

Once Upon a Spacetime

by

Bradford Skow

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Philosophy
New York University
September, 2005

Hartry Field

© Bradford Skow
All Rights Reserved, 2005

Acknowledgements

Thanks to my advisors Gordon Belot, Cian Dorr, and Hartry Field, for their comments, suggestions, and criticisms, and for providing models of what good philosophy can and should be. Thanks to my parents—all of them—for supporting me through graduate school. And special thanks to Peter Graham for the frequent meetings and for the discussion that led to the writing of chapter 4.

Abstract

This dissertation concerns the nature of spacetime. It is divided into two parts. The first part, which comprises chapters 1, 2, and 3, addresses ontological questions: does spacetime exist? And if so, are there any other spatiotemporal things? In chapter 1 I argue that spacetime does exist, and in chapter 2 I respond to modal arguments against this view. In chapter 3 I examine and defend supersubstantivalism—the claim that all concrete physical objects (tables, chairs, electrons and quarks) are regions of spacetime.

Four-dimensional spacetime, we are often told, ‘unifies’ space and time; if we believe in spacetime, then we do not believe that space and time are separately existing things. But that does not mean that there is no distinction between space and time: we still distinguish between the spatial aspects and the temporal aspects of spacetime. The second part of this dissertation, comprising chapter 4, looks at this distinction. How is it made? In virtue of what are the temporal aspects of spacetime temporal, rather than spatial? The standard view is that the temporal aspects of spacetime are temporal because they play a distinctive role in the geometry of spacetime. I argue that this view is false, and that the temporal aspects are temporal because they play a distinctive role in the geometry of spacetime and in the laws of nature.

Table of Contents

Acknowledgements	iv
Abstract	v
Table of Contents	vi
List of Figures	ix
1 An Argument for Substantivalism	1
1 The Standard Argument	1
2 Characterizing the Spatiotemporal Structure of the World	4
3 The Embeddability Criterion	7
4 Machian Spacetime	9
5 Account 1: One Mixed Fundamental Distance Relation	11
6 Arguing for PURITY	11
6.1 Clarifying the Premises	13
6.2 Defending the Premises	15
7 Account 2: Pure Distance Relations	18
8 PURITY and Substantivalism	20
9 Modal Relationalism	22
2 Modal Arguments against Substantivalism	27
1 Introduction	27
2 Defending The First Premise	29
2.1 The Leibniz Shift Argument	29

2.2	Back to the Hole Argument	34
2.3	Confusions About The First Premise	36
2.4	Defending the First Premise	38
2.5	Counterpart Theory	42
3	The Second Premise	43
3.1	Determinism: Examples and Counterexamples	43
3.2	Analyzing Determinism	45
3.3	A Formal Definition	48
4	Earman's and Stein's Argument	51
5	Conclusion	52
3	Supersubstantivalism	54
1	Introduction	54
2	Varieties of Supersubstantivalism	55
3	Is Dualism the Default View?	57
4	Cartesian Supersubstantivalism	59
5	Spatiotemporality	62
6	Newtonian Supersubstantivalism	64
7	Moderate Supersubstantivalism	67
8	The Modal Argument Against Supersubstantivalism	69
4	What Makes Time Different From Space?	72
1	Introduction	72
2	Spacetime Diagrams and Directions in Spacetime	73
3	Does Geometry Distinguish Temporal from Spatial Directions? . . .	78
4	Dimensionality	80
5	Laws of Nature	85
6	Testing Our Intuitions	88
7	Symmetric Laws and Indeterminacy	92
8	Laws That Govern in a Spacelike Direction	94
9	Alternative Views	98
9.1	The Causal Theory of Time	98
9.2	Three-Dimensionalism	99

9.3	The Brute Fact View	100
	Bibliography	103

List of Figures

1.1	22
4.1	74
4.2	75
4.3	Newtonian Spacetime	76
4.4	Minkowski Spacetime	82
4.5	88

Chapter 1

An Argument for Substantivalism

1 The Standard Argument

Distinguish relationalism about motion from relationalism about ontology. Relationalists about ontology deny that spacetime exists. In doing so they oppose substantivalists, who affirm that spacetime exists.¹ (When I use it without qualification, ‘relationalism’ means the same as ‘relationalism about ontology.’) Relationalists about motion assert that all motion is the relative motion of bodies. In doing so they oppose absolutists, who affirm that there are some states of motion that are not states of motion relative to this or that material object.

I will work in the context of pre-relativistic physics. The standard argument for substantivalism in this context, going back to Newton, has two premises:

¹Some philosophers claim that relationalists do not deny the existence of spacetime, but merely assert that spacetime is a logical construction from the spatiotemporal relations among material objects. (Forbes (1987) is an example.) I’ve always thought that to say that x ’s are logical constructions is just another way to say that x ’s don’t exist. Speaking in terms of ‘logical constructions’ obscures the debate.

- (1) Relationalism about motion is false.
- (2) If relationalism about motion is false, then spacetime exists.

To argue for (1), substantivalists claim that no adequate physical theory is consistent with relationalism about motion. Newtonian gravitational theory has a well-posed initial value problem: there is a unique future evolution from any complete state of the world at a time.² But no competing theory of gravitation that is consistent with relationalism about motion has a well-posed initial value problem. And so, given the availability of Newton's theory, no relationalist theory can be adequate.

The argument that no theory that is consistent with relationalism about motion has a well-posed initial value problem has premises. First we suppose (i) that our world, and all physically possible worlds, are worlds of massive point particles. Then if relationalism about motion is true, the history of the world is a history of inter-particle distances. We also suppose (ii) that Newtonian gravitational theory is empirically adequate and a good guide to physical possibility. That is, we suppose that one of the solutions to the equations of Newtonian gravitational theory captures the actual history of inter-particle distances; and that each solution captures a physically possible history of inter-particle distances. Finally we suppose (iii) that if relationalism about motion is true, then to give the complete state of the world at a time is to specify the distances between each pair of particles and the velocity of each particle relative to each other.

Given what we have supposed, there can be no alternative to Newtonian gravitational theory that both is consistent with relationalism about motion and has a well-posed initial value problem. For any complete state of the world at a time—for any specification of the distances between and relative velocities of each pair of particles—there is more than one physically possible future evolution. Here is an example (from Newtonian gravitational theory supplemented with Hooke's law). Consider a (physically) possible world that contains just two lead balls connected by a spring. At some given time the distance between them is one meter and the rate at which this distance is changing is zero. According to Newtonian mechanics,

²Setting aside problems with space invaders (Earman 1986) and collision singularities.

there are at least two physically possible futures: either the two balls are rotating around their center of mass at such a rate that the ‘centrifugal force’ exactly balances the force from the spring, and so they remain at relative rest for all time; or they are not rotating, and so the spring collapses and then expands, and the initial time was at the moment when the balls turn from moving apart to moving back together.³

(Absolutists diagnose this failure as follows: what relationalists about motion call ‘the complete state of the world at a time’ is not complete. There are distinct instantaneous states of the world that the relationalist cannot distinguish. In the example, there are at least two ways to ‘complete’ the relationalist description of the ball-and-spring system, completions that differ over the value of each ball’s absolute velocity.)

