial frames, to velocities and force-free motion. We can then run the generalisation of Knox's argument in superspace—what it is to be a spatiotemporal degree of freedom is to play a role in determining the super-inertial structure of a theory.

It is worth discussing Knox's programme in a little more detail, in light of arguments from chapter 2. I should make it clear that superspace makes for a different sort of counterexample compared those identified in chapter 2 (universally coupled massive scalar gravity and Gödel spacetime); whereas in those cases, I argued that theoretical spacetime identified the *wrong* operational spacetime structure, here, I contend that superspace cannot be given an opera-

[18] For more on these frames, see [23, Ch. 2].

tional meaning at all.

The set of dynamical symmetry transformations common to all superfields is, by construction, the *super-Poincaré group*—this is an example of a 'super-Lie group'. Super-Lie groups are a generalisation of Lie groups; whereas the latter are groups which are also smooth manifolds, the former are groups which are also superspaces (i.e., generalisations of manifolds with anti-commuting supernumber-valued coordinates). Recall that a theoretical spacetime is a geometrical structure determined by constructing geometrical objects that are invariant under transformations from the dynamical symmetry group common to all fields in question. By construction, therefore, superspace is a theoretical spacetime.

The Poincaré group is a subgroup of the super-Poincaré group, *a fortiori*, any Poincaré transformation is a dynamical symmetry of a supersymmetric theory. We therefore have a *symmetry-coincident* theoretical spacetime. If, in addition, we assume the light postulate, or some equivalent statement of an invariant quantity for these tangent spaces, we can construct a *cone-coincident* theoretical spacetime, as far as the Poincaré transformations go.

Knox's prescription identifies as 'inertial' those frames in which the laws take their simplest form. In the case under consideration here, these are those frames in which the dynamical laws governing the superfields are expressed in coordinates in which the super-Minkowski interval, determined by the super-Minkowski metric is invariant. Thus, the super-Minkowski metric plays the role of picking out 'inertial' frames (these frames are sometimes referred to as 'super-inertial frames'; see e.g. [23])—it therefore qualifies as cone-coincident theoretical spacetime.

To see if the super-Minkowski metric qualifies as an operational spacetime, however, we need to assume that one can construct the appropriate generalisation of rods and clocks. Let us call these objects *super-surveyors*. An example of a super-surveyor would be the analogue of a light clock, but one in which the oscillating material is governed by the supersymmetric version of Maxwell's equations, the so-called *super-Maxwell* equations. But, as is discussed in more detail in the next section, the extra SUSY dimensions *cannot* be surveyed by such matter field configurations.

3.4.2 Chronometry in superspace

What does it mean to survey a geometrical object in superspace? What does one gain by thinking of a supersymmetric theory as existing in superspace? A cynical, but not wholly inaccurate characterisation of this chapter is as an exploration of the repercussions of prefixing objects from discussions in the foundations of spacetime with the word 'super-'. Let us, therefore, look at 'super-rods' and 'super-clocks'—call such matter configurations *super-surveyors*. In discussions in the foundations of relativity theory, surveyors are generally treated as epistemologically convenient matter configurations whose behaviour is such that they allow us to measure lengths and durations as given by the underlying spacetime metric. Nothing in the foundations of relativity theory, therefore, requires surveyors to be primitive.

What is it to be spatial?

Minkowski famously made a proclamation in 1908 about the fate of space and time as distinct concepts:

> 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 [78].

There is a sense in which this is true—the symmetry group of special relativity includes transformations on coordinates that 'rotate' spatial coordinates into temporal ones and vice versa. However, this does not mean that physically, our best theories do not distinguish between spatial and temporal directions—they had better, given the starkly different ways in which spatial and temporal evolutions are treated (cf. e.g. [106]). Callender makes the important observation that 'although relativity banishes "time" and "space", there is nevertheless a sense in which it is committed at its very core to a difference between timelike and spacelike directions' [29, p. 124].

What Callender means by this is that, even in Minkowski spacetime, there is a physical distinction between spacelike and timelike directions. Moreover, this distinction survives a Poincaré transformation—there is no such transformation that will change a spacelike direction into a timelike one or vice versa. And this fact is encoded in the structure of the Minkowski metric tensor, η_{ab}.

In its most general form, a *relativity principle* asserts the equivalence of a description of physics in one frame of reference with that in another suitably related frame of reference. The Galilean relativity principle, for example, asserts that the form of dynamical equations is the same in all inertial frames, i.e. frames that are moving at a constant velocity with respect to each other. A relativity principle brings along with it a notion of invariance. Importantly, this invariant quantity does not need to be put in by hand; recall that in special relativity, for example, it follows from another dynamical postulate, the *light postulate*. For Galilean relativity, the invariant notion is of spatial distance—the length of a rod before and after any Galilean transformation is unchanged. What is distinctive about the special relativistic relativity principle is that the invariant quantity is a *speed*, not a spatial length. This is, of course, a trivial consequence of assuming the light postulate, which asserts that the two-way speed of light is isotropic and independent of the speed of the source. Without the light postulate, the existence of *some* invariant quantity is guaranteed by the relativity principle. The light postulate adds three physically salient elements to this—(i) that the invariant quantity is a *speed*, (ii) the actual value of the invariant speed and (iii) that some physical matter fields (i.e. Maxwell fields) exist with that speed of propagation.

Spatiotemporal ≠ spatial+temporal

A new complication arises in superspace for super-surveyors. The extra dimensions are not straightforwardly *spatial* (or temporal) in the way that the spatial dimensions in Minkowski spacetime are. In the definition of the translation rule for coordinates on Minkowski spacetime, there is no mixing of spatial and temporal coordinates—coordinate mixing arises from transformations in the Lorentz subgroup of the Poincaré group. In the superspace considered here, the symmetry group is the super-Poincaré group, which is the semi-direct product of the Lorentz group with the group of *supertranslations*. As equation (3.1) illustrates, a supertranslation leads to coordinate-mixing—all of the commuting coordinates, x^μ, pick up a non-trivial fermionic coordinate when supertranslated. These fermionic coordinates defy a straightforward interpretation as being purely spatial, thanks to their anticommuting nature. Therefore, if a symmetry transformation exists that takes a bosonic coordinate, like x^0 and maps it to a mixed bosonic-fermion coordinate, then that coordinate can also no longer be given a purely spatial or purely temporal interpretation. Of course, one could still argue that these dimensions are spatiotemporal in some sense (e.g. in the sense of a theoretical spacetime) given the preceding discussion about external dynamical symmetry coincidence. Thus, we find ourselves in the interesting position of having to countenance the idea that a dimension might be spatiotemporal without being explicitly (in a given frame) spatial or temporal.

There is, however, no generalisation of the light postulate to the entirety of superspace (of course, the original light postulate will still apply to the Minkowski subspace of super-Minkowski space). Whatever the invariant quantity is, according to the relativity principle in superspace, it is certainly not a speed. Indeed, the concept of speed is no longer well defined, given the absence of a supertranslation-invariant notion of time and space, as the transformation law for the temporal direction demonstrates. If we believe that our surveying devices have to be built out of SUSY-matter (and that seems reasonable; how else could they survey the superspace geometry?) then it is possible that this matter will not read off intervals of the super interval, even if, with respect to the Minkowski subspace of superspace, the strong clock hypothesis may be satisfied.

I do not intend to suggest that supersymmetric matter will somehow propagate faster than light on its Minkowski dimensions; the structure of the super-Poincaré group ensures that it will not. But generic features of superspace geometry will be accessible to super matter if it satisfies a generalisation of the clock hypothesis, the *super-surveyor hypothesis*:

Super-surveyor hypothesis: There exist matter configurations that read off generalised metric intervals along the trajectories of particles in any physically suitable state of motion. Further, this is the generalised metric with respect to which super-surveyors built out of other matter fields will record intervals along their worldlines.

We have good reason to suspect the clock hypothesis is approximately satisfied; this evidence comes from observation, not from any specific prediction from within general relativity. In other words, the clock hypothesis cannot be derived from within general relativity.[19] In the absence of empirical or extra-theoretical reasons to believe in the satisfaction of the super-surveyor hypothesis, we cannot make claims about super-surveyors as we do in special or general relativity. Our only alternative is to look for a mathematical indication from within supersymmetric field theory that a super-surveyor hypothesis could be satisfied, but indications from our attempts to do so in special and general relativity are not promising. In what follows, I provide two short heuristic arguments to cast doubt on the possibility of such a proof in a superclassical Yang–Mills theory based on disanalogies with Maxwellian electromagnetism in special relativity.

First, the clock hypothesis in special relativity (as it applied to light clocks) relied on the validity of the light postulate, which as we have seen, asserted a claim about the independence of the speed of light from the speed of the source of that light. If one looks for a justification from within Maxwellian electromagnetism, one eventually hits some different phenomenological postulate. For example, that the speed of light is constant in all inertial frames follows from the statement that the magnetic permeability, μ_0 of the classical vacuum and the electric permitivity of the classical vacuum, ϵ_0 are the same in all inertial frames. We can thus arrive at the light postulate in reverse, from assuming the frame-independence of the values of these constants. But ultimately those facts are merely postulated, and ultimately only justified by appeal to empirical

[19] This was one of the morals of chapter 1.

observation, the likes of which we have no access to in superspace.

The surveyors we have considered in this thesis so far have all been modelled as Langevin clocks—perfect reflectors with some oscillating material bouncing between them—which emphasise that the crucial link between clocks and chronometry is through the notion of an oscillation. Oscillations, whether of some material in a Langevin clock, or related to some quantum process (e.g. flavour oscillations of neutrinos [2, 66] or even related to the periods of binary systems of black holes, are a necessary component of measurements of proper time (or, more generally, parameters along trajectories that are frame-independent). But oscillations require a time parameter, and as we have seen, the supertranslation symmetry of superspace means that the previously temporal dimension picks up an fermionic component. And these extra dimensions, which, by the supertranslation transformation are equivalent to the ordinary Minkowski dimensions, have no finite extension.

Whether or not this means that the super-surveyor hypothesis is doomed in superspace is unclear, although the absence of a sensible notion of an oscillation means, at the very least, we will need to come up with a radically new method of surveying the supermetric if the hypothesis is to be satisfied. A more detailed analysis is beyond the scope of this chapter; the aim here was to assess the verdict that Earman, Brown and Knox would pass on the spatiotemporality of superspace. What is clear is that we cannot fall back on phenomenological principles in superspace in the way that we could in special relativity, in order to bootstrap our way up to a theory that satisfies the (super-)surveyor hypothesis. As we saw in chapter 2, the relativity principle on its own does not determine the invariant quantity. And the efficacy of any sort of surveyor hypothesis depends crucially on this. Simply knowing about the structure of, say, the super-Minkowski metric, does not immediately give us generic (i.e. dynamics-independent) chronometric facts about the behaviour of superfields—in superspace, as in certain general relativistic solutions, we do indeed need to get our hands dirty with the physical details of individual superfields and, in particular, with the details of the theory that point towards what, if any, the physically relevant invariant quantity is.

The upshot of this uncertainty about whether superspace is an operational spacetime is, however, that prescriptions due to Earman, Brown and Knox do not guarantee that superspace will count as an operational spacetime, even though they all unequivocally categorise superspace as a theoretical spacetime.

3.5 Conclusion

This chapter had two primary aims: (i) to argue that superspace (equipped with suitable geometry) is the correct theoretical spacetime setting for supersymmetric field theories and (ii) to demonstrate that merely identifying theoretical spacetime structure does not guarantee operational information about how measuring devices (and therefore matter fields in general) would behave. As a result of this severance, a large class of spatiotemporally-oriented metaphysical claims will have to be reassessed.

A broader secondary aim was to present a first attempt at identifying some of what might go into an interpretation of a spacetime compatible with (or derived from, as the case may be) a theory of quantum gravity. To argue that our final theory of quantum gravity, whatever it may be, will make it necessary to reject much of what we take for granted in such debates (geometric features of the space of independent coordinates, for example). But absent such a theory, our speculation and philosophical insights are best grounded in extant debates and points of view on the philosophy of spacetime physics. As Pooley [94, p. 183] maintains, 'one is likely to do more justice, not less, to the conceptual novelty of [general relativity] by seeking as much common ground with previous theories as possible.' The same, I think, applies to SUSY. Finding that common ground allows us to identify which assumptions that go into our metaphysical theorising are innocuous and which are not, arguably a noble aim for any scientifically-informed metaphysics.

Part II

The dynamical approach to spacetime

CHAPTER 4

Fibre bundles and the dynamical–geometrical debate

4.1 Introduction

Here are two competing views on the purpose of this project of casting general relativity as a Yang–Mills theory in the fibre bundle formalism: (i) we can take cues from Yang–Mills theories to help us interpret general relativity; (ii) we can take cues from general relativity to help us interpret Yang–Mills theories. Recent papers by Wallace [119] and Weatherall [122] have sought to argue for the utility of formulating general relativity using fibre bundles, but in starkly different ways and to starkly different ends. Weatherall opts to formulate general relativity as a theory of a dynamical connection on a frame bundle, which induces a connection on an associated vector bundle which is identified with the tangent bundle. This leads to a *prima facie* disanalogy with other Yang–Mills theories, which, while also theories of similarly induced connections on associated bundles, do not identify those bundles with the tangent bundle. Faced with this disanalogy, Weatherall argues against its physical significance. Based on this argument, he proposes a deflationary account of gauge transformations, arguing, in effect by analogy with the manner in which general relativity, although itself a fibre bundle theory, has traditionally been interpreted without resort to fibre bundles.[1]

Wallace, on the other hand, takes that disanalogy to mandate a different formulation of general relativity as a Yang–Mills theory. For him, general relativity is a theory of a dynamical connection on a principal *Poincaré* bundle, which induces a connection on an associated *affine* bundle, *not* identified with the tangent bundle.[2] In this chapter, I argue that Wallace and Weatherall's first-order disagreement over the importance of the disanalogies between general relativity and Yang–Mills theories is symptomatic of a broader disagreement, which is rooted in the dynamical–geometrical debate set up in the introduction—indeed, Weatherall explicitly describes his approach as providing a 'geometrical' interpretation of Yang–Mills theories. As a consequence of arguing for this view of the disagreement, this chapter will also present the dynamical–geometrical debate in a mathematically precise and novel manner, as a disagreement over whether or not to accept a particular restriction on the action of a certain geometrical object, the so-called *solder map*. The primary goal of this chapter is to argue against the viability of Weatherall's deflationary story about gauge.

This chapter is structured as follows, in §4.2, I introduce, in some detail,

[1] In this sense, Weatherall echoes a sentiment expressed by Trautman [112, p. 287]: 'Physicists were using concepts that are now part of the theory of fiber bundles before mathematicians introduced the notion of a bundle... In this respect the situation of physicists can be likened to that of Monsieur Jourdain *qui fait de la prose sans le savoir*. There thus arises the question whether it is worth while to learn the language and use the methods of fiber bundles since so far it has been possible to do without them.'

[2] Wallace is not the first person to propose reformulating general relativity as a gauge theory of the Poincaré group; This was first proposed by Kibble [65], and has been discussed more recently in the philosophy literature by Healey [55, Ch. 3.2]. However, both Kibble and Healey, unlike Wallace, view the base space as 'spacetime'. As a result, Wallace's presentation is better suited to the requirements of this particular chapter.

fibre bundles, which provide the mathematical formalism appropriate to the formulation of Yang–Mills theories. Then, in §4.3, I introduce Wallace's account of general relativity as a Yang–Mills theory. In §4.4, I introduce Weatherall's competing formulation of general relativity as a Yang–Mills theory, and argue that his deflationary proposal for an interpretation of gauge symmetries, based on his construal of general relativity as a Yang–Mills theory, is not viable. Finally, I use the setup from the previous sections to explicate the dynamical–geometrical dispute.

4.2 Parametrised field theory

In this section, I introduce the mathematics of fibre bundles, and in particular, the important relationship between so-called *principal G-bundles* and their *associated vector bundles*. I have decided to include definitions in the main text but have omitted theorems (relevant references are provided in footnotes). The definitions are presented as required in the sequence of the text, rather than as a list at the beginning, or in an appendix. In addition to definitions, I have included, as often as possible, brief descriptions of the physical intuitions and conveniences that these mathematical structures are intended to capture. I assume the reader's familiarity with some differential-geometrical tools that are common to discussions in the foundations of spacetime, including *manifolds, Lie groups* and *tensors.*[3]

This section is structured in such a way as to begin with the intuitive notion of a field as a map from a manifold into some (vector) space, and end with a picture of a Yang–Mills gauge theory as the theory of two objects: a *matter field*

[3] The definitions of these objects that I am tacitly working with are those found in Wald [116], while the bundle-related material is based on Nakahara [81].

as a section on a vector bundle and a *gauge field* derived from a connection on a principal G-bundle.

4.2.1 Field theory

Faraday suggested thinking of a magnetic field as being composed of lines of flux in space. This conception lends itself (but is not committed) to the picture of fields as being properties of 2-surfaces in spacetime.[4] The fact that interactions between different fields are localised (in some idealisation) to these same spacetime regions lends credence to the viability of this picture. So let us begin with this intuition, and look at its mathematical representation. Call this sort of field—i.e. a collection of properties of 2-surfaces—a *Faraday field.*

A Faraday field theory consists of a smooth manifold, $\mathcal{M}$, certain tensorial objects like metric fields, and a function from that manifold to some mathematical space, $\mathcal{V}$, most often assumed to have vector space structure. Most of the time, there are more mathematically possible values of the field than there are empirically distinguishable states of the system. This means that certain transformations relate empirically indistinguishable states of affairs. These are the *symmetries* of the theory, represented by certain structure-preserving maps from $\mathcal{V}$ to itself—these are its *automorphisms.* If $\mathcal{V}$ is a vector space, then the largest possible symmetry group is the group of invertible linear maps, denoted by $GL(n, \mathbb{K})$, where n is the dimension of $\mathcal{V}$ and $\mathbb{K}$ is its underlying field of scalars.