That is the argument for (1). The argument for (2) is shorter. If relationalism about motion is false and there are some states of motion that are not states of relative motion, then spacetime exists. For the only way to define these absolute motions is by reference to spacetime. (A particle is undergoing absolute acceleration at a time, for example, just in case its worldline is not tangent to a geodesic at that time.)

Relationalists like Leibniz and Mach claimed that the first part of this argument failed. But no relationalist produced a theory with a well-posed initial value problem. Until recently. Julian Barbour has developed such a theory.⁴

(So where does the argument that there can be no such theory go wrong? There are different ways to develop Barbour’s theory. Different developments reject different parts of the argument. On one way, relationalists reject (iii), and say instead that the complete state of the universe at a time is given by the distances between and relative velocities of all particles and, in addition, a holistic property of the entire universe at that time (holistic because not reducible to the properties of its parts), its angular momentum. (Relationalists who take this approach say some-

³I first heard this example from David Albert.

⁴This theory is discussed in detail in (Barbour 1999), (Belot 2000), (Butterfield 2002), (Pooley and Brown 2002), among other places. Belot notes that theories like this were discovered as early as 1924.

thing about why we do not need to believe in spacetime in order to make sense of angular momentum.) On another way, relationalists reject (ii), that Newton's theory is a good guide to physical possibility. They deny that worlds in which the entire universe is rotating, including one where two lead balls connected by a spring are rotating around their center of mass, are physically possible.)

With theories like Barbour's available, substantivalists are in a weaker position. No longer able to argue for (1) by pointing out that relationalists about motion have no physical theory, they must now argue that absolutist theories like Newton's are better than relationalist theories—and not better because more adequate to the empirical phenomena, for both theories do just as well on that score, but better for some more subtle reason. But it is not clear that the absolutist theory is better: Barbour (1999) and Pooley and Brown (2002) argue that the relationalist theory is better than Newton's because it predicts a phenomenon that in Newton's must be taken as a brute fact: namely, the fact that the universe is not rotating.

To argue for substantivalism I take a different approach. In the debate over the standard argument, it is widely assumed that we need spacetime to make sense of absolute acceleration, but that we do not need spacetime to make sense of the distances between and relative velocities of material objects.⁵ I think this is wrong. Even if relationalism about motion is true, and we can do physics while recognizing only distances between and relative velocities of material objects, we still need spacetime (I will argue) to make sense of the distances between material objects. Relationalism about motion demands substantivalism just as much as its denial. Any adequate characterization of the spatiotemporal structure of the world—even one consistent with relationalism about motion—must make reference to spacetime. Before outlining my argument in any detail, I will explain this terminology.

2 Characterizing the Spatiotemporal Structure of the World

Relationalism about motion is an answer to the question:

(3) What is the spatiotemporal structure of the world?

⁵Sklar (1974) disputes the first half of this claim.

Both substantialists and relationalists must answer this question, though they will not answer it the same way. It is easiest to see what a complete, detailed answer to this question will look like if we look at some substantialist answers first.

Substantialists believe in spacetime. Just as there are many possible geometries that space may have (it may be Euclidean, or hyperbolic, for example), there are many possible geometries that spacetime may have: it may have a neo-Newtonian structure, for example, or a full Newtonian structure. When substantialists disagree about the answer to (3), then, they are disagreeing about which of these geometries the spacetime in our world has.

A substantialist might answer (3) by saying that spacetime has a Newtonian geometry. How might a relationalist answer (3)? She denies that spacetime exists, so she cannot give the same answer. But she does not think that (3) is an empty question, and she might even think that, except for their ontological disagreement, the substantialist's answer is right.

To see the general form a relationalist answer may take, look more closely at a substantialist answer. Suppose some substantialist says that spacetime has a Newtonian geometry. We may ask: in virtue of what does it have that geometry? And the answer is: it has that geometry because the points of spacetime instantiate certain spatiotemporal relations in a certain pattern. Which spatiotemporal relations? There are many possible answers to this question. Here is one: spacetime has a Newtonian geometry because the points of spacetime instantiate *the spatial distance between x and y is r* and *the temporal interval between x and y is r* in a certain pattern. (That is, there are geometrical laws that these relations satisfy.)

A relationalist who says that the world has a Newtonian spatiotemporal structure might say, then, that it has this structure (at least in part) because the material objects, rather than the points of spacetime, instantiate certain relations in a certain pattern. (Maybe the relationalist uses the same relations the substantialist does; maybe not.)

So both substantialists and relationalists produce lists of spatiotemporal relations and laws governing those relations in order to characterize the spatiotemporal structure of the world. Now, for any given spatiotemporal structure the world might have there are many different sets of relations that characterize it. Euclidean space,

for example, may be characterized using the relation *the distance from x to y is r*; or using the relations *y is between x and z* and *x and y are as far apart as z and w*; and there are other choices as well. But I am interested in which relations are, according to substantialists and to relationalists, the fundamental relations that characterize it. Fundamental relations, I assume, have the following feature: facts about the instantiation of non-fundamental relations obtain in virtue of the facts about the instantiation of the fundamental relations.⁶ Facts about the instantiation of the fundamental relations, on the other hand, are brute, ‘bottom-level’ facts. (The non-fundamental relations, then, supervene on the fundamental relations. Since it is also true that the fundamental relations supervene on themselves, *every* relation supervenes on the fundamental relations.)

Now I can outline my argument in more detail. I will argue that even if the correct answer to (3) entails relationalism about motion, relationalist accounts of which spatiotemporal relations are fundamental are objectionable. So we should prefer substantialism to relationalism.⁷

In my argument I impose a constraint on which spatiotemporal relations may be fundamental. To state this constraint I need to introduce some terminology. Say that relations which relate abstract objects to concrete objects (*x is n years old* is an example) are mixed relations; other relations—relations that relate only abstract objects, or only concrete objects—are unmixed, or pure. The constraint is this:

PURITY: Necessarily, the fundamental spatiotemporal relations are pure.⁸

⁶Here and throughout I limit this claim to qualitative, non-modal relations.

⁷Again, I am working in the context of pre-relativistic physics. But the arguments generalize to relativistic spacetime structures.

⁸If one believes that all possible relations exist necessarily and that necessarily, if a relation is fundamental then it is necessarily fundamental, then the modal operator at the beginning of PURITY is redundant. In what follows I will make these assumptions.

PURITY presupposes that some fundamental relations are spatiotemporal relations. Leibniz denied this; resemblance nominalists (who think that the fundamental relations are relations of resemblance) presumably deny this; some advocates of string theory or of some other theory of quantum gravity may deny this. My arguments need not presuppose that such views are false. Even if no fundamental relation is

Here is an outline of the rest of this chapter. In section 3 I state a necessary condition on relationalist characterizations of the spatiotemporal structure of the world, and in section 4 I explain in detail what the spatiotemporal structure that corresponds to relationalism about motion looks like and how it differs from Newtonian and neo-Newtonian spacetime. In section 5 I look at the most obvious relationalist characterization of this spatiotemporal structure and reject it because it is inconsistent with PURITY. I argue for PURITY in section 6, and in the remainder of the chapter I look at relationalist accounts of the fundamental spatiotemporal relations that are compatible with PURITY. I argue that they are unacceptable.⁹

3 The Embeddability Criterion

Suppose a relationalist answers (3) by saying that the world has a Newtonian spatiotemporal structure—not a likely answer, but this is just an example—and that he tells us a story about the fundamental facts in virtue of which it has this structure (a story that, as I said, will include a list of the fundamental spatiotemporal relations that material objects instantiate and laws that those relations obey). I claim that this relationalist story is correct only if it guarantees that relationalist worlds in which it is true are uniquely embeddable in Newtonian spacetime, up to the symmetries of that spacetime. (And similarly for other possible spatiotemporal structures the world might have.) More precisely: let \mathbb{R}^4 with the spatial distance between any two points (x_1, x_2, x_3, x_4) and (y_1, y_2, y_3, y_4) given by the standard Euclidean formula $d_s(x, y) = \sqrt{(x_2 - y_2)^2 + (x_3 - y_3)^2 + (x_4 - y_4)^2}$ and temporal distance given by $d_t(x, y) = |x_1 - y_1|$ be our canonical model of Newtonian spacetime geometry. Then the relationalist laws are correct only if there is a one-to-one function f from spatiotemporal, still (I take it) there are some spatiotemporal relations that are *most* fundamental: no spatiotemporal relation is more fundamental than they are. PURITY may then be cast as a principle about the most fundamental spatiotemporal relations. In the body of this paper I will ignore this complication, since nothing turns on it.

⁹The argument I give here is similar to the one Field gives in ‘Can We Dispense with Spacetime?’ (1989). While I think that much of what I say is in the spirit of the arguments Field gives, there are important differences between our arguments. I will mention these differences in the relevant places.

the point-sized things in the relationalist world into \mathbb{R}^4 such that

- f preserves spatial and temporal distance: for any point-sized instantaneous¹⁰ things x and y in the relationalist world, $d_s(x, y) = d_s(f(x), f(y))$ and $d_t(x, y) = d_t(f(x), f(y))$, and
- every other such function differs from f by a symmetry of Newtonian spacetime.

(I use ‘ d_s ’ to name both the spatial distance function defined on the material objects in the relationalist world and the spatial distance function defined on points of \mathbb{R}^4 .) The symmetries of Newtonian spacetime are the one-to-one functions from spacetime onto itself that preserve spatial and temporal distance. There are similar definitions of symmetry and embeddability for other spacetimes. (Not all of these definitions need be given in terms of a spatial and temporal distance function.)

Of course, the relationalist’s story may not explicitly mention a spatial or temporal distance function. But these functions must be definable from the story he does tell, for certainly there are facts about the spatial and temporal separation of any two point-sized instantaneous objects in a Newtonian relationalist world (given a choice of units of measurement), whether or not those facts are fundamental.¹¹

The embeddability criterion is accepted by many who discuss the debate between relationalists and substantivalists.¹² I take it that it needs little defense. Clearly, if a relationalist says that a world has a Newtonian spatiotemporal structure, but there is no way to embed that world in Newtonian spacetime, then he is

¹⁰Relationalism is easiest to defend if is conjoined with the doctrine that every material object is composed of point-sized instantaneous parts. (So this doctrine entails the doctrine of temporal parts.) If this doctrine is true, then to specify the spatiotemporal relations among the material objects in the world it is enough to specify the spatiotemporal relations among all the point-sized instantaneous material objects.

¹¹In fact it is possible to formulate the arguments in this paper in terms of the definability of a Euclidean spatial distance function in relationalist worlds, instead of in terms of the embeddability of (parts of) relationalist worlds into Euclidean space. I think framing the arguments in terms of embeddability makes the issues clearer.

¹²(Belot 2000), (Earman 1989), and (Friedman 1983) are three examples.

lying. And if the embedding is not unique up to a symmetry of the spacetime, then the relationalist's story underdetermines the spatial relations among the material objects. (For example, if the relationalist says that the world has a Euclidean spatial structure, but there are three instantaneous pointsized material objects, all simultaneous, and there are two embeddings of them into Euclidean space, one which maps them to the vertices of an equilateral triangle, and one which maps them to the vertices of a triangle that is not equilateral, then the relationalist has not adequately characterized the spatial structure of the world.)

Embeddability is a necessary condition that a relationalist characterization of the spatiotemporal structure of the world must meet. But it is not sufficient. There are two commonly cited obstacles to its sufficiency.¹³ First, any relationalist world that is embeddable in a four-dimensional Newtonian spacetime is also embeddable in a five-dimensional (or higher) Newtonian spacetime with an extra spatial dimension. The embeddability criterion alone does not fix the dimensionality of space. Second, any relationalist world with (as a substantialist would say) large enough empty regions of spacetime that is embeddable in a four-dimensional Newtonian spacetime is also embeddable in a four-dimensional spacetime in which space is Euclidean in all the occupied regions, but is non-Euclidean in some of the empty regions. So the embeddability criterion alone does not always fix the geometry of unoccupied regions of spacetime.

I mention these problems with treating the embeddability criterion as sufficient only to set them aside. I will only appeal to its status as a necessary condition.

4 Machian Spacetime

To say that all motion is relative motion of bodies is to say something about the spatiotemporal structure of the world. I can now be more specific about just what this structure looks like. (To do so I speak as a substantialist.) Earman (1989) lists six possible (non-relativistic) spacetimes and orders them by how much structure they have. All the spacetimes Earman lists consist of a stack of three-dimensional Euclidean instantaneous spaces. The order in which they are stacked gives their

¹³(Earman 1989).

ordering in time. Machian spacetime is the weakest of these: *all there is* to that spacetime is a stack of three-dimensional Euclidean instantaneous spaces. We can say how far apart in space points on the same instantaneous space are, but not how far apart in space or in time points on different instantaneous spaces are. So in Machian spacetime the only meaningful questions are questions about how far apart particles are at any instant, and qualitative (not quantitative) questions about their relative motion: whether two particles are getting closer together, for example. Machian spacetime, then, is a spacetime in which all motion is the relative motion of bodies.¹⁴ Barbour designed his relationalist replacement for Newtonian mechanics to live in a Machian world. (By contrast, Newtonian spacetime has more structure: it comes with a way to determine how far apart in both space and time points on different instantaneous spaces are. In this spacetime, we can tell which trajectories through spacetime are trajectories of particles at absolute rest, and we can ask whether a single particle is in absolute motion, and if so what its numerical speed is.)