In general, the automorphisms of $\mathcal{V}$ are required to preserve more than

[4] The clarification that tensorial properties should be seen as properties specifically of 2-surfaces in spacetime, rather than straightforwardly properties of spacetime is made especially clear by Butterfield [28], as part of a larger project to excise from (classical) physics and geometry the idea that the fundamental quantities of a theory are intrinsic properties of spacetime points [25, 26].

just its linear structure, so the symmetry group of a dynamical theory will be a subgroup of $GL(n, \mathbb{K})$. This additional structure arises from the requirement that $\mathcal{V}$ be a carrier space of some representation of some group. Such spaces are also known as *representation spaces*.

Definition: A *group representation*, T, is a group-structure-preserving map from some group, G, to the group of invertible linear maps, $GL(n, \mathbb{K})$, i.e. (i) $\forall g_1, g_2 \in G$, $T(g_1 \circ g_2) = T(g_1) \cdot T(g_2)$, and (ii) $T(g^{-1}) = T(g)^{-1}$ where $\circ$ is the group operation in G and $\cdot$ is the composition of linear maps.

Physical intuition: Any element of $\mathcal{V}$ is a linear combination of basis vectors; this is most conveniently expressed as an n–tuple of $\mathbb{K}$–valued entries. A vector in three-dimensional space is represented by a real-valued column vector with three entries (in some coordinate system). Rotating the vector amounts to multiplying this column vector by some matrix to determine a set of rotated coordinates. So the action of an abstract group element is *represented* on the vector space by the action of a matrix on a column vector. The representation, T, is just the map from an abstract group element to a matrix that instantiates that transformation on an element of a vector space. What this means is that the subgroup of the group of linear maps is isomorphic to the abstract group, and in this sense, represent that elements of G.

A simple way, formally, of keeping track of dynamical symmetries is just to equip $\mathcal{V}$ with a representation of the dynamical symmetry group. This ensures that equivalence classes of elements of $\mathcal{V}$ under a symmetry transformation are treated as empirically indistinguishable.[5]

[5] In the rest of this thesis, until the final chapter, unless explicitly stated otherwise, our attention will be restricted to symmetries that form *Lie groups*, i.e. groups which are also smooth manifolds, and whose composition and inverse operations are smooth maps.

4.2.2 Yang–Mills theories on fibre bundles

In this subsection, I introduce two important mathematical objects—principal and associated bundles. I also demonstrate how so-called *gauge* fields (i.e. force fields) and ordinary matter fields are represented by different objects, with different transformation properties, that represent their different roles in encoding both relational and absolute properties of the field.

Definition: A *right G–action* is a map, $\triangleleft : B \times G \to B$, where B is a manifold and G a Lie group, such that $\forall p \in B$, (i) $p \triangleleft e = p$, where e is the identity element of G and (ii) $(p \triangleleft g_1) \triangleleft g_2 = p \triangleleft (g_1 \cdot g_2)$, where $g_1, g_2 \in G$ and $\cdot$ is the group action in G.

Definition: A right G–action is *free* iff $\forall g \in G$, $\forall p \in B$, if $p \triangleleft g = p$, then $g = e$.

Definition: The *orbit* of a point $p \in B$ under a right G–action is the set $\theta_p = \{q \in B | \exists g \in G : g \triangleleft p = q\}$

Definition: A *smooth bundle*, is a quadruple $\langle E, \pi, B, F\rangle$, where E, B and F are smooth manifolds, and $\pi : E \to B$ is a surjective map, the target of whose inverse is called the 'typical fibre', F, one copy of which is associated with each $p \in B$. Further, there exists a set $\{U_i\}$ of open sets in B, and diffeomorphisms, $\phi_i : U_i \times F \to \pi^{-1}(U_i)$ such that $\pi \circ \phi(p, f) = p$.

Physical intuition: Informally, a fibre bundle consists of a base manifold, B, (often, but not necessarily, taken to represent space or spacetime)[6] together with copies of some other manifold, called the fibres, F, associated

[6] In presenting the abstract setup, and discussing Wallace's fields-as-bodies picture later on, I use 'B' to represent the base/body manifold. This is to avoid the association that this manifold usually has with spacetime. When such an identification is required, as when discussing Weatherall's view, I switch to representing it as 'M', the more traditional way of representing the spacetime manifold.

with each point (sometimes an image of a fibre being 'glued on' to the base manifold is employed; I would like to avoid that since it leads to intuitions about some elements being 'closer' to the base manifold than others). A matter field, that might previously have been represented as a function on B, is now represented as a *section* of the bundle, i.e. a smoothly varying choice, at each point in B, of some single element of the fibre F above it. A scalar matter field is, therefore, a section of a bundle with scalar fibres—i.e. the fibres are copies of some algebraic field, usually $\mathbb{R}$.

So far, we have just made precise the intuition about associating a copy of the field-value space to each point in spacetime. Now we take a short detour to look a different kind of bundle, known as a principal G–bundle, in which the fibres are isomorphic to the group under which a theory is invariant. Although the motivation for introducing such an object might appear opaque at this point, we will soon discover the power of this construction.[7]

Definition: A *principal G-bundle*, (also known simply as a principal bundle), $\langle E, \pi, B\rangle$ is a smooth bundle where (i) E is a right-G–space (i.e. E is equipped with a right- G–action, $\triangleleft$) (ii) $\triangleleft$ is free and (iii) it is isomorphic as a bundle to $\langle E, \rho, E/G\rangle$, where $\forall \epsilon \in E$, $\rho(\epsilon) = [\epsilon]$, where $[\epsilon]$ is the equivalence class of elements of E under the relation 'lies in the same orbit as'. In this paper, I will represent a generic principal G–bundle as $G \to E \to^{\pi} B$.

There is a technical reason for imposing the condition of 'free action' on the principal bundle: it is only under a free action that the fibres in a principal bundle are isomorphic to the structure group G itself. There is a subtlety about what level of structure is preserved by this isomorphism. The fibres are

[7] All of the following definitions are based on Nakahara's [81, Ch. 9], with some notation changed to bring it in line with the conventions used in this thesis.

not isomorphic as groups to G—in particular, because they are right-action spaces, in which this action is free, there is no privileged 'identity' element. The isomorphism, then, is somewhere between a diffeomorphism and a Lie group isomorphism—it is an isomorphism of what is known as a *principal homogeneous space* or *torsor*.[8] At this point, one might wonder why we need a principal bundle. After all, it is clear that this is not the smooth bundle whose sections represent our field. The fibres are Lie groups—typically the wrong sort of space. The reason we define a principal bundle is that we are interested in imposing a representation of the group G on the fibres in the smooth bundle whose sections *do* represent our field. This bundle is known as the *associated vector bundle*. The associated vector bundle is just an example of a smooth bundle where the fibre, F is some vector space. A gauge theory can then be specified by defining a *connection* on the principal bundle, and using this to define a *covariant derivative* on the associated bundle. This is the covariant derivative familiar from standard presentations of Yang–Mills theories (see e.g. [81, p. 59]).

It is useful to keep in mind our goal of being able to compare properties (i.e. elements of the fibres in an associated bundle) of distinct points on B. In particular, we do not merely want to know if two points have the same property, but also if they differ, then we would like to know by how much. Maudlin's example of velocity comparison is useful here. Consider a velocity vector, v_p at some point $p \in B$. This is an element of the tangent space, T_pB at p. To determine whether it is the same as a velocity vector at some distant point, $q \in B$, we need to be able to *parallel transport* v_p to T_qB. Parallel transport is, intuitively, a way of sliding a vector along some path in such a way that, at

[8] I am grateful to Neil Dewar for clarifying this for me (twice).

every point, it 'points in the same direction as earlier.' In a generically curved spacetime, the same starting vector, v_p could be parallel transported from p to q along two different curves to two different vectors v_q and v'_q in T_qB. Since vector spaces have addition (therefore subtraction) structure, it is possible, not only to see if they coincide, but also to measure the amount by which v_q and v'_q fail to coincide.

Definition: A *pushforward* of a map, $f : M \to N$ is a linear map $f^* : T_pM \to T_{f(p)}N$ such that $X \mapsto f^*(X)$ as $f^*(X)g := X(g \circ f)$, where $X \in T_pM$ and g is a real-valued function on N.

Definition: A *connection* is an assignment, to every tangent space T_pE,[9] a vector subspace, H_pE, called the *horizontal subspace* of T_pE. This subspace has the following properties:

(1) $H_pE \oplus V_pE = T_pE$, where V_pE is the *vertical subspace* of T_pE, defined as the set of tangent vectors whose pushforward maps, generated by the projection map, π, maps them to the identity element of T_pB.

(2) $(\triangleleft g)_* H_pE = H_{p\triangleleft g}E$, where $(\triangleleft g)_*$ is the pushforward of the right action of the group element $g \in G$

(3) $\forall X_p \in T_pE$, $X_p = hor(X_p) + ver(X_p)$, where $hor(X_p) \in H_pE$ and $ver(X_p) \in V_pE$, this decomposition is unique, and, further, given a vector field X, the component vectors, $hor(X_p)$ and $ver(X_p)$ are part of the smooth vector fields, $hor(X)$ and $ver(X)$ respectively.

Definition: A *principal connection* is a connection defined on a principal bundle.

[9] Note that the spaces considered here are tangent spaces to the total space E, not just the base manifold B.

Definition: A *one-form*, ω, on a smooth manifold is smoothly varying choice of a cotangent vector at each point in the manifold.

Physical intuition: Very roughly speaking, a connection on a fibre bundle is a rule that identifies vectors in tangent spaces at 'neighbouring' points on the fibre bundle. This is achieved by identifying a subspace of the tangent space as 'vertical'—the vectors in this subspace are all pushed forward to the zero vector by the projection map, π. One can then define a horizontal subspace by choosing a subspace of T_pE whose elements, together with the elements of the vertical subspace, span the entire tangent space. There will be uncountably many choices available—a choice of a connection amounts to a choice of which subspace is horizontal, given a vertical subspace.

We are interested in constructing a parallel transport map for vectors defined on the base manifold, B. We begin with a vector, $v_p \in T_pB$. We then ask which vector in the tangent space, T_qB a neighbouring point, q counts as parallel. We can only make this judgement with respect to a curve connecting p and q. This curve, call it λ, will have an associated tangent vector field, T_λ. Each point, p' in the image of $\lambda(t)$ therefore has associated with it a privileged element of $T_{p'}B$, i.e. the tangent vector to $\lambda(t)|_{p'}$. What we need is a vector field along λ that includes v_p and identifies a privileged member of each tangent space at points along λ as 'being parallel' to v_p.

To get a sense of what is going on here, let us start with a simpler setup. Consider a curve, $\lambda : [0, 1] \to G$, where G is a Lie group. This curve will define a tangent vector field T_λ. It is easy to prove that the set of all (left-invariant) vector fields is a Lie algebra, with Lie bracket given by the commutator. Further, this Lie algebra is isomorphic to the tangent space at the identity, T_eG of the group. Associate, through this isomorphism, the tangent vector field, T_λ with a Lie algebra element. This is where the connection one-form comes in—it linearly maps the tangent vector field along λ to an element of T_eG (of course, it does more than this—it associates to each element of T_eG a vector field on all of G, not just a vector field along the image of a given curve). The connection one-form is chosen to have the following important property: $\omega_p(X_p^A) = A$, where A is an element of T_eG, and X^A is a vector field on G identified by the Lie-algebra isomorphism map between T_eG and the algebra of left-invariant vector fields on a manifold. We can now generalise this to vector fields on a principal bundle. Recall that a principal bundle consists of a collection of copies of such Lie groups (technically 'torsors'; see fn. 8). The vector fields defined on each fibre together define a vector field on all of the principal bundle.

For any curve between p and q, the Lie algebra-valued connection one-form associates the tangent vector at each point on the curve with some other element of the tangent space, T_eG, at the identity of the principal fibre, G, at that point. That Lie algebra element is associated with a vector field on the entire fibre, so, in particular with a specific element the tangent space at each point. The vector field generated in this way consists *only* of vectors that are parallel by the standards of that connection—call this the *parallel vector field* associated with a curve and a connection. It is easy to see that, in general, in addition to the choice of connection, the path-dependence affects the choice of the parallel vector field.

Thus, a connection determines a parallel transport rule which in general, depends on the path chosen between two points on the manifold. Given a path on the base manifold, B of the principal bundle, $G \to E \to^{\pi} \to B$, we can define a unique *horizontal lift* of the path into the total space, E. This is a particular kind of section on the principal bundle: it is a path in the total space whose tangent vector at each point lies in the horizontal subspace of each tangent space on the path. A horizontal section can only be specified on a principal bundle with a connection. If $\gamma : [0, 1] \to B$, where $\gamma(0) = a \in B$ and $\gamma(1) = b \in B$, is a curve on the base manifold, then its horizontal lift is a curve $\gamma^{\uparrow} : [0, 1] \to E$ where $\gamma^{\uparrow}(0) =: p \in \pi^{-1}(a)$, $\gamma^{\uparrow}(1) =: q \in \pi^{-1}(b)$ and $\forall t \in [0, 1]$, $\gamma^{\uparrow}(t) \in \pi^{-1}(\gamma(t))$.

This construction is important because it can be used to define a horizontal lift of paths in an associated bundle. This horizontal lift is uniquely determined by the principal connection, and conversely, the set of these horizontal lifts uniquely determines the principal connection.

Definition: A *left G–action* is a map, $\triangleright : G \times B \to B$, such that $\forall p \in B$, (i) $e \triangleright p = p$, and (ii) $g_1 \triangleright (g_2 \triangleright p) = (g_1 \cdot g_2) \triangleright p$.

Definition: An *associated bundle*, $\langle E_F, \pi_F, B\rangle$, to a principal G–bundle, $\langle E, \pi, B\rangle$ is the smooth bundle where F is a smooth manifold equipped with a left action, called the *associated fibre*, $\triangleright$, E_F is the quotient manifold $E \times F/\sim$ where $\sim$ is the relation on $E \times F$: $(p, f) \sim (p', f') \iff \exists g \in G : p' = p \triangleleft g, f' = g^{-1} \triangleright f$, and $\pi_F : E_F \to B$ maps $[p, f] \mapsto \pi(p)$. I will represent associated bundles later on as $F \to E_F \to^{\pi_F} B$.

Definition: Let $\gamma : [0, 1] \to B$ be a path in B, and $\gamma^{\uparrow} : [0, 1] \to E$ be its horizontal lift in the principal bundle, E. The *horizontal lift of γ in the associated bundle through* $[p, f]$ is the curve $\gamma^{\uparrow, E_F} : [0, 1] \to E_F$ where $\gamma^{\uparrow, E_F}(\lambda) := [\gamma^{\uparrow}(\lambda), f]$.

Physical intuition: If the manifold F is a vector space, then the construction of an associated bundle amounts to associating a copy of F, endowed with a representation of G, to each point on the base manifold, B. We can then define a matter field as a section of a vector bundle, and this matter field will be manifestly invariant with respect to G-transformations. In other words, if G is a proper subset of the space of linear maps, $GL(n, \mathbb{K})$, then the vector space, F will be endowed with further structure. For example, if $G = SO(1, 3)$, then F is a Minkowski vector space.

A horizontal lift of a curve in an associated bundle associates with each point in that path, an element of the fibre F. Of course, each path in B is still associated with a tangent vector field, and each tangent vector is a derivative that acts as an endomorphism on the space of $\mathbb{K}$–valued scalar functions on B.

Given a connection one-form on a principal bundle, we know that we can construct a unique horizontal lift (unique, that is, up to choice of starting point in the fibre) in the associated bundle, of a path in B. Consider a section, $\sigma : B \to E_F$, of the associated bundle. Let it map $p \in B$ to $\sigma(p) \in P_{F_p}$ and $q \in B$ to $\sigma(q) \in E_{F_q}$. We now parallel transport $\sigma(p)$ to E_{F_q} by pushing it forward along the horizontal lift in E_F. This will map $\sigma(p)$ to some vector $\sigma'(q)$ which is not necessarily the same as $\sigma(q)$. But now we can use the vector space structure of F_q to compare the values of $\sigma(q)$ and $\sigma'(q)$. If we parametrise our curve on the base manifold, B, using a parameter $\tau \in \mathbb{R}$, then we can define a *covariant derivative*:

$$D_{\pi_F^{-1}(q)}\sigma(q) = \lim_{\tau \to 0} \frac{\sigma(q) - \sigma'(q)}{\tau}$$

With respect to a particular choice of local section on the principal bundle, and a choice of local coordinates on the base manifold, the covariant derivative can be written as:

$$D_\mu = \partial_\mu + A_\mu$$

where A_μ, the *gauge field* is a connection one-form ω, on the base manifold of the principal bundle and ∂ is a flat derivative operator, whose form is also determined by the choice of section on the principal bundle.

This is the covariant derivative used to define all the dynamical equations of the field. The matter fields are represented by sections on the associated bundle, and the gauge field is a Lie algebra-valued connection one-form on the base manifold of the principal bundle. Such a theory is commonly known as a *Yang–Mills* theory. It is worth pointing out that Wallace's criteria, towards whose articulation we are working, require a further distinction between vector bundles whose fibres qualify as *internal* and those which do not. This requires an understanding of the notion of a vertical bundle automorphism (this particular statement is from Dewar [33]).