Can relationalists say what it is for the world to have a Machian spatiotemporal structure?

To guarantee the embeddability of the world's particles into Machian spacetime relationalists must guarantee that each 'time-slice' of the world (each equivalence class of the world's point-sized instantaneous material objects under the simultaneity relation) is embeddable in three-dimensional Euclidean space. It might seem at first like this is easy to do: all a relationalist needs to do specify the spatial distances between the particles in that time-slice, and ensure that those distances obey Euclidean laws. Then the time-slice will be uniquely embeddable into Euclidean space.

But what are the fundamental spatial relations that give the distances between the particle-slices? There are several choices here, and the problems relationalists face do not come out until we are more clear about which choice we have in mind.

¹⁴It is not the only such spacetime; Leibnizian spacetime, which differs from Machian spacetime only by having a temporal metric, is also a spacetime in which all motion is the relative motion of bodies.

5 Account 1: One Mixed Fundamental Distance Relation

One account of the fundamental spatial relations a relationalist might use to account for the distances between particles is this: say that there is one fundamental spatial relation, *the distance from x to y is r* . Then state laws for this relation that guarantee the embeddability (up to uniqueness) into Euclidean space of the particle-slices instantiating it. I suspect that when relationalists think that a Machian spatiotemporal structure is perfectly acceptable they do so because they think that this is all that need be done.

The problem with this account is that it is incompatible with PURITY:

PURITY: Necessarily, the fundamental spatiotemporal relations are pure.

It is now time to argue for this principle.

6 Arguing for PURITY

There is one group of philosophers who will find PURITY appealing: nominalists, those who deny that there are any abstract objects.¹⁵ For if PURITY is false, because (say) *the distance from x to y is r* is the only fundamental spatial relation, and if in addition there are no numbers, then no spatial relations are instantiated, and so the world is not spatial at all. Since nominalists do not think their view entails that spatiality is an illusion, they will embrace PURITY.

But I think that even anti-nominalists should accept PURITY. I myself find the following argument convincing: even granting that there are numbers, the following counterfactual is true:

- (4) If there were no numbers, the world would still be spatial, and in fact the spatial structure of the world would be just as it actually is.

But if PURITY is false because *the distance from x to y is r* is the only fundamental spatial relation, then (as before) if there were no numbers no spatial relations would

¹⁵Some nominalists deny that there are any properties or relations. They will deny that PURITY is literally true. But nominalists should have some fictionalist reading of sentences containing property talk, and they will accept PURITY when read that way.

be instantiated. And any world in which no spatial relations are instantiated is a world that is not spatial at all. And that means that (4) is false.¹⁶

Many anti-nominalists will be unconvinced by this argument, though. They will say that numbers exist necessarily, and so that (4) is a counterfactual with a necessarily false antecedent. So either (4) makes no sense at all, or it is vacuously true; in either case, there is no problem for an anti-nominalist who accepts PURITY.

So I will take a different approach. My argument for PURITY is this:

(P1) If PURITY is false, then the way the world's particles are spatially arranged at any one time is not intrinsic to the particles.

(P2) The way the world's particles are spatially arranged at any one time *is* intrinsic to the particles.¹⁷

(C) So PURITY is true.¹⁸

This argument is addressed to relationalists. I mean to convince relationalists to accept its conclusion. (I think substantivalists should also accept its conclusion, but for different reasons.¹⁹) I do not claim that its premises are true (its second premise in particular), just that relationalists should believe them. Before I explain why, I will clarify what the premises mean.

¹⁶If PURITY is false because some, but not all, fundamental spatiotemporal relations are mixed, then it does not follow that if there were no numbers the world would not be spatial. But in this case (4) is still false: if there were no numbers then the spatial structure of the world would not be just as it actually is.

¹⁷From now on I omit the reference to time.

¹⁸Field also tries to motivate something like PURITY by noting that a theory incompatible with PURITY does not give intrinsic explanations of phenomena (1989, pages 192-193). I take it that an intrinsic explanation is one that appeals only to the intrinsic properties of and relations among the things invoked in the explanation; so my focus on the intrinsicness of the way the world's particles are arranged is not, at bottom, different from Field's focus on intrinsic explanations. But Field just takes it for granted that (P1) is true. I have found many people willing to deny it; I provide an argument for it.

¹⁹Substantivalists should believe that properties characterizing the geometrical structure of spacetime (like the property of having a Newtonian geometrical structure) are intrinsic. Arguments similar to the ones I give below show that if PURITY is false, then these properties are not intrinsic.

6.1 Clarifying the Premises

I use ‘the way the world’s particles are spatially arranged’ as a name of a relation. I take it that talk of the way some particles are spatially arranged is familiar: one way for three particles to be spatially arranged, for example, is for any two of them to be as far apart as any other two. Then they stand at the vertices of an equilateral triangle. In this case, then, ‘the way the particles are spatially arranged’ names the three-place relation *x and y are as far apart as y and z, and x and y are as far apart as x and z*. (P2) is the claim that relations like this are intrinsic. But what does it mean to say that a relation is intrinsic?

We are familiar with the distinction between intrinsic and extrinsic properties. Intuitively, an intrinsic property is a property that characterizes something as it is in itself. What intrinsic properties something has in no way depends on what other things exist (things other than it or its parts) or how it is related to them. With extrinsic properties (properties that are not intrinsic), by contrast, other things can ‘get in on the act’ when it comes to determining whether something instantiates them.

Although it is unfamiliar, the notion of an intrinsic relation is a straightforward generalization of the notion of an intrinsic property. Just as an intrinsic property characterizes something as it is in itself, an intrinsic relation (with more than one argument place) characterizes some *things* as they are in *themselves*. For example, *x is as massive as y* is intrinsic (or, at least, seems intrinsic at first): if two things are equally massive, that is a matter of how those two things are in themselves. Whether they instantiate *x is as massive as y* does not depend on what else there is or what those other things are like.²⁰

²⁰If *x is as massive as y* is intrinsic, then it is plausible that it is also internal. Internal relations supervene on the intrinsic properties of their relata. More formally, a relation *R* is internal just in case: if *x* and *y* instantiate *R* and *x'* and *y'* are duplicates of *x* and *y* (respectively) then *x'* and *y'* also instantiate *R*. (*x'* and *y'* need not exist in the same possible world as *x* and *y*. The generalization to relations with more than two argument places is straightforward.) All internal relations are intrinsic, but not all intrinsic relations are internal. (Fundamental relations (with more than one argument place) are intrinsic but not internal. Some define ‘external relation’ as ‘relation that is intrinsic but not internal.’ (Lewis (1983) for example.)