Definition: A *bundle automorphism* on a smooth bundle $\langle E, B, \pi \rangle$ is a pair of diffeomorphisms, (α, β) where $\alpha : E \to E$ and $\beta : B \to B$ such that $\pi \circ \alpha = \beta \circ \pi$

Defintion: A *vertical bundle automorphism* is a bundle automorphism of the form (α, Id_B) where Id_B is the identity map on the base manifold.

Definition: The fibres of an associated vector bundle are *internal* just in case the bundle has a non-trivial set of vertical automorphisms.

Physical intuition: Internal fibres are fibres in which a transformation preserves all dynamically relevant structure. A Minkowski vector space, for example, exhibits a $SO(1, 3)$ internal symmetry. The tangent bundle, on the other hand, is not a bundle whose fibres are internal. There exists a preferred identification between its elements and infinitesimal curves through the point on which it is defined. There is no non-trivial vertical automorphism on the bundle that does not break this preferred identification.

The final piece of geometrical machinery we need in this chapter is the structure required to express dynamics. We will not specify what the dynamical equations look like. For the arguments presented here, all we need is the ability to write down these equations, along with some facts about their symmetries; the dynamical equations themselves play no further role. We have already encountered the one-form—a linear map from tangent vector fields to real-valued functions, by mapping tangent vectors, at each point on the manifold, to some real number (or, more generally, some mathematical space). We can generalise this to an *n-form*:

Definition: A (differential) *n-form* on a manifold is a $(0, n)$ tensor field, ω that is totally antisymmetric.

Intuition: If we take the tensor product of a tangent space with itself, we get a new, larger vector space. It is possible, then, to find a vector space whose elements are maps from this tensor product of tangent spaces to the space of real numbers. This is just what an $\mathbb{R}$–valued $(0, 2)$ tensor is. An n–form is a special class of this type of tensor which is antisymmetric in its entries. That is to say:

$$\omega(X_1, ...X_n) = \text{sgn}(\Pi)\, \omega(X_{\Pi(1)}, ...X_{\Pi(n)})$$

where sgn(Π) is the sign associated with the permutation, Π, and $X_1...X_n$ are vector fields.

Definition: A *wedge product* is a map $\wedge : \Omega^n(M) \times \Omega^m(M) \to \Omega^{m+n}(M)$ such that

$$(\omega \wedge \sigma)(X_1...X_{n+m}) := \frac{1}{m!}\frac{1}{n!}\sum_{\Pi\in\text{Perm}} (m+n)\text{sgn}(\Pi)\omega \otimes \sigma(X_1...X_{n+m})$$

where $\Omega^n M$ is the space of n-forms on the manifold, M, $\omega \in \Omega^n(M)$, $\sigma \in \Omega^m(M)$ and $\otimes$ is the tensor product.

Intuition: Taking a tensor product between two $n-$forms does not guarantee that the result will also be a form—the property of antisymmetry might not be preserved. The wedge product is a particular kind of product that preserves antisymmetry.

Definition: An *exterior derivative* is an operator $d : \Omega^n(M) \to \Omega^{n+1}(M)$ such that:

$$d\omega(X_1, ..., X_{n+1}) := \sum_{i=1}^{n}(-1)^{i+1}X_i(\omega(X_1...\hat{X}_i, ...X_{n+1}) + \sum_{i<j}(-1)^{i+j}\omega([X_i, X_j], X_1..., \hat{X}_i, ..., \hat{X}_j, ..., X_{n+1})$$

where that hat over a vector field indicates that it is removed from the argument of the $n-$form ω, $[\cdot, \cdot]$ is a Lie bracket and Ω^n is the (vector) space of $n-$forms.

Intuition: The exterior derivative is an operation on forms that produces another form (ordinary differentiation is not guaranteed to do that).

Definition: A *pullback*, $f_* : V \to W$, is a linear map from the space of covectors, V, on a manifold, M to the space of covectors, W, on a manifold N, induced by a homomorphism $f : M \to N$ defined as:

$$(f_*\omega(X))_x = \omega_{f(x)}(df_x(X))$$

where ω is a covector, X is a tangent vector, $x \in M$ and d is the exterior derivative.

The principal bundle construction is useful in a number of ways, but its primary drawback is that it is very difficult to do actual physics with. To translate this picture into one that is closer to the actual practice of physics, we use a *gauge fixed* form of a Yang Mills theory.

One avoids the necessity of choosing a section and dealing with principal bundles at all by defining an isomorphism, $f_p : F_p \to F$, between each fibre of the associated bundle, F_p, and a fixed copy of that fibre, F (technically, this makes it a base-space indexed choice of automorphisms on the fibres). This automorphism is known as a *gauge*. Whereas on the bundle picture, a field, $\tilde{\phi}(x)$ was a section of a bundle, i.e. $\tilde{\phi}(p) : B \to F_p$, a field is now a globally-defined entity: $\phi(p) = f_p \cdot \tilde{\phi}(x) \in F$.

We went to the bundle picture in the first place because we realised that comparison of properties at distant points on the base manifold was not automatically guaranteed—this is why we developed the notion of a connection. So in returning to a global picture, we must realise that there is a free choice of gauge. So two distant points might have the same property as determined by one gauge but not another; this reflects what we had earlier noticed about the path and connection-dependence of 'sameness of property' ascriptions in bundles.

A choice of gauge can be given by a choice of (flat) connection, so we can define a covariant derivative as above. The gauge freedom under transformations in the symmetry group, G, of the theory translates to the following conditions on the matter fields, ψ, the gauge fields, A_μ and the covariant derivative, D:

$$\begin{aligned} \psi(x) &\mapsto g(x)\cdot\psi(x) \\ A_\mu(x) &\mapsto \mathrm{Ad}(g(x))A_\mu(x) - \partial_\mu(x)g^{-1}(x) \\ D_\mu\psi(x) &\mapsto g(x)\cdot D_\mu\psi(x) \end{aligned}$$

where Ad is the adjoint action of the Lie group G on its Lie algebra,[10] g is an element of G. In order to perform actual calculations, one might have to go further and make a choice of gauge and coordinates. On this set up, then, the kinematics and dynamics are (almost) in the form that one might encounter in a textbook on electromagnetism, for example.

Equipped with these geometrical notions, Wallace's two (equivalent) criteria for a theory counting as a Yang–Mills (or gauge) theory can now be presented:

YM$_1$: A Yang–Mills theory is a theory of a vector-valued function together with a gauge field which is a one-form that takes its values in the Lie algebra of the symmetry group of the theory.

YM$_2$: A Yang–Mills theory is a theory of a dynamical principal connection together with a section of an associated *internal* vector bundle,

where an internal vector bundle is, as defined above, a vector bundle with

[10] The adjoint action of a Lie group on a Lie algebra is the pushforward of the adjoint map on a Lie group, i.e. if $\mathrm{Ad}_g(h) = ghg^{-1}$, then the pushforward of this map, Ad_{g_*} is the adjoint action on the Lie algebra.

non-trivial vertical bundle automorphisms.[11]

Of course, these correspond to just one choice of definition of a Yang–Mills theory; Weatherall [122] has a much more permissive definition. He drops the restriction to *internal* associated vector bundles, as a result of which, for him, general relativity is a theory of a vector bundle associated to a principal-$GL(4, \mathbb{R})$ bundle; this is discussed in more detail in §4.4.1. For Weatherall, this is not a problem—the base manifold is a *spacetime* manifold, and as a proponent of the geometrical approach, he endorses the idea that spacetime kinematically privileges some dynamical structure; we discuss this in more detail in §4.4.1.

Wallace shows that, on the standard picture, the tetrad and the teleparallel formulations, general relativity is not a Yang–Mills theory. Here is the brief demonstration of why the standard formulation of general relativity fails his criteria $\mathbf{YM}_2$ (incidentally, this is the formulation of general relativity as a Yang–Mills theory that Weatherall favours): the dynamical objects in play are the metric field, g_{ab} and the covariant derivative built from the Levi-Civita connection, ∇_a. The former is not a section on an appropriate internal vector bundle, the latter is not an appropriate connection. The metric tensor is a section of the $(0, 2)-$tensor bundle, which is *not* an internal bundle (for reasons to do with the preferred identification between tangent space elements and curves through a point). This is not, however, particularly damaging—not all dynamical objects in a theory are matter fields. And, as a result of assumed metric compatibility, if the connection is of the appropriate form, then the Einstein field equations, are really equations governing the dynamical *connection*, in exactly the same way as in electromagnetism and other Yang–Mills theories.

So is the connection of the appropriate type? Things look promising to

[11] The question of whether these two definitions are truly equivalent is subtle and interesting, but not one with which I will engage in this chapter; for the rest of this chapter I focus on $\mathbf{YM}_2$.

begin with—it, too, is defined on a principal bundle of the appropriate type. This connection gives rise to a covariant derivative that acts on the associated bundle. However, in order that this covariant derivative has a direct link to the behaviour of scalar, vector and spinor fields in *spacetime*, this associated bundle has to be identified with the tangent bundle. And, as we have seen, the tangent bundle is not an internal bundle; thus we see a conflict with $\mathbf{YM}_2$.

Weatherall bites the bullet at this point, and argues that this is not a damning disanalogy. In effect, he is disagreeing with Wallace's $\mathbf{YM}_2$, in particular the requirement that a Yang–Mills theory requires the fibres of the associated bundle to be *internal*. If we stay on track with Wallace, and try to see if general relativity, formulated using tetrads, is any closer to being a Yang–Mills theory on Wallace's criteria, we discover that, although the tetrad formulation does not quite meet the criteria of $\mathbf{YM}_2$, its failure to do so is instructive in coming up with the required formulation. This is discussed in §4.2.3.

4.2.3 The tetrad formulation of general relativity

In general relativity, we deal exclusively with facts concerning tangent spaces (and linear functionals on products of tangent spaces) defined on Lorentzian manifolds. The entire empirical content of general relativity is presented in terms of vectorial and spinorial tensors, all of which are built up out of linear functionals on tangent spaces. Given a particular choice of coordinates on the manifold, one can make a natural choice of basis of the tangent space, known as the *coordinate basis*. Intuitively, this corresponds to the basis of the tangent space being composed out of directional derivatives that act along the coordinate axes at a given point. The use of coordinate bases is so common in general relativity as to be almost invisible. Indeed, in the preceding sections, the flat derivative operator, ∂ and the connection one-form, A_μ were both defined in the coordinate-induced bases of their respective spaces.

On a generic smooth manifold, the tangent spaces do not have enough structure to define a notion of orthonormality. However, a metric on the manifold provides the structure to define orthonormality. On (metrically) flat manifolds, one can define a chart on any open region such that the corresponding coordinate bases are orthonormal, but this is no longer true on curved manifolds. A *tetrad field* is a smooth assignment of four orthonormal basis vectors to the tangent spaces to four-dimensional metric manifolds. Of importance to this chapter is the fact that they make clear the connection between general relativity and Yang–Mills theories.

If we represent a coordinate basis vector in T_pB as $\hat{e}_\mu := \partial_\mu$, a generic vector, v, can be expressed as $v = v^\mu \hat{e}_\mu$. If we were to use different basis of the tangent space, say $\hat{e}_\alpha$, then the same vector would be expressed in this basis as $v = v^\alpha \hat{e}_\alpha$. The rule to transform between two tetrads is as one might imagine: $\hat{e}_\mu = e^\alpha_\mu \hat{e}_\alpha$. (Carroll [30, p. 484] draws attention to a sloppiness in accepted terminology, whereby both the orthonormal bases $\hat{e}_\mu$ and their components, e^α_μ in some other basis, $\hat{e}_\alpha$ are referred to as tetrads. Here, I refer to the former as tetrads, and the latter as tetrad components).

The components of the tangent vector in the different bases are then related as: $v^\alpha = e^\alpha_\mu v^\mu$. In this example, the tetrad component e^α_μ is merely imposing a change of basis on the same space. More generally, however, a tetrad can be seen as acting on a tangent space to a point on a non-metric manifold, and linearly mapping each element of that tangent space to a vector in a Minkowski vector space, $\mathbb{M}$. As such, it is formally a *Minkowski-vector space-valued one-form*— $e^\alpha{}_\mu : T_pB \to \mathbb{M}$.[12] It is useful to define an *inverse tetrad component*, $e_\alpha{}^\mu : \mathbb{M} \to T_pB$,

[12] In this rest of this thesis, to avoid confusion with abstract indices, tetrad indices will be

which satisfies $e^{\alpha}_{\mu} e_{\beta}^{\ \mu} = \delta^{\alpha}_{\beta}$.

The standard of orthonormality is determined by the metric, i.e. $g_{\mu\nu} = e^{\alpha}_{\mu} e^{\beta}_{\nu} \eta_{\alpha\beta}$, where $\eta_{\alpha\beta} = \text{diag}(-1, 1, 1, 1)$. In the case of a Lorentzian manifold, the symmetry group of the tangent space is the Lorentz group, $SO(1,3)$. So tetrads are related by the following transformation law:

$$\hat{e}_{\alpha'} = \Lambda^{\alpha}_{\ \alpha'} \hat{e}_{\alpha} \tag{4.1}$$

where $\Lambda \in SO(1,3)$. This means that we can build 'mixed index' tensor component equations. Consider, for example, a (2, 2) tensor of the form $T^{\alpha\mu}_{\ \ \ \beta\nu}$. Its transformation law would be:

$$T^{\alpha'\mu'}_{\ \ \ \ \beta'\nu'} = \Lambda^{\alpha'}_{\ \alpha} \frac{\partial x^{\mu'}}{\partial x^{\mu}} \Lambda^{\beta}_{\ \beta'} \frac{\partial x^{\nu}}{\partial x^{\nu'}} T^{\alpha\mu}_{\ \ \ \beta\nu} \tag{4.2}$$

The covariant derivative of general relativity, when expressed in some coordinate basis (and demonstrated here as acting on a vector, $v = v^{\alpha}\hat{e}_{\alpha}$) takes the form:

$$\nabla_{\mu} v^{\nu} = \partial_{\mu} v^{\nu} + \Gamma^{\nu}_{\sigma\mu} v^{\sigma} \tag{4.3}$$

where $\Gamma^{\nu}_{\sigma\mu}$ are known as the *connection coefficients*. This particular expression is tied to the use of a coordinate basis. In an orthonormal basis, the covariant derivative is expressed in terms of a *spin connection*, $\omega_{\mu\ \ \beta}^{\ \ \alpha}$. The same covariant derivative, now expressed in a mixed basis takes the following form:

$$\nabla_{\mu} v^{\alpha} = \partial_{\mu} v^{\alpha} + \omega_{\mu\ \ \beta}^{\ \ \alpha} v^{\beta} \tag{4.4}$$

The introduction of a tetrad field brings with it some new degrees of freedom.

represented by lower-case Greek letter from early in the alphabet.

As a result, there is more information that needs to be specified in order to make 'tetrad general relativity' equivalent to standard general relativity. In particular, the spin connection is uniquely determined by the condition that the following equation holds:[13]

$$de^{\alpha} + \omega^{\alpha}_{\ \beta} \wedge e^{\beta} = 0 \tag{4.5}$$

where d is the *exterior derivative* and $\wedge$ is the *wedge product*.

Tetrad general relativity does come a little closer to the form of a Yang–Mills theory as specified in §4.2.2. The spin connection and the tetrad components both have exactly the right sort of transformation properties to be considered as a Yang–Mills connection and matter field, respectively.[14] However, the tetrad is not the right sort of object—it is a one-form, not a section of an associated bundle.

But, *qua* connection, one might wonder whether the tetrad is the right sort of connection. Perhaps the tetrad and the spin connection are both *components* of a single connection one-form on a principal bundle whose structure group G includes the Lorentz group as a proper subgroup? The obvious move is to interpret it as the connection of the translation component of the Poincaré group. This turns out to be the solution: if it is the translation-invariance that is scuppering our attempt to see general relativity as a theory of a connection on a principal Poincaré bundle because the tangent bundle is not an appropriately associated bundle, then let us change the fibre in our associated bundle to one

[13] In §4.3.2, we will see that this equation also expresses the fact that the tetrad can also be interpreted as the translation component of a Poincaré connection.
[14] They transform as follows:

$$e^{\alpha}_{\mu}(x) \to \Lambda^{\alpha}_{\beta}(x) e^{\beta}_{\mu}(x)$$

$$\omega^{\alpha}_{\beta\mu}(x) \to \Lambda^{\alpha}_{\gamma}(x)\omega^{\gamma}_{\delta\mu}(x)(\Lambda^{-1})^{\delta}_{\beta}(x) - \partial_{\mu}\Lambda^{\alpha}_{\beta}(x)$$

this *is* invariant under translations: an affine space.

4.3 Fields as bodies

This section introduces Wallace's construction, pausing to spell out some details that will be illuminated by the material presented in the preceding subsections. In §4.3.1, I discuss the most conceptually significant step—the introduction of a *location* field and demonstrate how the associated connection is a Yang–Mills connection in the sense of §4.2.2. Finally, §4.3.2 demonstrates how the kinematical structure of vacuum general relativity is an instance of the kinematical structure of a Yang–Mills theory.

4.3.1 Location fields

The most important move, conceptually speaking, is not to give in to the intuition that the base manifold represents spacetime (specifically locations in spacetime). To allow us to model spacetime locations in terms of the fibres, we need to introduce the concept of an *affine space*.

Definition: An *affine space* is a triple $\langle \mathcal{X}, \mathbb{V}^n, + \rangle$, where $\mathbb{V}^n$ is an n-dimensional vector space, $\mathcal{X}$ is a set, and $+ : \mathcal{X} \times \mathbb{V}^n \to \mathcal{X}$ is an action that satisfies the following two criteria: (i) $\forall a, b \in \mathcal{X}, \exists \vec{c} \in \mathbb{V}^n : a + \vec{c} = b$ and (ii) if $\exists a \in \mathcal{X} : a + \vec{c} = a + \vec{d}$, then $\vec{c} = \vec{d}$,, where $\vec{c}, \vec{d} \in \mathbb{V}^n$.