Here is another way to think about intrinsic relations. Intrinsic relations usually correspond to intrinsic properties of fusions. Consider the property *x has two parts that are equally massive*. This property is intrinsic. Any intrinsic duplicate of something with two equally massive parts would itself have two equally massive parts. And something instantiates this property just in case it has two parts that instantiate *x and y are equally massive*. And the latter (as I said) is an intrinsic relation. In general, if an n -place relation R is intrinsic then the property of having n parts that instantiate R is also intrinsic.²¹

Using the distinction between fundamental and non-fundamental relations I can give a more precise characterization of ‘intrinsic relation.’ All relations supervene on the fundamental relations. But this is global supervenience: if some things instantiate a relation R , they do so in virtue of the global pattern of instantiation of the fundamental relations. An intrinsic relation, by contrast, supervenes ‘locally’ on the fundamental relations. That is, a relation is intrinsic only if it supervenes on the fundamental properties of, and fundamental relations among, its relata and its relata’s parts.

Local supervenience is necessary for a property to be intrinsic. But if there are any things that exist necessarily it is not sufficient. Suppose God exists necessarily; then the property of coexisting with God supervenes on any set of properties. But it is not intrinsic.

The solution is to demand more: a relation $R_{x_1 \dots x_n}$ is intrinsic just in case it can be *analyzed* in terms of the fundamental relations $x_1 \dots x_n$ and their parts instantiate. An analysis is a kind of ‘definition’ of the non-fundamental relation: to give an analysis is to display an open sentence that expresses that relation (and so has as many free variables as that relation has argument places) such that every predicate in that open sentence expresses a fundamental relation. We can tell whether a relation is intrinsic by looking at its analysis: $R_{x_1 \dots x_n}$ is intrinsic just in case every

The distinction between internal and non-internal relations will not play a role in my arguments.

²¹I believe the converse is true for fundamental relations, but can fail for non-fundamental relations. The (non-fundamental) relation $\exists z (x \text{ and } y \text{ compose } z)$ is an example: it is not an intrinsic relation, but the property of having two parts that compose something is an intrinsic property.

quantifier in its analysis is restricted to $x_1 \dots x_n$ and their parts.²²

6.2 Defending the Premises

I have now explained what the premises mean, and given a more precise characterization of ‘intrinsic relation.’ I turn now to arguing that relationalists should accept the two premises.

Here is my argument for (P1). Suppose PURITY is false. Suppose in addition that there is just one fundamental spatial relation, *the distance from x to y is r* . Then the way the world’s particles are spatially arranged may be analyzed in terms of this relation.

If this is true, then intuitively, the particles get to be spatially arranged in a certain way in virtue of being related in the right way to numbers. So their being arranged in that way is not a matter of the way they are, in themselves; so the way they are arranged is not intrinsic.

To go through this argument in more detail, fix on a particular way the particles might be arranged, and look at its analysis. Suppose there are only three particles and that at a certain time they stand at the vertices of an equilateral triangle. Then the way they are arranged may be analyzed as follows:

$\exists r$ (the distance from x to y is r & the distance from y to z is r & the distance from z to x is r)

This analysis contains a quantifier that ranges over something other than x, y, z , and their parts; namely, a quantifier that ranges over numbers. So the relation analyzed is not intrinsic.

Someone might object that numbers (and abstract objects generally) should receive a special dispensation in the definition of ‘intrinsic.’ That definition bans

²²The definition of ‘intrinsic’ I am working with is related to the one given by David Lewis in *On the Plurality of Worlds* (1986b). It follows from Lewis’s definition that the fundamental properties and relations are intrinsic. So this definition is appealing only if we accept this consequence. But I don’t think this consequence should be controversial. It is part of our conception of fundamental properties that they are intrinsic: it is usually said that the fundamental properties make for similarity among their instances, and that they carve nature at its joints (Lewis 1983).

quantifiers not restricted to the parts of the things instantiating the relation from analyses of intrinsic relations. But perhaps it should allow quantifiers that range over only abstract objects into such analyses; so long as no quantifier ranges over *concrete* things other than the parts of the things instantiating the relation, the relation is intrinsic. (The above analysis can easily be rewritten so that the first quantifier is restricted to numbers.)

We should reject this revised definition of ‘intrinsic.’ To illustrate why it is wrong, notice that it leads to incorrect results about the intrinsic properties of numbers themselves. If we suppose that numbers exist, then it is plausible to suppose that *x is the successor of y* is a fundamental relation that natural numbers instantiate. (Other relations among the natural numbers—like addition and multiplication—can be defined in terms of successor.) Now the property of being the smallest natural number, like the property of being the shortest person in the room, is not intrinsic. Whether a number has this property depends on what other numbers there are, and whether it is the successor of any of them. But this property’s analysis—‘ $\neg\exists y(y$ is a number and x is the successor of y)’—contains just one quantifier restricted to numbers, and so the revised definition of ‘intrinsic’ says that it is intrinsic.

In this argument for (P1) I’ve assumed that if PURITY is false then *the distance from x to y is r* is the fundamental spatial relation. It should be clear from my discussion that nothing turns on using this particular mixed relation. The argument works equally well no matter which mixed relation we choose. This is important because, if one were going to choose a mixed spatial relation to regard as a fundamental relation, one would not be likely to choose *the distance from x to y is r*. It has seemed to many that this relation isn’t really fundamental, but is analyzed in terms of some other fundamental mixed spatial relations. For example, some have thought that the fundamental spatial relation concerns distance ratios, rather than distances. That is, there is just one five-place fundamental spatial relation, *x and y are r times as far apart as z and w*. Distances between points are then analyzed in terms of the ratios of their distances to the distance between two reference objects (say, the ends of the standard meter). This theory is appealing because it does away with seemingly-mysterious facts about which things are distance one unit apart. (We are very good at determining when two things are one meter apart, and when they are one foot

apart; but these facts look like facts about the ratio between the distance between the two things in question and the distance between the endpoints of the standard meter, or the standard foot. Asked to determine whether two things are one unit apart, in some absolute sense that is independent of any standard of measurement, and we do not know what to do.)²³ But for all this view's virtues, my argument for (P1) works just as well if it (or some other theory of mixed fundamental spatial relations) is true.

So much for (P1). What about (P2)? Not everyone will accept (P2). Some substantialists, in particular, will reject it. For some substantialists say that the way the world's particles are spatially arranged is derived from facts about where in space each particle is located.²⁴ The way the particles are arranged, then, is not just a matter of how those particles are, in and of themselves, but is also a matter of the way they are related to space. So their spatial arrangement is not intrinsic.

I claim that relationalists should accept (P2). For (P2) is, I think, one of the claims that motivates relationalism in the first place. Distinguish two motivations for relationalism. (There may be others.) One way to motivate relationalism is to complain that points of space and instants of time are unobservable entities, and assert that for that reason we should not believe in them. But another (and better) way to motivate relationalism is to complain that space and time are irrelevant and unnecessary. As I said in the previous paragraph, one role that space plays for substantialists is to 'ground' spatial relations among material objects. I may be a mile above Los Angeles; according to substantialists, what it is for me to be a mile above Los Angeles is for me to be located at a certain region of space, and Los Angeles to be located at a certain region of space, and for those regions to be a mile apart. That is, substantialists say that material objects inherit their spatial relations from the spatial relations among the regions of space they occupy. But this detour through spatial relations among regions of space looks unnecessary. Why can't the

²³After reading Kripke (1980) you might doubt that '*x and y are one meter apart*' expresses the relation *x and y are as far apart as a and b* where *a* and *b* are the end points of the standard meter. I think the view that best fits with Kripke's picture is account 2 (section 7 below; see also footnote 25).

²⁴Not all substantialists say this: supersubstantialists, for example, do not. Supersubstantialism is the subject of chapter 3.

‘one mile apart’ relation hold directly between me and Los Angeles? Why does it need to be derived from spatial relations among regions of space? And if it does not need to be so derived, then space is irrelevant. The way the world’s particles are spatially arranged is not a matter of the way they are embedded in space; instead, it is a matter of how those particles are, in and of themselves. It is intrinsic to the particles. Here, for example, is Barbour expressing this motivation:

What is the reality of the universe? It is that in any instant the objects in it have some relative arrangement. If just three objects exist, they form a triangle. In one instant the universe forms one triangle, in a different instant another. What is to be gained by supposing that either triangle is placed in invisible space? (1999, pp.68-69).

(Not for nothing does Barbour call his theory ‘intrinsic particle dynamics.’) I think this, and not appeals to problems with unobservable entities, is the best motivation for relationalism. It is a motivation that I myself feel, even though I am no relationalist. But—what is important for my purposes—this line of thought motivates relationalism by appealing to (P2). So relationalists should accept (P2).

There is, of course, another well-known motivation for relationalism: the Leibniz shift argument. But I do not think that this motivation is very good. (I discuss this argument, and its successor the hole argument, in more detail in the next chapter.)

I have argued that everyone should accept (P1) and relationalists should accept (P2). So relationalists should accept (C): they should accept PURITY.

Let me be clear about what this means. If numbers do exist, then material objects surely do bear mixed relations like *the distance from x to y is r* to numbers. I only claim that it cannot be relations like these that *give* the world its spatial structure. When it comes to the fundamental facts about spatial structure, numbers are strictly irrelevant.

7 Account 2: Pure Distance Relations

Relationalists cannot characterize the spatial structure of the world by saying that the three-place spatial relation *the distance from x to y is r* is fundamental. Here

is an alternative theory: say that there are infinitely many two-place fundamental spatial relations. These relations have names like ‘*the distance from x to y is one,*’ ‘*the distance from x to y is two,*’ and so on. But we should not be misled by the names: although one relation’s name contains the letters ‘o-n-e,’ we are not to think of these as a separate syntactic unit that serves to name a number. We are not to think of these pure relations as derived from the three-place relation *the distance from x to y is r* by filling in its third argument place with a number.²⁵

If these relations are fundamental then the way the world’s particles are spatially arranged is intrinsic. So this theory is compatible with PURITY and respects the motivation for relationalism. And so long as the inter-particle distances do not violate Euclidean laws, the world’s time-slices will be uniquely embeddable into Euclidean space.

The problem with this theory is that it just cannot be true. For according to this theory there are necessary truths that we must think are brute and inexplicable, but which we ought to be able to explain. For example, it is necessary that if two things instantiate *the distance from x to y is one* then they do not also instantiate *the distance from x to y is two*. And not only is this necessary; each distance relation excludes the other distance relations in this way. Or again, it is necessary that if a and b instantiate *the distance from x to y is two* and b and c also instantiate *the distance from x to y is two* then a and c do not instantiate *the distance from x to y*

²⁵This seems to be the view Melia prefers at the end of his (Melia 1998). It also fits best with Kripke’s discussion of the standard meter (1980). Kripke suggests that we refer to the standard meter stick only to fix the reference of ‘one meter apart.’ Two things can be one meter apart without bearing any relation to the endpoints of the standard meter; they can be one meter apart even in possible worlds in which the standard meter stick does not exist. Account 2 provides the right sorts of distance relations to be the semantic values of expressions like ‘one meter apart,’ as Kripke understands them. So does a version of account 1: if the fundamental spatial relation is the three-place relation *the distance from x to y is r* then ‘one meter apart’ expresses the derived two-place relation *the distance from x to y is n* for some number n . But as I mentioned above, while on this second interpretation we can (if asked) tell which things bear *the distance from x to y is r* to the same number that the endpoints of the standard meter actually do, if asked which things bear this relation to the number one, we have no idea how to find out the answer.

is five. And not only is this necessary; other instances of the triangle inequality are true as well.

The problem is not that these necessary truths are not logical necessities. I do not claim that all necessities involving fundamental relations are logical (as some versions of the combinatorial theory of possibility do). The problem is that these necessities exhibit a striking pattern, and we ought to be able to explain why they exhibit this pattern. But we cannot.²⁶

This argument generalizes to other attempts to turn theories of mixed fundamental spatial relations into theories of pure fundamental spatial relations by substituting infinite families of pure $n - 1$ place relations for a single n place mixed relation. For example, a theory that says that relations like *x and y are two times as far apart as z and w* are fundamental fails for the same reasons.

8 PURITY and Substantivalism

So far I have been using PURITY to beat up on relationalists. But (as I remarked in footnote 19), substantivalists should also accept PURITY. It may fairly be demanded whether and how substantivalists can produce a theory of the fundamental spatial relations that is compatible with PURITY.

Substantivalists have no problem doing this. There is a synthetic axiomatization of Euclidean geometry using just two primitive predicates of points of space, ‘ x, y are congruent to z, w ’ and ‘ x is between y and z .’ (Actually there is one such axiomatization for two-dimensional Euclidean geometry, and one for three-dimensional Euclidean geometry, and so on; let’s stick to the three-dimensional

²⁶Field argues that we would not want to accept a physical theory that used two-place *predicates* like ‘the distance from x to y is one’ as semantic primitives, because it would be unlearnable (since it contains infinitely many primitive predicates) and unusable (since it contains infinitely many axioms that cannot be captured in a finite number of axiom schemas). But this is not enough to show that we cannot accept that two-place relations like *the distance from x to y is one* are fundamental: for the proponent of this metaphysical view may deny that the physical theory scientists learn and calculate with needs to have semantic primitives that express fundamental relations.

case.) These two predicates clearly express pure spatial relations.²⁷ Field (1980) has produced a similar synthetic axiomatization of the geometry of neo-Newtonian spacetime. A substantialist in search of a PURITY-friendly account of which relations are fundamental will be attracted to synthetic axiomatizations of geometry like this. He need only accept that these synthetic axioms are true of spacetime and make the further claim that the predicates which appear as semantic primitives in the theory express fundamental relations.²⁸

But this account of the fundamental spatiotemporal relations is not available to relationalists. The first problem is that it is not obvious what laws the relationalist will propose for betweenness and congruence. He cannot use the same laws that substantialists use. Those laws entail, among other things, that there are infinitely many points of space; so if a relationalist accepted them and rewrote them to mention particles rather than points of space, he would have to accept that if the world has a Euclidean spatial structure then there are infinitely many particles. And no relationalist wants to accept that.