Physical intuition: An affine space is a vector space that has 'forgotten its origin.' Given any two points, one can define a vector between them. As a space of vectors (though, of course, not a *vector space*), its elements can be acted on by (images of) representations of group elements. One can therefore construct an associated bundle out of affine fibres, and all of the features of a Yang–Mills theory can continue to be expressed.

We are used to thinking of the Poincaré group as a 'spacetime' symmetry group. Recall the common intuition that spacetime symmetries are related to transformations on the base space of a bundle. This is precisely the source of the tension in the $\mathbf{YM}_2$ interpretation of general relativity—the tangent bundle is not an internal bundle, and translations are not linear transformations. The parametrised field theory formalism indicates how this impasse could be resolved—move the 'spatiotemporal' degrees of freedom into an internal bundle. We therefore interpret the base manifold, B, as representing an extended body whose spatiotemporal and non-spatiotemporal properties are represented by sections of a fibre bundle. Note that this means it is still eligible to be called Yang–Mills theory according to $\mathbf{YM}_2$, which does not restrict fibres to being vector spaces; all it required was that it was the sort of space that could carry a representation of a Lie group. The sort of fibre that carries a representation of the Poincaré group is an *affine Minkowski* space, AM i.e. an affine space equipped with a Minkowski metric.

Briefly returning to a global (i.e. non-bundle) picture of a field theory, we note that a standard picture of a field, $\tilde{\phi}(x)$ in affine Minkowski spacetime, $\tilde{\phi}(x) : AM \to \mathbb{V}$, with Lagrangian (density) $\mathcal{L}(\phi, \partial_\mu \phi)$ could be thought of as being defined by $\phi : B \to AM \oplus \mathbb{V}$, where $\oplus$ is the *direct sum* operation.[15] A field, ϕ is now *both* an assignment of an element, $\rho^\alpha \in AM$ (a spatiotemporal location) and an element $\psi \in \mathbb{V}$ (a field value) for all $x \in B$. We interpret the ρ^α as a *location*

[15] A direct sum is a way of 'joining' two vector spaces together in such a way that linear structure is preserved in the larger space. More precisely, if V and W are two vector spaces, then $V \oplus W$ is constructed out of the Cartesian product, $V \times W$: if $v_1, v_2 \in V$ and $w_1, w_2 \in W$, then $(v_1, w_1), (v_2, w_2) \in V \times W$. The operations of addition and scalar multiplication are defined pointwise: (i) $(v_1, w_1) + (v_2, w_2) = (v_1 + v_2, w_1 + w_2)$ and (ii)$\alpha(v, w) = (\alpha v, \alpha w)$, where $\alpha \in \mathbb{K}$.

field which assigns a spatiotemporal location to each part of B. The Lagrangian now takes the form $\mathcal{L}(\rho, \partial_\mu\rho, \psi, \partial_\mu\psi)$. Note that the indices $\mu, \nu...$ continue to label coordinates on the base manifold, but no longer have an interpretation as spacetime indices. The spacetime indices are $\alpha, \beta, ...$ which index the affine Minkowski fibre.

The affine Minkowski space admits an orthonormal basis—an inverse tetrad, e_α, in virtue of its Minkowski metric structure. The dynamics (i.e. the Lagrangian, and associated Euler-Lagrange equations) needs to be expressed in terms of spatiotemporal quantities. By inspection, $\partial_\mu\rho^\alpha = e_\mu{}^\alpha$. We can therefore express the Lagrangian as follows:

$$\mathcal{L}(\rho, \partial_\mu\rho, \psi, \partial_\mu\psi) = \mathcal{L}(\psi, e^\mu_\alpha\partial_\mu\psi)\det(e^\alpha_\mu) \tag{4.6}$$

4.3.2 General relativity as a Yang–Mills theory

The fibre bundle formalism of Yang–Mills theories focusses on the dynamical connection on the principal bundle. But general relativity is, in its standard presentation, at least, a theory about a dynamical metric field. So construing general relativity as a Yang–Mills theory shifts the focus away from the metric, and onto the Levi-Civita connection. This is, mathematically, not a problem at all; the metric compatibility condition is implicit in the Einstein field equations (in other words, the derivative operator that ensures satisfaction of the Bianchi identities is the metric-compatible Levi-Civita derivative operator). So metric and affine geodesics coincide. One might even go so far as to argue that the affine structure is more fundamental, since vacuum general relativity (i.e. solutions where the stress-energy tensor vanishes) is describable as a pure affine theory, without using a metric at all.[16]

In this chapter, following both Weatherall and Wallace, I discuss only the kinematical structure of Yang–Mills theories. That is to say, a theory is a Yang–Mills theory if its kinematical structure is sufficiently similar to the standard set by $\mathbf{YM}_2$; there is no restriction on the dynamics that the connection ought to satisfy. Therefore, to see that the fields-as-bodies setup adequately describes general relativity, it is sufficient to demonstrate that one has enough structure to be able to write down the dynamics of the Levi-Civita connection. In particular, I shall demonstrate that the Yang–Mills curvature two-form of the rotational part of the connection on the principal Poincaré bundle is the Riemann curvature usually associated with the metric g_{ab}.

Consider a connection, ω^α on a principal Poincaré bundle. This induces a covariant derivative on the associated affine bundle, which allows us to compare the values of the location field, ρ^α at two neighbouring body manifold points. With respect to a choice of gauge (i.e. a smooth assignment of a map from each affine Minkowski fibre to a single copy of affine Minkowski space), the covariant derivative induced by the connection is:

$$\nabla_\mu \rho^\alpha = \partial_\mu \rho^\alpha + (A_\mu \cdot \rho)^\alpha \tag{4.7}$$

where A_μ is the Poincaré-algebra valued connection one-form. If we choose an origin, this connection can be broken up into a rotational component, $\omega^\alpha_{\mu\beta}$ and a translational component, θ^α.

One can then take the *covariant exterior derivative* of this one-form, defined

[16] This is because in a vacuum spacetime the Ricci tensor vanishes, i.e. $R_{ab} = 0$, and this only requires an affine connection in order to be defined. Of course, this comes with its own problems; for example it is not clear that such a theory is truly relativistic, given that in the absence of a metric, there is no distinction between temporal and spatial dimensions. This is touched upon in [97, §5].

as follows:

Definition: A *covariant exterior derivative* of an H-valued k–form, ϕ on a principal bundle with connection one-form, ω is a map $D\phi(X_i, ...X_k) := d\phi\big(hor(X_1), ..., hor(X_k)\big)$,

where d is the exterior derivative, $X_1, ..., X_k$ are vector fields, $hor(X_i)$ is the horizontal component of the ith vector field and H is some appropriate target space for the k–form map.

Recall that the choice of horizontal subspace of a tangent space was enforced by the connection one-form, ω. Thus a covariant exterior derivative can only be defined after a choice of connection has been made. It takes a one-form, valued in some space F, and outputs a two-form valued in the same space. In particular, the covariant exterior derivative of the Lie algebra-valued connection one-form itself is a very important object—it is referred to as the *curvature*, $\mathcal{R}$, of the connection.

In general, the covariant exterior derivative of a generic one-form, ϕ is given by the following formula, which just re-expresses the above definition:

$$D\phi^a_b = d\phi^a_b + \omega^a_c \wedge \phi^c_b \tag{4.8}$$

The rotational component of the Poincaré connection is a Lorentz algebra-valued one-form, ω^a_b. Taking the covariant exterior derivative of this one-form with respect to itself, one gets an expression for curvature:

$$D\omega^\alpha_\beta =: \mathcal{R}^\alpha_\beta = d\omega^\alpha_\beta + \omega^\alpha_\gamma \wedge \omega^\gamma_\beta \tag{4.9}$$

The Levi-Civita connection is the Lorentz algebra-valued connection on the associated bundle induced by a connection on the spin frame bundle (i.e. the principal $SO(1,3)$– bundle). In other words, it is the connection induced by the rotation component of the Poincaré connection: the expression for the curvature given in equation (4.9) just is the Riemann curvature tensor of general relativity. So the fields-as-bodies approach does reproduce the kinematics of general relativity.[17]

4.4 The dynamical–geometrical dispute

There is a tradition in the literature in the foundations of spacetime of ascribing to spacetime geometry the status of being an explanans of certain facts about the behaviour of physical fields, in virtue of these fields being associated, in some sense, with a spacetime geometry. The disagreement between the dynamical and geometrical camp is over what counts as an explanation of this behaviour.

The issue of the arrow of explanation is tricky, and was discussed at some length, in the introduction. For now, it is worth restating the three distinct roles that a metric, a connection and perhaps a more general class of geometric objects can play in a physical theory:

A: a role in *articulating* the dynamics of matter fields.

B: a role in articulating the dynamics of matter fields and, in addition, ensuring that symmetries of the matter fields are also symmetries of the metric *and vice versa.*

C: a role in articulating the dynamics of matter fields and, in addition, ensuring that symmetries of the matter fields are also symmetries of the metric

[17] In fact it does a whole lot more, reproducing teleparallel gravity and tetrad general relativity; for details of this, see [119, §V].

and vice versa and, further, ensuring that dynamical fields survey the associated metric field globally.

A matter field or particle 'surveys a metric' just in case it can be used to read off intervals of proper time, as given by that metric, along its worldline (which need not be a geodesic). This can be achieved by reference to, for example, some periodic process instantiated by a rigid system (with an appropriate stipulation of 'rigidity').[18]

One can situate the dynamical and geometrical views at different points on the spectrum defined by the above classification of roles. **A** is just a standard component of physical theories; both the dynamical and geometrical approaches accept it unproblematically. I take the proponent of the dynamical approach to be committed only to **A**, while the geometrical approach as espoused by Weatherall and Maudlin, is committed to **C**.

This section has two related aims. The first is to discuss and rebut an argument from Weatherall, about the interpretation of gauge symmetries, which is based on his version of the geometrical approach (which I take to entail that the Levi-Civita connection plays all three of the above roles). Based on this discussion, I intend to make precise the dispute between the dynamical and geometrical views, in terms of the fibre bundle geometry introduced in this chapter.

Weatherall's own take on the geometrical approach is helpfully expressed in [122], in which he argues that the philosophical significance of expressing general relativity as a Yang–Mills theory is that we can take our cues from general relativity when trying to interpret Yang–Mills theories, rather than the other way round (as is more commonly argued, in e.g. [55]). He argues that a gauge

[18] These ideas were discussed at length in Chapter 1, in particular in §1.2.1 of that chapter.

transformation, ordinarily thought of as a transformation induced by a shift in the principal bundle, should instead be thought of merely as a change of basis in the fibres of the associated vector bundle.

I disagree with this for two reasons. First, it is not clear that Weatherall's, as opposed to Wallace's, construal of general relativity as a Yang–Mills theory is appropriate. Consequently, since this option is not straightforwardly available for Wallace's approach (affine spaces do not have bases in the same way as vector spaces do), it is not obvious that this is an acceptable interpretation of gauge.

Second, and more importantly, the disanalogy between general relativity as a Yang–Mills theory in Weatherall's sense, and a standard Yang–Mills theory, far from being innocuous, actually sinks Weatherall's deflationary claim about gauge symmetries. In §4.4.1 I state, in detail, Weatherall's preferred formulation of general relativity as a Yang–Mills theory. I argue that, while this option is available *in general relativity* to the proponent of the geometric approach, it is not appropriate to the dynamical approach. I compare Weatherall's and Wallace's approaches in §4.4.2, and present a novel restatement of the dynamical–geometrical dispute in terms of fibre bundles.

4.4.1 Weatherall's version of Yang–Mills general relativity

In this section, I set out Weatherall's proposal in detail, using the machinery introduced in §4.2. Tangent spaces are supposed to be spaces of velocity vectors of particles and field locations, so play a central role in ascribing proper times along curves. In particular, if the base manifold is taken to represent (part of) spacetime, then it is only tangent spaces that have a link to spatiotemporal distances and velocities. With this in mind, in this section, I switch to representing the base manifold as 'M', the more standard notation for 'spacetime manifold'. In fact, the tangent space requires no more structure to be defined than the smooth structure of a manifold. A tangent vector is just an equivalence class of curves passing through a point on the manifold; it inherits its linear structure from the linear structure of the space of curves.

The tangent space (to a point on a four-dimensional manifold) is often thought of as a fibre of a bundle associated to the principal $GL(4, \mathbb{R})$ bundle, also known as a *frame bundle*.[19] The typical fibre of a frame bundle is the group of transformations of bases of the tangent spaces. The frame bundle is the appropriate bundle to encode these transformations, because of the existence of a *canonical solder form* (at least this is how the standard story goes; in §4.4.1, I define the term and discuss Weatherall's response to this claim).

One of the most obvious differences between an affine connection and a generic Yang–Mills connection is that a torsion can be defined very naturally on the former. Torsion is related to the lack of symmetry in the lower indices of the connection coefficients, $(\Gamma_\mu)^\alpha_\beta$. But these two indices, in a generic theory, index different spaces—μ is a one-form coordinate basis index, while α and β are Lie-algebra indices. How can we speak intelligibly of their being symmetric? The answer is that as coordinates on the associated bundle to the frame bundle, they happen to index the same space, so can be symmetrised. To see why this is, consider how one might construct the associated bundle to a principal frame bundle.

Recall from the definition of the associated frame bundle that a typical fibre

[19] The term 'frame bundle' (here, we restrict ourselves to four-dimensional base manifolds) refers to the bundle, LM, of bases to the tangent spaces to points on the manifold, $L_xM; x \in M$. More precisely, $L_xM = \{(e_1, ..., e_4) \mid ((e_1, ..., e_4)$ is a basis of $T_xM\}$. This space is isomorphic as a vector space (and, being finite dimensional, also isomorphic as a manifold) to the space of real-valued 4×4 matrices, i.e. $GL(4, \mathbb{R})$.

is induced by the left-action of the structure group of the principal bundle. So if one begins with a principal frame bundle, then the fibre of the associated bundle is constructed out of the quotient manifold $E\times F/\sim$, where F is a manifold equipped with a left action. More concretely, consider the right action of $GL(4, \mathbb{R}$ on the frame bundle:

$$\triangleleft : GL(4, \mathbb{R}) \times LM \rightarrow LM$$

$$(e_1 ... e_4) \triangleleft g = (g_1^m e_m, ... g_4^m e_m) \tag{4.10}$$

where m is the dimension of the base manifold. Now, to construct the associated vector bundle, we begin with a vector space, F. We need to consider the left action of $GL(4, \mathbb{R})$ on F:

$$\triangleright : GL(4, \mathbb{R}) \times F \rightarrow F$$

An element $f \in F$ is just a quadruple of real numbers. The left action of g on f gives us:

$$(g \triangleright f)^a = (g^{-1})^a_b f^b \tag{4.11}$$

We have thus constructed an associated bundle, $\langle LM_{\mathbb{R}^4}, B, \pi_{\mathbb{R}^4}\rangle$, to the frame bundle. The connection will take its values in the Lie algebra $gl(4, \mathbb{R})$, and because the base manifold is four dimensional, all real-valued indices, whether one-form coordinate basis indices, or Lie-algebra indices, run over the same values. So, to reiterate, the connection $(\Gamma_\mu)^\alpha_\beta$ is such that μ, α and β all index the same space. So it makes sense to define a term like $(\Gamma_\mu)^\alpha_\beta + (\Gamma_\beta)^\alpha_\mu$. This is an important part of being able to define torsion. However, we are not yet done.

Torsion is a property of a spacetime connection, and all we have constructed is an associated bundle whose fibres are real vector spaces. To have a link to spacetime, these fibres must be *tangent* spaces.

It turns out that the associated bundle that we just constructed is isomorphic, as an associated bundle, to the tangent bundle.[20] Moreover, any associated bundle constructed similarly from p copies of $\mathbb{R}^4$ and q copies of its dual, will be isomorphic to the (p, q) tensor bundle over B. This isomorphism comes for free—we did not have to specify any further structure on our frame bundle in order to get out the tangent bundle.

Weatherall's view is that the principal bundle, and the connection(s) defined on it, serve one overarching purpose: to '[coordinate] derivative operators acting on different, but systematically related, vector bundles associated with different kinds of matter influenced by the same forces.' [122, p.2391]. By being receptive, in theory, to the possibility that these derivative operators might be uncoordinated, Weatherall's view concedes something significant to the proponent of the dynamical approach.

If, as is standard, we were to interpret Weatherall's view from the perspective of a principal bundle formulation of a Yang–Mills theory, then general relativity as a Yang–Mills theory is just the theory described above—a theory of a frame bundle whose associated bundle is isomorphic, as an associated bundle, to the tangent bundle. However, to look at general relativity from the perspective of principal bundles is to miss the point of Weatherall's argument.

Weatherall argues that all Yang–Mills theories can be seen as theories of covariant derivatives on vector bundles whose fibres represent some fields (e.g. matter fields, spacetime vector, spinor or tensor fields). Mathematically speak-

[20] This is discussed in more detail in §4.4.1.

ing, it is, of course, true that these covariant derivatives are associated with a connection on a principal bundle. But this, says Weatherall, overstates the relative importance of the principal bundle, as is demonstrated by looking at general relativity. From the principal bundle perspective, one could see this structure as associated with some structure on the principal (frame) bundle, specifically, a choice of section and a principal connection. A choice of section, σ, is a choice of basis to the tangent space at every point on M, and the pull-back of the connection, ω, along this section, σ, is the connection coefficient associated with that frame. If the frame field is a tangent space basis built out of a choice of local coordinates on some open set $U \in M$ (this condition is known as *holomorphicity*), then these just *are* the connection coefficients of the Levi-Civita connection, realised as a covariant derivative operator on the tangent bundle. So why not just start with the Levi-Civita connection and bypass the entire principal bundle setup?

This is precisely what Weatherall does:

> Given the status of frame fields and Christoffel symbols on modern geometrical approaches to general relativity, the remarks above suggest that [the frame bundle] and the structure defined on it are really auxiliary. In other words, there is a reason that one need not mention a principal bundle or principal connection for most purposes in general relativity which is that it is the induced covariant derivative operator on M that matters... The frame bundle merely provides an alternative... way of encoding information about this derivative operator. [122, p. 2400]

This passage is helpful in understanding Weatherall's 'geometrical approach', especially when seen in light of the following quote taken from a later paper of his, on the dynamical approach:

> To posit a given geometry for spacetime is to limit the resources available for expressing dynamical laws; whether a given geometry is suitable for expressing the actual laws is a highly contingent matter. [123, p. 14]

The fact that the geometry in question in general relativity is a *spacetime* geometry is what makes the difference. By positing a particular spacetime geometry, in this case, a Lorentzian metric geometry, the claim that the proponent of the geometrical approach is pushing, says Weatherall, is that no further geometrical resources are required to state the laws for *all* dynamical fields in the theory. In a sense this is not a particularly bold claim, at least in light of role **B** for the Levi-Civita connection, introduced above. It does, however, require some extra-theoretic reason for positing some geometrical structure for spacetime; if not, then it just collapses into the dynamical approach, which says that the geometrical structure of spacetime just codifies the local symmetry structure of matter field dynamics. In this section, I demonstrate that Weatherall's deflationary gauge argument only works if we accept **B**, as it relates to *all* Yang–Mills connections. I argue that the construction of the relevant associated bundle isomorphism mandates that we should not accept **B**, though we may accept **A**, as it applies to generic Yang–Mills connections.[21] Thus Weatherall's argument rests on a false premise.

As a consequence of Weatherall's view of spacetime geometry, if one were to find a dynamical field that required, in addition to a particular Lorentzian metric field, some other geometrical objects, then, by modus tollens, the Lorentzian

[21] I also think that this is true of the Levi-Civita connection, but will not argue for that in this chapter.

metric geometry would not qualify as a spacetime geometry. This is particularly significant in the context of the Gödel solution discussed in chapter 1; on Weatherall's understanding, that solution would not count as a spacetime. I will have more to say about this in §4.4.1. Before that, I would like to demonstrate why Weatherall's deflationary story does not work, even on the assumption that his is the appropriate way to view general relativity as a Yang–Mills theory.

Spacetime vector fields play a very different role in physical theories than internal vector fields—the former are always taken to be sections of a tangent bundle, while the latter are sections of some internal vector bundle. The difference between a tangent bundle and an internal vector bundle, as we discussed in §4.3.2, is that the elements of the former are identified with infinitesimal paths through the manifold. These infinitesimal paths can also be thought of as determining directional derivatives on the manifold. What are these directional derivatives of? The answer is simple: other fields. Spacetime vector (and other tensor) fields are therefore not (necessarily) matter fields themselves, rather they determine velocities of trajectories of material particles or wave packets of matter fields. In general relativity, all spacetime derivatives are the same, and determined by the Levi-Civita connection. This fact—that the derivative operator is a spacetime derivative, and, from the quote above, positing a spacetime geometry limits the resources for expressing dynamical laws—means that the Levi-Civita connection, by definition, coordinates the behaviour of distinct matter fields: i.e. it determines the geodesic behaviour for all of them.

What if we had a different, non-spatiotemporal, derivative operator? For example, the gauge covariant derivative operator of electromagnetism. Weatherall suggests that the first step is to think of the structure groups of the principal bundles as subgroups of the linear group. So far, this is not a problem; it is a mathematical fact that $SO(3)$ is a subgroup of $GL(n, \mathbb{R})$, where $n \geq 3$. Just as the tangent bundle associated to a spin frame bundle (i.e. a principal $SO(1, 3)$ bundle) has extra structure (i.e. a Minkowski inner-product), the extra structure of, say, an internal $SU(2)$ bundle is encoded in the fact that the principal bundle to which it is associated is a subbundle of a $GL(n, \mathbb{C})$ bundle. The analogy that Weatherall exploits, then, is between the principal bundle to which a general relativistic tangent bundle is associated—a frame bundle—and the principal bundles to which other Yang–Mills theories' matter field vector bundles are associated—*subbundles* of the (possibly complex) frame bundle. Now, the important question is the following: how do we ensure that this coordinates different charged matter fields? Weatherall offers the following:

> [T]hat electrons, muons, etc. are all represented by vector bundles associated to the *same* $U(1)$-principal bundle, and have covariant derivative operators induced by the same principal connection, provides the precise sense in which these different kinds of particles all respond to the same electromagnetic influences. In effect, the principal bundles in Yang–Mills theory coordinate frames [i.e. bases] across different vector bundles. [122, pp. 2404-2405]

This leads to an immediate concern: grant, for the sake of argument, that the principal bundles are not needed to coordinate the action of the Levi-Civita derivative operator because it is a *spacetime* derivative operator. Surely the fact that different matter fields respond to the same electromagnetic field is evidence for the necessity of the principal bundle for those Yang–Mills theories. Weatherall acknowledges the disanalogy, but strives to underplay it:

> One might worry that the role that I have just ascribed to the principal bundles... is robust enough that it is misleading to call it "auxiliary",

> since being "auxiliary" may suggest that the structure is eliminable. In any case, I hope that I have been clear enough above about what I take the roles of various bundles to be...that the sense of "auxiliary" I have in mind is clear. It is the sense in which a coach is auxiliary to the players on the field. [122, fn. 34]

I do not believe that this does enough. In the case of general relativity, the principal bundle was *dispensable*—it gave us no information that was not already encoded in our choice of Levi-Civita derivative operator as a spacetime derivative operator. In the case of electromagnetism, the principal bundle is *indispensable*. To make this point more vivid, it is necessary to delve briefly into the technicalities of the setup. The point that I intend to spell out in the rest of this section can be summarised as follows: the tangent bundle, TM is isomorphic, as an associated bundle, to the associated vector bundle, $F_{GL(4,\mathbb{R})} \to E_F \to^{\pi_F} M$, to the frame bundle, LM only because the solder form defined on LM is preserved by a principal bundle automorphism that generates the required associated bundle isomorphism between TM and $F_{GL(4,\mathbb{R})} \to E_F \to^{\pi_F} M$. Crucially, all associated bundle isomorphisms are generated by a principal bundle automorphism (this will be discussed in §4.4.1). In the absence of a solder form on the frame bundle, it would not be possible to identify which principal bundle automorphism to use to generate the associated bundle isomorphism. All of this information is already encoded in a choice of Levi-Civita derivative operator; for a normal gauge covariant derivative operator, this needs to be specified, and it cannot be done systematically without a principal bundle. In other words, any systematic account of the 'coordination' of matter fields to a Yang–Mills covariant derivative operator is tantamount to a specification of a principal bundle automorphism; the principal bundle is equally indispensable in both general relativity as well as Yang–Mills theories.

Coordination and bundle morphisms

The tangent bundle is *an* associated bundle to the frame bundle. Recall, from the definition, that there are several bundles that can be associated with a principal bundle. It is important to note that one does not directly define the tangent bundle as the associated bundle to the frame bundle. Instead, one constructs the appropriate vector bundle, $F \to E_F \to^{\pi_F} M$, and then demonstrates that it is isomorphic, as an associated bundle, to the tangent bundle. This last clarification is important; the two bundles could be isomorphic at various levels of structure— manifold, smooth bundle or associated bundle. It is only in virtue of this *associated* bundle isomorphism that one can refer to the tangent bundle as the associated bundle to the frame bundle.

Defintion: If two Lie groups G and G' act on manifolds M and M' respectively, and with a homomorphism $\rho : G \to G'$. then a map $f : M \to M'$ is said to be ρ–*equivariant* if $\rho(g)f(p) = f(gp)$.[22]

Definition: A *bundle map*, (u, h) between two bundles, $E \to^{\pi} M$ and $E' \to^{\pi'} M'$, is a pair of smooth maps, $h : M \to M'$, $u : E \to E'$, such that $u \circ \pi = \pi' \circ h$.

Definition: A *principal bundle map*, (u, h) between two principal G–bundles, $G \to E \to^{\pi}\to M$ and $G' \to E' \to^{\pi'}\to M'$ is a bundle map which, in addition, is ρ–equivariant on the fibres.

Definition: An *associated bundle map*, (u', h') between two associated bundles,$F \to E_F \to^{\pi}\to M$ and $F' \to E_{F'} \to^{\pi'}\to M'$, which may be associated to different

[22] This particular definition is taken from Isham [61, pp. 178-179].

principal bundles, is a bundle map constructed from a principal bundle map, (u, h) between the underlying principal bundles as $u'([p, f]) = [u(p), f]$, where $p \in M$ and $f \in F$.

An associated bundle isomorphism, then, can only be constructed from a principal bundle isomorphism. Let us see how this works in the case of a frame bundle. To do so, we need one further piece of structure, that will also be important in setting up the distinction between the dynamical and geometrical approaches—the *solder form*:

Definition: A solder form, θ, on a principal bundle $G \to E \to^{\pi} M$, is a $\mathbb{V}$-valued one-form, where $\mathbb{V}$ is a linear representation space of the structure group, G, of the principal bundle and, (i) $\theta(ver(X)) = 0$, where $ver(X)$ is the vertical subspace of the tangent bundle to M; (ii) $g \triangleright (\triangleleft g)^{*}\theta = \theta$ and (iii) $E_V \simeq TM$.

From the definition, it is clear what the role of the solder form is—to provide the isomorphism between the tangent bundle and the associated vector bundle to the frame bundle. This object is often taken to be canonically defined, and, indeed, it almost is.

Consider a solder form, θ on the frame bundle, LM. This is a map from the space of tangent vector fields on E to some vector space, $\mathbb{V}$. Let $\mathbb{V} = \mathbb{R}^4$. Define the solder form:

$$\theta_e(X) := (u_e^{-1} \circ \pi_*)(X) \tag{4.12}$$

where $u_e : \mathbb{R}^4 \to T_{\pi(e)}M$, where $u_e(0, 0, 1...) \mapsto e_i$, $(0, 0, 1...)$ is a basis vector of $\mathbb{R}^4$ with a 1 in the ith entry, e is a frame in the fibre at a point $\pi(e)$ of the base manifold M and $X \in T_{\pi(e)}M$. Note the pushforward of the tangent vector from the total space, E to the base manifold, M.

This solder form is canonical in the sense that no further structure needs to be added to the bundle in order to define it. However, as Weatherall points out, that the solder form 'is only canonical relative to a number of prior choices, including a choice of basis for V [the fibre of the associated vector bundle] and an identification of points on LM (which, after all, is just some manifold) with ordered n tuples of vectors in TM' [122, p. 2409]. This is undeniably true, and another instance of a particular choice of conventions being so natural as to be invisible. He goes on, somewhat polemically, to say 'there is a sense in which one has solder forms on *all* frame bundles, not just frame bundles associated to the tangent bundle'.[122, p. 2410] Of course there is the obvious rebuttal: *those isomorphisms are not solder forms!* A solder form generates a bijection between the frame bundle and the *tangent* bundle, not the frame bundle and some vector bundle—the latter object isn't even a one-form (it has the wrong domain).

Quibbles about nomenclature aside, Weatherall's point, about the (actual) solder form not being canonical without further qualifications is an important one (I will return to this in §4.4.1). Equally important is his point regarding associated bundle isomorphisms between different internal spaces. Replace 'solder form' in the above quote with a more appropriate term like 'solder map', and the point is clear. The coordination between different internal vector spaces is achieved by having a solder map between the vector space $\mathbb{V}$ of the associated bundle and the vector space $\mathbb{V}_j$, corresponding to the field value space of the jth matter field. And, of course, this is only achievable if the vector bundles are isomorphic as associated bundles *associated to the same principal bundle*.

So, to summarise the result of this section, note that, on the geometrical approach, the tangent spaces automatically coordinate the spatiotemporal behaviour of all matter fields. This allows us to think of the gauge symmetry of general relativity as being equivalent to a freedom of choice of frame at each tangent space to the base manifold, and this can be done without recourse to talk of principal bundles. Weatherall argues that this move can be made for a generic Yang–Mills covariant derivative, modulo the issue of coordination. I argued that the issue of coordination is central to the deflationary interpretation of gauge transformations as changes of internal vector bases. The coordination can only be achieved by an associated bundle isomorphism generated by a principal bundle automorphism. Thus the principal bundle is completely indispensable when it comes to coordination, whether it is mentioned explicitly or not.

On the geometrical view, the solder form on the principal bundle does not have to be expressed explicitly. In the next subsection, I will demonstrate why this is a shortcoming of the geometrical view—and, indeed any view that coordination by symmetry constraints alone fixes the global trajectories of matter—they exclude the possibility of matter fields' dynamics requiring, in Weatherall's words, 'more resources'.

Canonical solder forms

Weatherall argues that the mathematical difference between internal vector and tangent spaces make no physical difference; consider $\xi^A_{|p}$ and $\xi^A_{|q}$, where p and q are two distinct finitely separated points on the manifold, and $\xi^A_{|p} \in V_p$, $\xi^A_{|q} \in V_q$. We have one way of ascertaining whether $\xi^A_{|p}$ points in the same direction as $\xi^A_{|q}$, but that does not involve distant comparison of field values, rather, it involves parallel transporting $\xi^A_{|p}$ to V_q along some curve. Further, there is no relationship between $\xi^A_{|p}$ and the tangent vector at p to the curve along which ξ^A is parallel transported—they are in different spaces, so the tangent field along the curve cannot provide a standard of rotation.

However, when we parallel transport a tangent vector along some curve, then, at each point, we can meaningfully ask whether it has rotated with respect to the tangent field along the curve, because now all relevant vectors at a point are in the same space. Accordingly, ‘there is a sense in which one can say a bit more about how vectors change when parallel transported along open curves, relative to a standard of constancy given by the tangent to the curve along which one has transported.’ [122, p. 2411]. Weatherall goes on to argue, *pace* Anandan [4], that this extra information is physically dispensable: ‘it is far from clear that, in the general case, the change in angle of a parallel transported vector relative to the (arbitrary) curve along which it is parallel transported measures any quantity of special interest, or could stand in as a notion of “holonomy” along open curves’ [122, p. 2412].

This argument is worth discussing for two reasons. First, it is incorrect; there are examples of cases in which it would constitute a ‘quantity of special interest.’ Second, it reiterates what the geometrical approach is for Weatherall. If Weatherall is correct in his judgement about the dispensability of this ‘open curve holonomy’ information, then he is committed to the idea that for all physical fields, the proper time measured by some physical process always corresponds to proper time intervals along its worldline. In short, he assumes the truth of the *clock hypothesis*:

Clock hypothesis: For each kind of matter, there exist configurations whose dynamics is such that they can be used to record, in an inertial or suitably gently accelerated state of motion, intervals of proper time along their worldlines, as determined by some metric. Further, this is the metric with respect to which clocks built out of other matter fields will record proper times along their worldlines.

Consider a timelike curve corresponding to the following setup: two perfect mirrors reflecting a photon between them. In order for this clock to measure intervals of proper time along its worldline (that is to say, our light clock is so small that any spatial extension is negligible for our purposes), the light ray must traverse a geodesic of constant velocity between them. This appears not to be a problem, as long as light traverses null geodesics, which are straight lines in spacetime. But, as was discussed in chapter 1, it is demonstrably not always the case that light traverses null geodesics in all spacetimes.[23]

The precise details, presented in that chapter, need not concern us here. All we need to rebut Weatherall's claim is the fact that the trajectory of light in a particular Gödel spacetime, a legitimate exact solution to the Einstein field equations, is not a null geodesic of g_{ab}. This means that the tangent vector corresponding to the trajectory of light is *not* the tangent vector that corresponds to a null geodesic of g_{ab} at that point. So integrating the line element as given with respect to g_{ab} between p and q does not given us the actual length of the trajectory traversed by light in that spacetime. So deviations from null behaviour as determined by g_{ab} *are* physically relevant. And we get this extra information from a choice of solder form; this is why it cannot be assumed once and for all, to be common to all dynamical fields.

4.4.2 Wallace, Weatherall and the dynamical–geometrical debate

It is instructive to use this disanalogy, as Wallace does, to frame general relativity as a different kind of Yang–Mills theory. Weatherall's and Wallace's formulations are similar to more familiar Yang–Mills theories in some ways, and dissimilar

[23] The discussion in that chapter is based on a recent paper by Asenjo and Hojman [7].

in others. Weatherall's is similar in the sense that the structure group on the principal bundle is compact, and the associated fibres are vector spaces, while Wallace's is similar in the sense that the associated fibres are truly 'internal' on the definition given in §4.2.

In this section, I discuss the relationship between the two formulations. This makes very clear what Weatherall's conception of the geometrical approach to general relativity is, and allows me to state the geometrical–dynamical debate in a novel and hopefully perspicuous way. I should mention that, while Weatherall is committed in his paper(s) to the geometrical approach, Wallace is less easy to pin down, at least in print.

Recall from the last section the importance of 'coordination' of the behaviour of matter fields. What made the Levi-Civita derivative operator special was that it was the derivative operator with respect to which all fields' directional derivatives were expressed. I argued, *pace* Weatherall, that the coordination is determined, among other things, by the structure group of the principal bundle. The solder form in general relativity played a crucial role in this coordination (although the Gödel solution demonstrated that more needs to be said even after coordination is established).