Relationalists could weaken the laws, only requiring that no set of particles instantiates (at a time) betweenness and congruence in a pattern that is contrary to the substantialist laws. This law amounts to a guarantee that there is at least one embedding of any time-slice of a relationalist world into Euclidean space. (Whether this law can be stated directly, without mentioning embeddings, is a good question, but I set it aside.)

The problem now is that there are (in general) *too many* embeddings of the particles into abstract Euclidean space. Here is an example.


Suppose that relationalism is true and that there are only three particles A, B

²⁷There is a theorem that establishes that every model of these synthetic axioms is isomorphic to \mathbb{R}^3 with a Euclidean geometry. This theorem is only true for the second-order axiomatization; Tarski sketches a proof of a similar result for the first-order theory in (Tarski 1959).

²⁸And this *is* an additional claim. If he does not make it, then it is consistent with all that has been said that the unmixed relation expressed by ‘ x, y are congruent with z, w ’ is analyzed in terms of mixed fundamental relations. (Perhaps it is analyzed as *there is an r such that the distance from x to y is r and the distance from z to w is r .*) But then PURITY is false.

and C , and the facts about the instantiation of betweenness and congruence at a time (other than the trivial ones²⁹) are as follows: betweenness is nowhere instantiated; congruence is nowhere instantiated. Any embedding of this relationalist world into \mathbb{R}^3 maps these particles to the vertices of a triangle that is not an isosceles or an equilateral triangle; and for *any* such triangle there is an embedding that maps the particles to the vertices of a triangle like that one. So, for example, there is an embedding of the particles that maps them to the vertices of triangle displayed in figure 1.1 on the left; and also one that maps them to the vertices of the triangle displayed below on the right. Clearly, no isometry of Euclidean space maps one of

Figure 1.1:



these triangles onto the other. So on this version of relationalism, the embedding of this world into Euclidean space is not unique up to isometry. And this is not a problem with this particular choice of pure fundamental spatial relations; the same problem will arise if relationalists use some other choice of primitives for synthetic geometry as a guide to which spatial relations are fundamental.³⁰

9 Modal Relationalism

The betweenness and congruence theory has a lot going for it. It escapes the problems that face both of the previous accounts I discussed. Why not use modality to

²⁹The trivial facts are those like A is between A and A and A, B are congruent to A, B , the universal generalizations of which are theorems.

³⁰Other choices of primitives are: congruence alone; the three-place predicate ‘ x, y , and z form a right triangle at y ’; and the three-place predicate ‘the distance from x to y is less than or equal to the distance from y to z .’ For each choice there are distinct spatial arrangements of three particles that are isomorphic with respect to the primitive predicates. (Royden 1959) surveys choices of primitives for synthetic geometry.

patch up the relationalist version of the betweenness and congruence theory and secure the unique embeddability of time-slices into Euclidean space? After all, some relationalists characterize the view in modal terms to begin with: Leibniz, the arch-relationalist, wrote: ‘space denotes, in terms of possibility, an order of things that exist at the same time’ (Leibniz and Clarke 2000, page 14).

The idea is to define distance relations in terms of betweenness, congruence, and a modal operator, and to ensure that these distance relations obey Euclidean laws. Then the world’s time-slices will be uniquely embeddable into Euclidean space.

One way to do this is to look at how a substantialist who accepts the betweenness and congruence theory can define four-place pure distance-ratio relations like *x and y are twice as far apart as z and w*, and try to convert those into modal definitions. Substantialists can define this relation in terms of betweenness and congruence as:

(S_2) $\exists v$ (*v is between x and y and x and v are as far apart as z and w and v and y are as far apart as z and w*)

A relationalist might simply prefix this definition with a modal operator:

(R_2) $\diamond \exists v$ (*v is between x and y and x and v are as far apart as z and w and v and y are as far apart as z and w*)

There is an analogous definition (R_3) for *x and y are three times as far apart as z and w*, and so on.

There are two related problems with this approach.

First, there are problems with the meaning of the modal operator. It cannot express logical or metaphysical or physical possibility. It is logically and metaphysically and physically possible for the distance ratio between two pairs of particles to be anything you please. If the modal operator expresses logical or metaphysical or physical possibility, then, two given pairs of particles will satisfy both (R_2) and (R_3). But there must be a unique distance ratio between the pairs of particles.

It cannot express metaphysical possibility restricted to worlds in which the non-modal spatial relations among the four points are as they actually are. Fixing

the betweenness and congruence relations fixes all the non-modal spatial relations. Suppose there are only four particles and that no particle is between any of the others and no two are as far apart as any other two. It is (metaphysically) consistent with this description that the first two particles are twice as far apart as the second two. It is also (metaphysically) consistent with this description that the first two particles are three times as far apart as the second two. So the four particles again instantiate both (R_2) and (R_3) , which they should not.³¹

Relationalists can say that the modal operator expresses ‘geometrical possibility.’³² But what is geometrical possibility? It must be some restriction of metaphysical possibility. I’ve looked at three restrictions; none of them work. I can think of no other candidate descriptions of this restriction. We may wonder whether we understand what this operator means at all.

There are also problems with necessary connections. I rejected account 2 because it asked us to accept necessary truths involving distance relations as brute and inexplicable. This account faces the same problem. Why is it that four particles can only instantiate one of (R_2) , (R_3) , (R_4) , and so on? We cannot derive these necessary truths from the way the modal operator used is analyzed, since it has no analysis.

The following is the best way for relationalists to respond to these problems:

I admit that I cannot ‘define’ the geometrical possibility operator by finding some sentence and then saying that the geometrical possibility operator acts like a quantifier over just those metaphysically possible worlds in which that sentence is true. But I can tell you something about which worlds (in addition to the actual world itself) are geometrically possible relative to the actual world.

Among the metaphysically possible worlds are worlds in which Machian spacetime exists, all and only the actual particles exist, and the non-modal spatiotemporal relations among the particles are as they actually are. For reasons given above, these worlds do not all agree on

³¹Field (1989) makes these points, and extends the argument to all spatial relations, modal or not.

³²(Belot 2000).

the distance ratios between pairs of particles at each time. But the set of these worlds divides up into equivalence classes, where members of any one equivalence class *do* agree on the distances between the particles at each time. One of these equivalence classes is special. The members of that class and no other are geometrically possible relative to the actual world. But I can't tell you which class is special, or why it is special and the others aren't.