On the fields-as-bodies approach, Weatherall's argument cannot be formulated, since the associated bundle fibres are affine spaces, which do not have bases. Recall that the fibres were affine spaces because of the translation component of the Poincaré group, which Wallace took to be the structure group of the principal bundle. The connection on the principal Poincaré bundle could be split into two components—a Lorentz connection, ω and a translation connection, θ. A choice of gauge for the translation symmetry is tantamount to breaking the translation symmetry. This breaks the overall symmetry from Poincaré to

Lorentz—Wallace refers to this as *stationary gauge*. In effect, this is choosing an origin for each internal space (these are the spaces in which the location field is valued), making it a Minkowski vector space.

Recall that spacetime vector and tensor fields had a special property—they assigned directional derivatives for all fields. So spacetime vector fields, on the geometrical approach, are just tangent vector fields to the spacetime manifold. On the fields-as-bodies view, the tangent space to the base manifold is not the spacetime tangent space. So in order to distinguish 'spacetime' vector fields from 'internal' vector fields, one needs to identify an element of this Minkowski vector space with an element of the affine Minkowski space in which the location field is valued. We need a solder map. There is a convenient way to define one—just choose an origin for each affine Minkowski space (the location field space), and it turns into a Minkowski vector space in which the 'spacetime' vectors are valued. In other words, go into stationary gauge. This makes the translation connection into a solder map (note, not a solder *form* because it bypasses the link to the tangent bundle TB), and coordinates the internal spaces (of which the 'spacetime' location field space is one) in just the way that the solder map did in the previous section.

We can see how Wallace's picture is more general than, and reduces in two steps to, Weatherall's. First, the affine Minkowski space is soldered onto the 'spacetime' Minkowski vector space by a choice of solder map generated by a choice of translation gauge (i.e. origin). Second, if we make that choice, once and for all, for all matter fields, and then decree that the Minkowski vector space is also the tangent space to the base manifold, then we arrive at Weatherall's picture. But, on Wallace's picture, the choice of origin is a choice of gauge, so has no physical significance. Different dynamical fields might have different 'spacetime location' behaviour, and this is reflected in a different choice of solder map.

The dynamical–geometrical debate

From the discussion above, I take the geometrical approach as espoused by Weatherall to be committed to the following:

(1) There is only one location field, ρ^a.

(2) All material fields, represented by tensors or spinorial tensors are assigned locations in the affine Minkowski space in which the location field is valued; for all fields, this affine Minkowski space is soldered to the Minkowski vector (tangent) space *by the same solder map*.

The dynamical approach asserts that (2) might be false, while being agnostic about (1)—as long as different material field spaces are soldered differently to the affine Minkowski space, it does not matter whether one uses one or several copies of the affine Minkowski space. This characterises the two positions purely in terms of properties of geometrical objects to which each side is committed, rather than expressing the dispute in terms of metaphysically abstruse concepts like explanation.

4.5 Conclusion

This chapter had two aims: (i) to introduce the formalism of parametrised field theory in order to articulate precisely the dynamical–geometrical dispute in general relativity and (ii) to use this to argue against Weatherall's recent deflationary account of gauge symmetries. On the parametrised field theory setup, we saw that the geometrical view could be characterised by two things, first, a commitment to a single location field, and second, a commitment to all material fields' 'external' spaces being soldered to the affine Minkowski space in the same way. The dynamical view eschews the latter commitment. The example of the light clock in Gödel spacetime gave us reason to support the dynamical view over both the traditional as well as Weatherall's version of the geometrical view, and also broke the analogy between general relativity and Yang–Mills theories that Weatherall relied on to support his interpretation of gauge symmetries.

CHAPTER 5

The kinematical structure of special relativity

5.1 Introduction

The dynamical approach to spacetime theories zeroes-in on a correlation between facts about geometry and facts about the behaviour of dynamical fields (such as those constitutive of rods and clocks). The position asserts, among other things, that facts about physical geometry are grounded in, or explained by, facts about dynamical fields, rather than the other way round. The converse position—that geometry is explanatory of matter field behaviour—is ubiquitous, and is the orthodox position on physical geometry and spacetime structure.[1]

John Norton has been a persistent critic of the dynamical approach, and, in 2008 [83], articulated the following objection to the view: the proponent of the dynamical approach, the so-called *constructivist*[2] is illicitly committed

[1] Examples in the physics literature include the classic [79] by Misner, Thorne and Wheeler [79] and the more recent *Modern Classical Physics* by Thorne and Blandford [110]; among philosophers, Friedman's influential *Foundations of space-time theories* [45] and Maudlin's *Philosophy of Physics: Space and Time* [77] both embrace a geometrical perspective.

[2] This is Norton's neologism for a proponent of the dynamical approach—in this chapter, when setting up the position, in order to remain consonant with the terminology of the rest of this thesis, I use the term 'dynamical approach', and when discussing Norton's criticism, I use the term 'constructivism' and 'constructive relativity' to mean the same thing as 'dynamical approach'.

to spatiotemporal presumptions in 'constructing' spacetime from facts about dynamical symmetries.

The image of construction brings with it an image of building blocks. The standard way of representing (at least some of) the building blocks of a theory is in terms of *kinematical* structure. Broadly speaking, this is the structure that needs to be presupposed in order to build models of physical goings-on. These models are known as *kinematically possible models (KPMs)*, and are usually expressed in the 'angle bracket' notation introduced in chapter 3 as $\langle M, g_{ab}, \Phi^i \rangle$, where M is a smooth manifold, g_{ab} is a Lorentzian metric and Φ^i is a place-holder for matter fields.

Theories are associated with collections of models specified by picking out some subset of a space of KPMs.[3] On their standard construal, including on Norton's way of thinking about them, these models implicitly have an extra bit of structure—a link between the manifold and the fields. Specifically, for any point in M, distinct fields, $\phi_1, \ldots, \phi_n$, can be 'evaluated' at that point, and their values can then be taken (as e.g. Field does [42]) to represent properties of the same spacetime point. Call a KPM in which this link does exist a *KPM of the first kind.* Constructing KPMs in this way is so intuitive that it is easy to lose sight of the fact that *nothing in the mathematical structure of the theory makes the link necessary.* The reason that the link is not necessary is that elements of the set M do not have to be assumed to have primitive identity.

I should point out that I make no commitment to whether or not it is mathematically or representationally necessary that the points of a set have some primitive distinguishing feature; all I claim is that one must be aware of whether

[3] This is what is constitutive of a theory on the so-called *semantic view* of scientific theories, due primarily to Suppes [108], and popularised by van Fraassen [113, Ch. 3]. This is the view of scientific theories that I assume in this chapter.

or not such an assumption is being made at the mathematical level, and what its justifications might be. In what follows, we will be able to interpret Norton's physical requirement as justifying his commitment to the primitive identity of spacetime points in his models.

When we express fields as 'functions of spacetime', in the form '$\phi(x)$', they are composite functions acting on a set of points, M, of the form $(\phi \circ x)(p) : M \to \mathbb{V}$, where $\mathbb{V}$ is some space in which the fields take their values and $x : M \to \mathbb{R}^4$ is a coordinate function.There is a notational ambiguity here. It is sometimes the case that 'x' is used to represent the coordinate function assigned to a point in M, its value, and also the pre-image of a point in $\mathbb{R}^n$ under that function. In this chapter, I use 'x' to represent the function, and p to represent the pre-image of a point in $\mathbb{R}^4$ under x, i.e. $x(p) \in \mathbb{R}^4$, where $p \in M$.

Given the arbitrariness of the coordinate function, there is no *a priori* need for the coordinatisation component of two separate composite functions, $\phi \circ x$ and $\psi \circ x$ to coincide. In other words, there is no *a priori* reason to require that the x in $\phi(x)$ is the same as the x in $\psi(x)$—our use of the same variable to denote both is indicative of a failure to allow for this mismatch. Call models in which the p in $\phi(p)$ and $\psi(p)$ can represent different points in the coordinate function's target space $\mathbb{R}^4$, *KPMs of the second kind.*

As we will see in §5.4, Norton's criticism is based on assuming that the only sort of KPMs appropriate for special relativity are of the first kind. The central aim of this chapter is to demonstrate that KPMs of the second kind can be constructed for special relativity. This chapter is not the first response to Norton's claims—Brown and Pooley themselves have responded. Their response, exemplified by Pooley in [92, §6.3.2], and discussed in §5.2 is to accept that the constructivists were, in fact, committed to some of these presumptions, but to argue that such presumptions are not illicit. I believe that, insofar as their response embraces the use of KPMs of the first kind, it concedes too much.

In this chapter, using the fibre bundle setup detailed in chapter 4, I construct kinematically possible models of special relativity of the second kind (in §5.3). I then demonstrate, in §5.4, that the dynamical approach viewed in these terms is immune to the charge that Norton levels against it. First, though, I present a short recapitulation of the relevant aspects of the dynamical approach in §5.2.

5.2 The dynamical approach

The dynamical approach to special relativity takes seriously the fact that measuring devices like rods and clocks are themselves entities whose constituent matter fields are governed by dynamical laws. Consequently, any explanation of their behaviour must make reference to facts about the dynamical equations that govern them. Falling back on a *geometrical* explanation of the form 'rods and clocks contract because they are embedded in Minkowski spacetime' is, in the eyes of the proponent of the dynamical approach, unsatisfactory, unless the geometry itself is taken to be a convenient shorthand for the symmetry structure of the dynamical field equations.[4]

A specially relativistic field theory is standardly presented in terms of dynamically allowed maps from a metric manifold into some mathematical space. The metric structure on the Minkowski manifold encodes certain facts about dynamical symmetries—these are transformations on the fields that leave the form of

[4] This is what grounds the ontological reduction of the Minkowski metric field in (and, indeed, any theory with a non-dynamical background spacetime). No such reduction is possible in metric theories like general relativity, in which the metric has associated degrees of freedom. But the moral of the dynamical approach—that the behaviour of matter fields is the ultimate explanans of the behaviour of matter—is equally applicable there. For a discussion of the dynamical approach in general relativity, see e.g. [18, Ch. 9], [97] and the introduction to this thesis.

the dynamical equations unchanged. According to the dynamical approach, the metric structure attributed to this manifold is nothing more than a reflection of this dynamical structure of matter fields. In particular, there is no sense to be made of the claim that the Minkowski manifold constrains or explains the symmetry structure of the dynamical laws. Further, the position engenders an *ontological* claim—the Minkowski metric is thought of as a 'glorious non-entity' [20], in the sense that it is not ontologically independent of matter fields.[5, 6]

This raises an immediate question: even if one accepts that the metric is a non-fundamental entity, what is the status of the smooth manifold of spacetime locations? After all, the smooth manifold is the space of independent variables in terms of which the field theory is expressed. If that too is a non-entity, then how can one speak coherently of a field theory? How can one even define a field? Pooley's response [92, §6.3.2] is that he and Brown never intended to tell a *reductive* story about *all* putatively 'spatiotemporal' structure.[7] They still presuppose irreducible smooth structure, but, importantly, attribute this structure to the matter fields themselves. As such, they do not presuppose illicit *spatiotemporal* smooth manifold structure in the manner in which Norton suggests. The dynamical symmetries of the fields then ground the further metric structure of the manifold.

The account I spell out in this chapter is closely related to Pooley's suggestion; indeed, it can be seen as one way of making sense of the topological structure

[5] As mentioned in the introduction to this thesis, this is not to say that proponents of the dynamical approach are anti-realists about the Minkowski metric; just anti-fundamentalist).

[6] Roughly speaking, a substance is ontologically independent if it is such that its existence does not require the existence of something else. One potentially useful way of cashing out this idea of ontological independence is in terms of essence [43], but here I make no commitment to a specific way of doing so.

[7] I use the term 'putatively spatiotemporal' to avoid having to commit one way or another to a particular piece of structure's being categorised as spatiotemporal. It is a plausible inference, based on Norton's own work on the subject (in e.g.[38]), that he considers the bare manifold to be a candidate for spatiotemporal structure, so manifold points are putatively spatiotemporal.

of extended matter. In particular, I demonstrate how the proponent of the dynamical approach can tell a story about topological and smooth structure which is not itself spatiotemporal. Rather it is the primitive structure of the 'body' manifold described in the previous chapter. The picture of spacetime which falls out of this account paints spacetime location as a property of a field.

5.3 Fields as bodies

In §5.3.1 I briefly recapitulate the 'fields-as-bodies' set up detailed in chapter 4. Following that, in §5.3.2, I discuss the link between the two types of KPMs, and the fields-as-bodies picture.

5.3.1 Fields as bodies revisited

Recall, from chapter 4, that a fibre bundle is is a quadruple $\langle E, \pi, B, F\rangle$, where E, B and F are smooth manifolds, and $\pi : E \to B$ is a surjective map, the target of whose inverse is called the 'fibre', F, above $p \in B$. Further, there exists a set $\{U_i\}$ of covering manifolds to M, and an associated diffeomorphism, $\phi_i : U_i \times F \to \pi^{-1}(U_i)$ such that $\pi \circ \phi(p, f) = p$. Intuitively speaking, a fibre bundle is a manifold built up of a 'base manifold' each of whose points is associated with its own copy of some other manifold, known as a fibre. These fibres, as long as they are manifolds, can be endowed with further structure—vector space structure turns out to be particularly useful in setting up Yang–Mills theories.

On a standard fibre bundle picture of a classical Yang–Mills theories, the base manifold is the spacetime manifold, and the fibres are matter field-valued spaces carrying a representation of some symmetry group. The force (or gauge/Yang–Mills) field associated with the interactions of these matter fields, under the force described by the theory, is represented by a Lie-algebra valued connection one-form. This connection can be thought of as being defined on a principal G−bundle, where G is the symmetry group associated with the force in question (e.g. $U(1)$ for electromagnetism, $SU(3)$ for the strong force).

This connection, on the entire principal bundle, is pulled back along a section, σ, to the base manifold, on which it is taken to define a force field on the base manifold. This force field is Lie-algebra-valued, and, in effect, provides a systematic correction term to the flat derivative operator such that the dynamical equations of the theory, when expressed with respect to this 'corrected' derivative operator (the so-called covariant derivative) are gauge-invariant.

This formal setup makes contact with physics that describes the real world by way of an interpretation of some of its constituent parts as representing physical bodies or entities. In particular, the base manifold is often taken to represent spacetime (this is what makes it appropriate to think of the base manifold as a metric manifold, i.e. a manifold which specifies an inner-product on its tangent spaces), the connection is taken to represent the force field and sections of the matter field bundle (usually an associated vector bundle carrying a representation of the symmetry group of the theory; for more details, see §4.2 of chapter 4) represent matter field configurations.

This way of doing this has a number of advantages—most importantly, it makes clear some useful formal analogies between different forces. But if we wish to think of gravity as a force (and there is a long history of distinguished people who did and continue to do so[8]) whilst conceding that general relativity is our gravitational theory of choice, then we run up against a problem. The problem is that, while other Yang–Mills theories are unequivocal about what

[8] It is a central pillar of the string theoretic approach to general relativity as detailed in e.g. [58, 91]. Even outside that context, there exist proponents of the so-called spin-2 approach to general relativity, e.g. Weinberg [125].

counts as a connection and what counts as the appropriate gauge group, general relativity is not. In particular, its central dynamical connection, the Levi-Civita connection, is not mathematically analogous to the dynamical connections of Yang–Mills theories like electromagnetism. As a result there is disagreement over how (if at all) general relativity should be seen as a Yang–Mills theory.

An appropriate reformulation is presented by Wallace [119] and described in detail in chapter 4. The trick is to stop thinking about the base manifold as representing spacetime. Instead, very roughly speaking, we think of attaching a copy of spacetime to each point on a base manifold which we refer to as the 'body' manifold'. This manifold is just a mathematical representation of the entirety of a physical field. We are allowed to construe a field in this way, since nothing in the mathematical set up of Yang–Mills theories precludes it; but what we do need to provide is a story about how this picture is consistent with extant physics. Again, this was discussed in detail in chapter 4. The upshot of that discussion is the following.

Conceive of a field as consisting of a body, B, each point of which is associated with a number of properties, for example, spatiotemporal location or charge-related. Each of these properties can take some numerical value; so a mathematical space of all of these possible values is associated with each point in B. The spatiotemporal properties are encoded in the 'location field', ρ^a, which takes values in an *affine* Minkowski space (this is because the location field, now treated as a *matter* field, takes values in a mathematical space endowed with a representation of the Poincaré group). The translation component is what mandates our choice of an affine fibre, while the Lorentz component mandates our choice of metric—the Minkowski metric—on the fibre.

Whereas on the standard picture, one might represent a field, say $\phi(p)$ as taking, as its argument, a point in *spacetime*, on this picture, a field takes its value at some point on the body manifold. That point on the body manifold will be associated *both* with a spatiotemporal location property (determined by the value of the location field) and the relevant field property. In effect, this equips each field with its own location field space, and the links between these individual spacetimes are provided by a choice of solder map (this was the moral of chapter 4). What this means is that, absent some dynamical reason for choosing a way of identifying points across different affine Minkowski fibres, we have constructed kinematically possible models of the second kind. In the next subsection, I discuss this in more detail.

5.3.2 Kinematically and dynamically possible models

If we begin with the set of KPMs, which represents the set of metaphysically possible models consistent with a specification of certain degrees of freedom necessary to write down dynamical laws, we arrive at the set of *dynamically possible models (DPMs)*. This is a proper subset of the set of KPMs, specified by some dynamical equations that the objects of the KPMs must satisfy.

Consider a dynamical field, $\phi(\bar{x})$, as specified by a KPM of the first kind. This sort of field (let us assume, for dialectical clarity, that it is scalar) is a function of the following kind:

$$\bar{\phi} : AM \rightarrow V$$

where AM is an affine Minkowski space and V is the space in which the field takes its values. This same field is expressed in the parametrised field theory formalism as follows:

$$\phi : B \to AM \oplus V$$

where B is the body manifold and $\phi(p) = (\rho^a(p), \lambda(p))$, where we refer to $\phi(p)$ as a *material field* and, as before ρ^a is the location field, and $\oplus$ is the direct sum operator. All of the dynamics can be encoded in a suitably reparametrised Lagrangian function (for details, see chapter 4 or [119]). We can now consider two fields, $\psi(p)$ and $\chi(p)$, both of which attribute *different* properties to some point on the body manifold. Importantly, we can attribute distinct location field values to $\psi(p)$ and $\chi(p)$ in exactly the same way as we can attribute distinct values to any other field. The spatiotemporal property is, recall, just another property of the body manifold. As a matter of convenience, we might choose the location field values of each point on the body manifold to be the same—this would then be a special case of the base manifold being in one-to-one correspondence with spacetime location. However, this is just a choice of gauge (or solder map), and the translation component of theory's Poincaré invariance means that any choice is as good as any other. This strips the points of the location field space of primitive identity—only differences between location field-values are meaningful.[9]

We thus see that the tight link between KPMs and primitive spacetime point coincidence of field arguments is broken. In denying that the proponent of the dynamical approach presupposes these coincidences, I do not claim to have done away with the need for kinematical structure. Rather, I claim that the kinematical content of the theory can be specified without primitive spacetime point coincidences—in other words, that KPMs of the second kind can be

[9] It might be argued, although I will not do so in this thesis, that this gauge freedom allows for a distinct realisation of the position of *sophisticated substantivalism* advocated in various forms by Brighouse [17], Butterfield [24], Hoefer [56] and Pooley [92].

constructed. In particular, there is no assumption of a set of *spacetime* points, structured as a manifold, providing the space of independent variables of which field values are predicated.

5.4 Norton's criticism

Norton claims that the dynamical approach to relativity (or, to switch to his term, constructive relativity) fails. He advances the following criticism of the constructive relativist project: that it 'only succeeds if constructivists antecedently presume the essential components of a realist conception of spacetime' [83, p. 821]. Norton is explicit about what, for him, constitutes a realist conception of Minkowski spacetime:

> (1) There exists a four-dimensional spacetime that can be coordinatized by a set of standard coordinates (x, y, z, t) related by the Lorentz transformation.
>
> (2) The spatiotemporal interval s between events (x, y, z, t) and (X, Y, Z, T) along a straight [footnote suppressed] line connecting them is a property of the spacetime, independent of the matter it contains, and is given by
>
> $$s^2 = (t-T)^2 - (x-X)^2 - (y-Y)^2 - (z-Z)^2$$
>
> (3) Material clocks and rods measure these times and distances because the laws of the matter theories that govern them are adapted to the independent geometry of this spacetime. [83, p. 823]

Before dealing with the specific challenges Norton levels against the constructive project, it is worth addressing two ways in which this conception of spacetime realism misrepresents the constructive project's aims (insofar as it asserts that the constructive project must deny them). First, that the constructivist must deny the existence of a 'four-dimensional spacetime that can be coordinatized by a set of standard coordinates.' Nothing in the setup of the dynamical approach commits its adherents to the denial of the existence of spacetime—all that is required is that spacetime structure (to the extent that it can be identified) be reducible to facts about the dynamics of fields. In other words, the constructivist need only be committed to the non-fundamentality of spacetime, *in the sense that it is ontologically dependent on matter fields.*[10]

Second, that the spatiotemporal interval is a 'property of the spacetime independent of the matter it contains.' This notion of independence is ambiguous. It could either mean (i) that the interval is a property of spacetime that does not require matter fields in order to be intelligible (call this the *ontological* reading) or (ii) that the interval still requires matter fields in order to be intelligible, but is neutral with respect to idiosyncratic details of *which* of several distinct matter fields is being considered (call this the *epistemological* reading). The constructivist denies the ontological reading, but accepts the epistemological reading. Thus the spacetime realist position that the constructive denies is more accurately characterised as follows:

> (1) There exists a **fundamental** four-dimensional spacetime that can be coordinatized by a set of standard coordinates (x, y, z, t) related by the Lorentz transformation.
>
> (2) The spatiotemporal interval s between events (x, y, z, t) and (X, Y, Z, T) along a straight [footnote suppressed] line connecting them is a prop-

[10] cf. fn. 6.

erty of the spacetime, **ontologically** independent of the matter it contains, and is given by

$$s^2 = (t-T)^2 - (x-X)^2 - (y-Y)^2 - (z-Z)^2$$

(3) Material clocks and rods measure these times and distances because the laws of the matter theories that govern them are adapted to the **ontologically** independent geometry of this spacetime.

Norton's characterisation of the spacetime realist position captures the implicit assumptions present across the literature in the foundations of spacetime. Take Anandan, for example:

> In classical physics... independent evidence is provided for the events, which constitute space-time, by particle trajectories and collisions between particles. In special relativity an additional reason for assuming the existence of space-time points is provided by the fact that in this theory there are fields that are functions of space-time which have independent reality. [3, p. 610]

Anandan is explicit about his assumptions in the context of spacetime geometries associated with classical particle theories. It is worth dwelling on these assumptions, because of their similarity to Norton's, despite Anandan's project having a number of similarities to Brown and Pooley's: Anandan's project is based on the assumption, explicit in the work of Brown [18], that spacetime geometry is nothing more than a Kleinian geometry associated with the dynamical symmetries of a matter theory. The precise form of this Kleinian geometry is, obviously, theory-relative, and Anandan's paper [3] proceeds to spell this out in detail for a number of classical and quantum theories.

In the context of classical particle physics, Anandan takes as primitive 'the notion of particles and a set of states associated with each particle called its dynamically possible states. In classical physics, it is assumed that each of these states can be represented by [subsets of M] called the space-time and whose elements are called events. These subsets, called particle trajectories, are not all disjoint' [3, p. 606]. Since, at this stage, Anandan has not specified the dynamics, what he describes above corresponds to our notion of kinematically possible models. It follows, from the trajectories and particles being primitive, and their being not all disjoint, that the points of collision are also primitive. If spacetime is composed out of events, and these events correspond to actual or possible collisions, then the connection between points of M and the space of events is also primitively established—we are, therefore, dealing with kinematically possible models of the first kind.

But now that we have moved the spacetime location values into an internal fibre, the notion of spacetime coincidence (in particular, the choice of a solder map) becomes a matter of dynamics of the material and location fields *taken together*. With each point on the body manifold, B,[11] equipped with its own affine Minkowski location-field space, we can explicitly construct KPMs of the second kind: $\langle B, R^i, \Phi^i \rangle$, where Φ^i is a place-holder for all material fields and R^i for all location fields (recall from the argument in the last chapter that the proponent of the dynamical approach is happy to countenance such a scenario).

5.4.1 Claim one

> The construction project must tacitly assume an already existing spacetime endowed with topological properties, so that it can in-

[11] Recall the distinction between the use of M for a spatiotemporal base manifold and B for a body manifold from chapter 4.

> troduce spatiotemporal coincidences, and a unique set of standard coordinates (x, y, z, t). [83, p. 824]

Norton here assumes that no sense can be made, in purely constructive terms, of a spatiotemporal coincidence, as a result of which, it needs to be presupposed. But once we realise that 'spacetime point', in the sense required to get a dynamical theory off the ground, is simply equivalent to 'location field value', then it is clear that Norton's claim is not true. Granted, the fields-as-bodies setup did presuppose topological properties on the body manifold, B, but this is not the structure that Norton claims is illicit. The topological properties associated with spacetime locations and coincidences are derived wholly from the dynamical symmetry group—this is not presupposed.

Spacetime points do not enter into the formalism at all—except insofar as they are related to the coincidence of location field and material field values at some point on the body manifold. It then becomes a matter of dynamics, determined by the Euler-Lagrange equations of the parametrised field Lagrangian, whether location field values coincide with material field values at distinct points on the body manifold. Norton conflates a difficulty in having epistemic access to spacetime coincidences with an inability to define them; the constructivist can perfectly easily make sense of an appropriate *definition* of a spacetime coincidence, without having to have a distinct conception of a spacetime point.

Consider, now, a coordinate system on which we have good reason to believe that, say, a Gaussian wave packet of the $\phi(x)$ field bounced off a Gaussian wave packet of the $\psi(x)$ field in the neighbourhood of some point (in spacetime). On the fields-as-bodies picture, this corresponds to the values of two distinct material fields evaluated on the body manifold, both of which are associated with the same location field (in the terminology of the last chapter, this implies the choice of a solder map).

There will be a class of location field valuations which assign locations in the affine Minkowski fibre (i.e. 'spacetime' locations) to material field values in such a way that the dynamics of each field determines that a collision took place in the vicinity of some spacetime location, and another class of location field assignments on which the kinks in the trajectories (or, more generally, some fact about the dynamical interaction) of each particle do not take place at the same spacetime location. On pragmatic grounds, the former class of location field assignments, hence solder maps, might be preferable. For free fields, on the other hand there simply is no operational sense of spatiotemporal coincidence, although they can still be defined by an arbitrary choice of solder map.

This observation points to an arbitrariness in our definitions of 'free' and 'interacting' particles. The trajectory of a particle of the ϕ field might be such that it contains a change in direction that, on an interacting field picture, we attribute to it having 'collided' with another field, ψ at some point. We can always absorb these kinks into a complicated dynamical picture of free fields, and we lose no descriptive power. But on this picture, each field, in a very rough sense 'inhabits its own spacetime'. In such a theory, the means of having epistemic access to a particular point (or arbitrarily small region of spacetime) is through the interaction dynamics—but this is not how we need to *define* a spacetime point coincidence And Norton's criticism of the constructivist is based on their purported inability to define such a point coincidence.

5.4.2 Claim two

> The Lorentz covariance of all matter theories asserts an adaptation between matter and this spacetime akin to the realist's [(3)], although

> without [(3)]'s presumption of the direction of the adaptation of matter to spacetime. [83, p. 824]

This second criticism loses all of its bite once it has been shown that spacetime coincidences do not have to be presupposed by the constructivist. The structure of the location field fibre is determined by the choice of symmetry group—for a locally Poincaré-invariant theory, it is the Poincaré group. The choice of symmetry group follows from Brown and Pooley's truncated Lorentzian pedagogy, as the largest non-trivial intersection of the dynamical symmetry groups of the fields in one's theory. This can be realised in various ways—Pooley [92] and Stevens [107] do it in Humean terms, extending work done in the context of Newtonian mechanics by Huggett [57]. Wallace [120] follows Anandan's [3] Kleinian approach to much the same end. However this process may be carried out, ultimately, the structure of the location field fibre merely reflects the symmetry properties of the dynamical fields. In that sense, the (metric) structure of spacetime is, indeed, adapted to the matter fields *in virtue of having been derived from them*, and it could not have been any other way. There is no 'direction of adaptation' between spacetime and matter any more than there is a direction of adaptation between 'bachelor' and 'unmarried man'.

5.4.3 Claim three

> Constructivists must accept that spatial distances and times elapsed are properties of spacetime as asserted in [(2)], on pain of failing to reconstruct traditional spacetime geometry and also having to accept an extreme form of operationalism in which quantities have values only if they are actually measured. [83, p. 824]

Once again, this line of argument from Norton builds on the success of previous lines of argument in his paper. Having dealt with those in previous sections, it becomes easier to dispense with this final argument.

To bolster his claim that the constructivist is committed to 'an extreme form of operationalism', Norton presents an example of a region of spacetime that is either (A) devoid of matter or (B) hosts a static matter distribution. On this, Norton offers the following:

> In this part of spacetime, we can select two noncoincident timelike-separated events A and B such that nothing changes as we pass along the straight segment of spacetime connecting them. In the ordinary realist's conception, we would say that some time elapses between them. What can a constructivist say? There are no material clocks actually present measuring the time elapsed, for there is either no matter present or no change in the matter present as we pass from A to B. So the constructivist has no material basis for the recovery of a time change. If times elapsed are to supervene on matter, or more vaguely to be a result of the properties of matter, then the absence of any change in the matter entails that there is no change in times elapsed. [83, pp. 831-832]

But as Pooley points out, 'for the constructivist there is literally nothing in an empty region and so nothing whose geometrical properties might be indeterminate. The constructivist does not believe in the existence of an independently existing spacetime!' [92, §6]. This corresponds, then, to certain regions of the body manifold not having spatiotemporal locations at all. In the fields-as-bodies picture, this would mean that there exists no dynamically privileged solder map. The location field would continue to take values in an affine Minkowski space, but this space would no longer correspond to spatiotemporal locations. This is a consequence of not endowing the location field-value space with a primitive spatiotemporal interpretation (by stipulating a solder map) and choosing, instead, to allow the dynamical material fields' Lagrangian (which, recall, includes the location field) to do so.

What about 'static' regions? These would correspond to regions of the body manifold with the same material field values across different location field values. Two responses suggest themselves; I am inclined to adopt the second. The first is to adopt the operationalist response that Norton explicitly rejects: confined to those regions of the body manifold where the material field values are static, the material field could not be used to construct clocks (there is, by definition, no periodic process allowed). However, this misses the point of Norton's criticism, for even for the proponent of the dynamical approach, there still remains a matter of fact about, say, the proper time elapsed along some trajectory between A and B as determined by the Minkowski metric, *under the assumption that the material field dynamics is Poincaré invariant.* And this is the sort of thing that the proponent of the geometrical approach would like to assert.

The second response is to double down on the dynamical approach—if the only material field in the universe is one which takes the same value at all spatiotemporal points, then there is no reason to suggest that its 'dynamics' should be merely Poincaré invariant—surely then they would be invariant under the drag-along of any diffeomorphism. This would simply mean that there was no matter of fact about 'proper time' between A and B, the latter property requiring a metric to be sensible. So, Norton is completely right when he says 'the absence of any change in the matter entails that there is no change in times elapsed', even though he intends it to be the unpalatable conclusion of a *reductio ad absurdum*.

5.5 Conclusion

The dynamical approach to physical geometry is, in effect, a call for explanatory rigour. It asserts that an explanation about the behaviour of matter fields requires, as its ultimate explanans, dynamical facts about those matter fields. In this chapter, I presented an slightly extended version of this approach, based on the observation that the kinematically possible models can be characterised without requiring primitive spacetime point-coincidence of fields. Once we are aware of the implicit link in our KPMs between spacetime points and 'points of independent variable coincidence of field values', it is important to determine whether or not this structure is epistemically warranted and dynamically necessary. I argued that it is neither.

Using the parametrised field theory formulation of classical field theories, which did away with the kinematical requirement that fields be defined as maps from a spacetime manifold, I demonstrated how 'KPMs of the second kind' could be constructed. These KPMs (and their associated Poincaré-invariant DPMs) were then be used to counter Norton's criticism of the dynamical approach.

Conclusion

Spacetime and matter fields both come with geometries, but the relationship between them is fraught. Part of the problem with a straightforward geometrical interpretation of physical theories is that, in virtue of its using 'fewer resources', it is unable to capture important subtleties related to idiosyncratic behaviour of matter fields. The power that spacetime geometry has to simplify our physical theorising comes at a price—we have to be careful not to demand results that it is not in a position to deliver. Such results include the common, unequivocal assertion that light travels on null geodesics of the metric of general relativity; that well-constructed clocks can be made to survey this metric arbitrarily accurately; and that this metric determines which degrees of freedom suffice to infer facts about the behaviour of all matter fields.

The metric, though, does not always deliver on promises that it is sometimes taken to have made. Some diagnosis of these failures is called for; providing one was a central concern of the first part of the thesis. As we saw in chapters 1-3 there are several reasons: approximation procedures (chapter 1), failures to account for ineliminable non-chronometric roles for the metric (chapter 2), and a shift of the link between algebraic and geometrical structure (chapter 3), to give a few examples.

But diagnoses are most helpful when they come with a prescription for a cure; one might worry that the cure in this case will be something radical—a new theory, or a perhaps a bold new interpretation of the old theory. The second part of the thesis shows that neither was required. The formal setup of relativity theory is sufficiently robust and flexible to lend itself to an interpretation that accounts for the observations detailed in the first part, and such an interpretation has existed in the literature, in various forms, at the very least since Bell's 1987 paper on the so-called *Lorentzian pedagogy*, although it can arguably trace its lineage, in part at least, as far back as Einstein himself.[12] This line of thought culminates in Brown and Pooley's *dynamical approach*. The formalism of relativity theory turns out to be a self-repairing organism.

[12] Brown cites the following passage from Einstein's 1949 *Autobiographical Notes*: 'But one must not legalize the mentioned sin [of treating rods and clocks as primitive] so far as to imagine that intervals are physical entities of a special type, intrinsically different from other variables ('reducing physics to geometry', etc.)' [41].

Bibliography

[1] Pablo Acuña. Minkowski spacetime and Lorentz invariance: The cart and the horse or two sides of a single coin? *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 55:1–12, 2016.

[2] Dharam Vir Ahluwalia and C Burgard. Gravitationally induced neutrino-oscillation phases. *General Relativity and Gravitation*, 28(10):1161–1170, 1996.

[3] Jeeva Anandan. On the hypotheses underlying physical geometry. *Foundations of Physics*, 10(7-8):601–629, 1980.

[4] Jeeva Anandan. Remarks concerning the geometries of gravity and gauge fields. In B. L. Hu, M. P. Ryan, and C. V. Vishveshwara, editors, *Directions in General Relativity*, pages 10–20. Cambridge University Press, New York, 1993.

[5] James L Anderson. *Principles of relativity physics*. Academic Press, New York, 1967.

[6] Frank Arntzenius. *Space, Time and Stuff*. Oxford University Press, Oxford, 2012.

[7] Felipe A Asenjo and Sergio A Hojman. Do electromagnetic waves always propagate along null geodesics? *Classical and Quantum Gravity*, 34(20):205011, 2017.

[8] David John Baker. On spacetime functionalism. MS, available at: http://philsci-archive.pitt.edu/14301/, 2018.

[9] Yuri Balashov and Michel Janssen. Presentism and relativity. *The British journal for the philosophy of science*, 54(2):327–346, 2003.

[10] Jeffrey A Barrett. *The quantum mechanics of minds and worlds*. Oxford University Press, Oxford, 1999.

[11] Katrin Becker, Melanie Becker, and John H Schwarz. *String theory and M-theory: A modern introduction*. Cambridge University Press, 2006.

[12] Jacob D Bekenstein. An alternative to the dark matter paradigm: relativistic mond gravitation. *arXiv preprint astro-ph/0412652*, 2004.

[13] JD Bekenstein. Relativistic gravitation theory for the MOND paradigm. *Physical Review D*, 70:083509, 2004.

[14] John S Bell. *How to teach special relativity*, page 61. Cambridge University Press, Cambridge, 1987.

[15] Gordon Belot. Geometry and motion. *British Journal for the Philosophy of Science*, 51(4):561–595, 2000.

[16] Gordon Belot. *Geometric possibility*. Oxford University Press, Oxford, 2011.

[17] Carolyn Brighouse. Spacetime and holes. In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume 1994, pages 117–125. Philosophy of Science Association, 1994.

[18] Harvey R Brown. *Physical Relativity*. Oxford University Press, Oxford, 2005.

[19] Harvey R Brown and Oliver Pooley. The origin of the space-time metric: Bell's 'lorentzian pedagogy'and its significance in general relativity. In Craig Callender and Nick Huggett, editors, *Philosophy Meets Physics at the Planck Scale*. Cambridge University Press, Cambridge, 2001.

[20] Harvey R Brown and Oliver Pooley. Minkowski space-time: a glorious non-entity. *Philosophy and Foundations of Physics*, 1:67–89, 2006.

[21] Harvey R Brown and James Read. Clarifying possible misconceptions in the foundations of general relativity. *American Journal of Physics*, 84(5):327–334, 2016.

[22] Harvey R Brown and James Read. The dynamical approach to spacetime. In Eleanor Knox and Alastair Wilson, editors, *The Routledge Companion to Philosophy of Physics*. Routledge, forthcoming.

[23] Ioseph L Buchbinder and Sergei M Kuzenko. *Ideas and methods of supersymmetry and supergravity*. Institute of Physics Publishing, Bristol, 1998.

[24] Jeremy Butterfield. The hole truth. *The British journal for the philosophy of science*, 40(1):1–28, 1989.

[25] Jeremy Butterfield. Against pointillisme about geometry. *arXiv preprint physics/0512063*, 2005.

[26] Jeremy Butterfield. Against pointillisme about mechanics. *The British journal for the philosophy of science*, 57(4):709–753, 2006.

[27] Jeremy Butterfield. Reconsidering relativistic causality. *International Studies in the Philosophy of Science*, 21(3):295–328, 2007.

[28] Jeremy Butterfield. Against pointillisme: a call to arms. In *Explanation, Prediction, and Confirmation*, pages 347–365. Springer, 2011.

[29] Craig Callender. *What Makes Time Special?* Oxford University Press, Oxford, 2017.

[30] Sean M Carroll. *Spacetime and geometry. An introduction to general relativity*, volume 1. Addison Wesley, Boston, 2004.

[31] Raymond Chiao. Superluminal phase and group velocities: A tutorial on sommerfeld's phase, group, and front velocities for wave motion in a medium, with applications to the" instantaneous superluminality" of electrons. *arXiv preprint arXiv:1111.2402*, 2011.

[32] David Deutsch. Quantum theory of probability and decisions. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 455, pages 3129–3137. The Royal Society, 1999.

[33] Neil Dewar. *Symmetries in physics, metaphysics and logic*. PhD thesis, University of Oxford, 2016.

[34] Bryce DeWitt. *Supermanifolds*. Cambridge University Press, 1992.

[35] Paul AM Dirac. *General theory of relativity*. Princeton University Press, Princeton, NJ, 1975.

[36] John Earman. *World Enough and Space-Time: Absolute versus Relational Theories of Space and Time*. MIT Press, Cambridge, MA, 1989.

[37] John Earman. *Bangs, crunches, whimpers, and shrieks: Singularities and acausalities in relativistic spacetimes*. Oxford University Press, 1995.

[38] John Earman and John Norton. What price spacetime substantivalism? the hole story. *British Journal for the Philosophy of Science*, pages 515–525, 1987.

[39] Jürgen Ehlers and Robert Geroch. Equation of motion of small bodies in relativity. *Annals of Physics*, 309(1):232–236, 2004.

[40] Albert Einstein. On the electrodynamics of moving bodies. *Annalen der Physik*, 17(10):891–921, 1905.

[41] Albert Einstein. Autobiographical notes. In Paul Arthur Schilpp, editor, *Albert Einstein: Philosopher-Scientist*, volume 1, pages 1–94. Harper and Brothers, New York, 1949.

[42] Hartry Field. Can we dispense with space-time? In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume 1984, pages 33–90. Philosophy of Science Association, 1984.

[43] Kit Fine. Ontological dependence. In *Proceedings of the Aristotelian society*, volume 95, pages 269–290. JSTOR, 1995.

[44] Samuel C Fletcher. Light clocks and the clock hypothesis. *Foundations of Physics*, 43(11):1369–1383, 2013.

[45] Michael Friedman. *Foundations of space-time theories: Relativistic physics and philosophy of science.* Princeton University Press, Princeton, 1983.

[46] Michael Friedman. *Dynamics of reason.* Csli Publications Stanford, 2001.

[47] Mathias Frisch. Principle or constructive relativity. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 42(3):176–183, 2011.

[48] JC Garrison, MW Mitchell, RY Chiao, and EL Bolda. Superluminal signals: causal loop paradoxes revisited. *arXiv preprint quant-ph/9810031*, 1998.

[49] Robert Geroch and Pong Soo Jang. Motion of a body in general relativity. *Journal of Mathematical Physics*, 16(1):65–67, 1975.

[50] Robert Geroch and James Owen Weatherall. The motion of small bodies in space-time. *arXiv preprint arXiv:1707.04222*, 2017.

[51] Michel Ghins and Tim Budden. The principle of equivalence. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 32(1):33–51, 2001.

[52] Kristina Giesel, Frederic P Schuller, Christof Witte, and Mattias NR Wohlfarth. Gravitational dynamics for all tensorial spacetimes carrying predictive, interpretable, and quantizable matter. *Physical Review D*, 85(10):104042, 2012.

[53] Kurt Gödel. An example of a new type of cosmological solutions of einstein's field equations of gravitation. *Reviews of modern physics*, 21(3):447, 1949.

[54] Herbert Goldstein. *Classical mechanics.* Pearson Education India, New Delhi, 2011.

[55] Richard Healey. *Gauging what's real: The conceptual foundations of contemporary gauge theories*. Oxford University Press on Demand, 2007.

[56] Carl Hoefer. The metaphysics of space-time substantivalism. *The Journal of Philosophy*, 93(1):5–27, 1996.

[57] Nick Huggett. The regularity account of relational spacetime. *Mind*, 115(457):41–73, 2006.

[58] Nick Huggett and Tiziana Vistarini. Deriving general relativity from string theory. *Philosophy of Science*, 82(5):1163–1174, 2015.

[59] Nick Huggett and Christian Wüthrich. Emergent spacetime and empirical (in) coherence. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 44(3):276–285, 2013.

[60] James Isenberg and James Nester. Canonical gravity. In *General Relativity and Gravitation. Vol. 1. One hundred years after the birth of Albert Einstein. Edited by A. Held. New York, NY: Plenum Press, p. 23, 1980*, volume 1, page 23, 1980.

[61] Chris J Isham. *Modern differential geometry for physicists*, volume 61. World Scientific Publishing Company, 1999.

[62] Ted Jacobson and David Mattingly. Gravity with a dynamical preferred frame. *Physical Review D*, 64(2):024028, 2001.

[63] Michel Janssen. Drawing the line between kinematics and dynamics in special relativity. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 40(1):26–52, 2009.

[64] Pankaj S Joshi. Global aspects in gravitation and cosmology. *Int. Ser. Monogr. Phys., Vol. 87,*, 87, 1993.

[65] Tom WB Kibble. Lorentz invariance and the gravitational field. *Journal of mathematical physics*, 2(2):212–221, 1961.

[66] Eleanor Knox. Flavour-oscillation clocks and the geometricity of general relativity. *The British Journal for the Philosophy of Science*, 61(2):433–452, 2010.

[67] Eleanor Knox. Newton–Cartan theory and teleparallel gravity: The force of a formulation. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 42(4):264–275, 2011.

[68] Eleanor Knox. Effective spacetime geometry. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 44(3):346–356, 2013.

[69] Eleanor Knox. Newtonian spacetime structure in light of the equivalence principle. *The British Journal for the Philosophy of Science*, 65(4):863–880, 2013.

[70] Eleanor Knox. Physical relativity from a functionalist perspective. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 2017.

[71] Vincent Lam and Christian Wüthrich. Spacetime is as spacetime does. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, https://doi.org/10.1016/j.shpsb.2018.04.003, 2018.

[72] Marc Lange. Laws and meta-laws of nature: Conservation laws and symmetries. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 38(3):457–481, 2007.

[73] Dennis Lehmkuhl. The Equivalence Principle(s). In Eleanor Knox and Alastair Wilson, editors, *The Routledge Companion to Philosophy of Physics*. Routledge, forthcoming.

[74] David Lewis. Putnam's paradox. *Australasian Journal of Philosophy*, 62(3):221–236, 1984.

[75] David B Malament. *Topics in the foundations of general relativity and Newtonian gravitation theory*. University of Chicago Press, 2012.

[76] John Byron Manchak. What is a physically reasonable space-time? *Philosophy of Science*, 78(3):410–420, 2011.

[77] Tim Maudlin. *Philosophy of physics: Space and time*. Princeton University Press, Princeton, NJ, 2012.

[78] Hermann Minkowski. Space and Time. *Annual Report of the German Mathematician Association, vol. 18, p. 75-88*, 18:75–88, 1909.

[79] Charles W Misner, Kip S Thorne, and John Archibald Wheeler. *Gravitation*. Macmillan, London, 1973.

[80] Wayne C Myrvold. How could relativity be anything other than physical? *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 2017.

[81] Mikio Nakahara. *Geometry, topology and physics*. CRC Press, Boca Raton, FL, 2003.

[82] John Norton. What was einstein's principle of equivalence? *Studies in history and philosophy of science Part A*, 16(3):203–246, 1985.

[83] John D Norton. Why constructive relativity fails. *The British journal for the philosophy of science*, 59(4):821–834, 2008.

[84] Barrett O'Neill. *Semi-Riemannian geometry with applications to relativity*, volume 103. Academic press, 1983.

[85] Thanu Padmanabhan. *Gravitation: foundations and frontiers*. Cambridge University Press, Cambridge, 2010.

[86] Roger Penrose. Structure of space-time. In Cécile DeWitt and John Wheeler, editors, *Batelle Rencontres: 1967 Lectures in Mathematics and Physics*. Benjamin, New York, 1968.

[87] Roger Penrose. *The Emperor's New Mind*. Penguin, New York, 1989.

[88] Roger Penrose and Wolfgang Rindler. *Spinors and space-time Volume 1: two-spinor calculus and relativistic fields*, volume 1. Cambridge university press, Cambridge, 1988.

[89] Volker Perlick. *Ray optics, Fermat's principle, and applications to general relativity*, volume 61. Springer Science & Business Media, 2000.

[90] J Brian Pitts. Space-time constructivism vs. modal provincialism: Or, how special relativistic theories needn't show Minkowski chronogeometry. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 2017.

[91] Joseph Polchinski. *String theory: Volume 1*. Cambridge university press, Cambridge, 1998.

[92] Oliver Pooley. Substantivalist and relationalist approaches to spacetime. In Robert W Batterman, editor, *The Oxford Handbook of Philosophy of Physics*. Oxford University Press, 2013.

[93] Oliver Pooley. Background independence, diffeomorphism invariance, and the meaning of coordinates. In Dennis Lehmkuhl, editor, *Towards a Theory of Spacetime Theories*. Birkhauser, 2017.

[94] Oliver Pooley. *The Reality of Spacetime*. Oxford University Press (Work in progress. Sent in personal communication), (forthcoming).

[95] Dennis Rätzel, Sergio Rivera, and Frederic P Schuller. Geometry of physical dispersion relations. *Physical Review D*, 83(4):044047, 2011.

[96] James Read. Explanation, geometry and conspiracy in relativity theory. In Claus Beisbart, Tilman Sauer, and Christian Wüthrich, editors, *Thinking About Space and Time:100 Years of Applying and Interpreting General Relativity, vol. 15 of the Einstein Studies series*. Birkhauser, Basel, 2018 (forthcoming).

[97] James Read, Harvey R Brown, and Dennis Lehmkuhl. Two miracles of general relativity. *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics*, https://doi.org/10.1016/j.shpsb.2018.03.001, 2018 (forthcoming).

[98] Alice Rogers. *Supermanifolds: theory and applications.* World Scientific, Hong Kong, 2007.

[99] Simon Saunders. Physics and Leibniz's principles. In Katherine Brading and Elena Castellani, editors, *Symmetries in physics: Philosophical reflections*, pages 289–307. Cambridge University Press, 2003.

[100] Simon Saunders. Derivation of the born rule from operational assumptions. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 460, pages 1771–1788. The Royal Society, 2004.

[101] Simon Saunders. Rethinking Newton's Principia. *Philosophy of Science*, 80(1):22–48, 2013.

[102] Erwin Schrödinger. *Space-time structure.* Cambridge University Press, Cambridge, 1985.

[103] Frederic P Schuller and Christof Witte. How quantizable matter gravitates: A practitioner's guide. *Physical Review D*, 89(10):104061, 2014.

[104] Frederic P Schuller and Mattias NR Wohlfarth. Geometry of manifolds with area metric: multi-metric backgrounds. *Nuclear physics B*, 747(3):398–422, 2006.

[105] Bradford Skow. Review of 'Harvey R. Brown: Physical relativity'. *Notre Dame Philosophical Reviews*, 2006.

[106] Bradford Skow. What makes time different from space? *Noûs*, 41(2):227–252, 2007.

[107] Syman Stevens. The dynamical approach as practical geometry. *Philosophy of Science*, 82(5):1152–1162, 2015.

[108] Patrick Suppes. A comparison of the meaning and uses of models in mathematics and the empirical sciences. In *The concept and the role of the model in mathematics and natural and social sciences*, pages 163–177. Springer, 1961.

[109] J L Synge. *Relativity: The General Theory.* North-Holland Publishing Co., Amsterdam, 1960.

[110] Kip S Thorne and Roger D Blandford. *Modern Classical Physics: Optics, Fluids, Plasmas, Elasticity, Relativity, and Statistical Physics.* Princeton University Press, 2017.

[111] Roberto Torretti. *Relativity and Geometry: Foundations and Philosophy of Science and Technology Series.* Elsevier, 2014.

[112] Andrzej Trautman. Fiber bundles, gauge fields, and gravitation. In A Held, editor, *General Relativity and Gravitation.* Plenum Press, New York, 1980.

[113] Bas C Van Fraassen. *The scientific image.* Oxford University Press, Oxford, 1980.

[114] J. Viaclovsky. Topics in Riemannian geometry. University of Wisconsin lecture notes.

[115] Vladimir P Vizgin. *Unified Field Theories: in the first third of the 20th century*, volume 13. Springer Science & Business Media, 2011.

[116] Robert M Wald. *General relativity.* University of Chicago press, 2010.

[117] David Wallace. Everettian rationality: defending deutsch's approach to probability in the everett interpretation. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 34(3):415–439, 2003.

[118] David Wallace. *The emergent multiverse: Quantum theory according to the Everett interpretation.* Oxford University Press, Oxford, 2012.

[119] David Wallace. Fields as bodies: a unified presentation of spacetime and internal gauge symmetry. *arXiv preprint gr-qc 1506.03512*, 2015.

[120] David Wallace. Who's afraid of coordinate systems? an essay on the representation of spacetime structure. *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics*, 2017.

[121] David Wallace. Fundamental and emergent geometry in Newtonian physics. *British Journal for the Philosophy of Science*, forthcoming.

[122] James Owen Weatherall. Fiber bundles, Yang–Mills theory, and general relativity. *Synthese*, 193(8):2389–2425, 2016.

[123] James Owen Weatherall. Conservation, inertia, and spacetime geometry. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 2017.

[124] S Weinberg. *The Quantum Theory of Fields vol. III: Supersymmetry.* Cambridge University Press, Cambridge, 2000.

[125] Steven Weinberg. *Gravitation and cosmology: principles and applications of the general theory of relativity*, volume 1. Wiley New York, 1972.

[126] Steven Weinberg. *Dreams of a final theory*. Vintage, New York, 1994.

[127] Julius Wess and Jonathan Bagger. *Supersymmetry and supergravity*. Princeton university press, Princeton, 1992.

[128] Barton Zwiebach. *A first course in string theory*. Cambridge university press, Cambridge, 2004.