(Which worlds are geometrically possible varies from world to world, and since there are worlds that are non-modal duplicates—have the same history of betweenness and congruence relations—but differ in distance ratios, there are worlds which are non-modal duplicates but differ over which worlds are geometrically possible.)

This helps with the problems I mentioned: we know how the geometrical possibility operator works, and the necessary connections between distance ratios, while still brute, are not so mysterious.

Still, we should reject this proposal. For one thing, it asks us to believe in brute modal differences: worlds which are isomorphic with respect to betweenness, congruence, and temporal ordering, and so are the same, as far as their non-modal facts are concerned, but which differ over distance ratios—which are modal facts—between pairs of particles.

But the more important problem with this proposal is that it looks like cheating. Relationalists need to guarantee the unique embeddability of the world's particles into Machian spacetime; the solution here is to use a modal operator to do it by brute force.

If this is how relationalists about motion hold on to relationalism about ontology, then what is attractive about relationalism about motion in the first place? There are supposed to be epistemic and (related) metaphysical advantages: interparticle distances and relative velocities are easier to observe, and a theory that relies only on them does not need to postulate spacetime. But if we can use these strange modal operators, then relationalism about motion no longer has these advantages over its denial. As for metaphysical advantages: let the world have a Newtonian spatiotemporal structure, so that talk of absolute states of motion makes

sense. We can still be relationalists; there is no need to appeal to spacetime to make sense of absolute motion. Just use a new modal operator to guarantee the unique embedding of the world's particles into Newtonian spacetime, and use this embedding to define states of absolute acceleration. As for epistemic advantages: if inter-particle distances and states of absolute acceleration are equally funny modal facts, why should one be easier to observe?

Like accounts 1 and 2, modal relationalism is a bad move. Even if relationalism about motion is true, relationalists cannot adequately characterize the spatiotemporal structure of the world. We should be substantivalists.

Chapter 2

Modal Arguments against Substantivalism

1 Introduction

Two famous arguments against substantivalism—the Leibniz shift argument and the hole argument—turn on substantivalism’s (alleged) modal commitments. Both contain premises to the effect that substantivalists must believe that it is possible that things be just as they are, qualitatively, while differing in some non-qualitative respect.¹ Leibniz argued that these the possibilities conflict with some a priori metaphysical and theological principles. Defenders of the hole argument argue that the possibilities in question lead to a failure of determinism. Responding to the hole argument is harder, so I will focus on it. I will say something about how we should think about Leibniz’s argument in the course of responding to the hole argument, though. The hole argument looks like this:

¹Qualitative properties are those that ‘make no reference’ to particular individuals. So the property of being red is qualitative, while the property of being in the same room as Ralph Nader is non-qualitative. If we have a language the predicates of which express qualitative properties, and which contains no proper names, then we cannot distinguish between qualitatively indiscernible worlds using this language: every sentence true in one of the worlds is also true in the other. I say more about what the qualitatively indiscernible worlds at work in the Leibniz shift argument and the hole argument look like below.

- (1) If substantivalism is true, then (assuming that general relativity is the true theory of the world²) it is physically possible that everything be just as it actually is, except that some spacetime points in the (absolute) future ‘play different roles.’
- (2) If there is more than one physically possible future consistent with the way things are now, then determinism is false.
- (3) Therefore, if substantivalism is true (and general relativity is the true theory of the world), determinism is false.

I’m going to take it for granted (as most do) that accepting the conclusion—that general relativity is not a deterministic theory—is an embarrassment for substantivalists. So there are two ways for substantivalists to respond to this argument. They can deny the existence of the qualitatively indiscernible possibilities mentioned in the first premise; or they can deny the second premise, and so deny that such possibilities lead to a failure of determinism.

I advocate the second kind of response. Making this response work involves getting clear on just what determinism is. I defend an analysis of determinism that blocks the hole argument and is independently plausible. Although I defend this analysis of determinism in the context of responding to the hole argument, its value is not limited to the use to which I put it. Understanding determinism is something worth doing for its own sake. And, more importantly, determinism shows up elsewhere in the debate between relationalists and substantivalists. For example, Earman (1989), following Stein (1977), argues relationalism about motion entails relationalism about ontology. In the previous chapter I argued that relationalism about motion requires substantivalism. Unless relationalism is inconsistent, we cannot both be right. Earman’s argument contains a premise about determinism; it presupposes the same analysis of determinism as the one at work in the hole argument. So like the hole argument, Earman’s argument fails for relying on (what I argue is) an incorrect analysis of determinism. I briefly discuss this argument in

²There is some debate over whether the hole argument applies to theories other than general relativity. Earman and Norton (1987) argue that it does, while Earman (1989) argues that it does not. This controversy won’t matter for my discussion.

section 4.

Before discussing determinism, though, I'll explain why I accept the first premise of the argument. It has become more common lately for substantialists to deny the first premise; I'll discuss why I think this is unwise, and why I think much discussion of the first premise is misdirected.

2 Defending The First Premise

Let's begin with the first premise. What reason is there to believe it? To answer this question I will first discuss the first premise of the hole argument's predecessor, the Leibniz shift argument. It is easier to first distinguish between good defenses and bad defenses of the first premise of this other argument, and then apply the lessons learned to the evaluation of the first premise of the hole argument.

2.1 The Leibniz Shift Argument

The first premise of the Leibniz shift argument is this:

- (L1) If substantivalism is true, then (assuming that some theory, like Newtonian gravitational theory, that lives in neo-Newtonian spacetime is true) it is physically possible that at each time, each material object be one foot to the left of where it actually is at that time.³

How might one argue for (L1)? Here's one way. On one way of formulating Newtonian mechanics, we write down some equations which pick out a set of models. These models are n -tuples of mathematical objects. A typical model may include \mathbb{R}^4 and several curves through \mathbb{R}^4 (functions from \mathbb{R} to \mathbb{R}^4).⁴ The models represent (or correspond to) physically possible worlds: by looking at the models we may figure out what sorts of arrangements of particles in spacetime are permitted by the theory. But of course *by themselves* the models tell us nothing about what is physically possible. Only when we also put in place some principles of interpretation, principles which assign representational properties to the models, do they yield in-

³Let's pretend that 'to the left' picks out a determinate direction in space.

⁴These curves must satisfy certain constraints, but they won't be important.

formation about physical possibility. Consider, for example, the model that consists of \mathbb{R}^4 and just one curve that assigns to each real number r the point $(r, 0, 0, 0)$ in \mathbb{R}^4 . What would the world be like, if this model correctly represented it? We cannot tell, until I specify some principles of interpretation. But when I tell you that \mathbb{R}^4 represents spacetime, that the first component of points in \mathbb{R}^4 represents temporal location and the other three represent spatial location, that some of the geometrical relations between points of \mathbb{R}^4 represent geometrical relations between points of spacetime, and that curves in \mathbb{R}^4 represent the careers of point particles in spacetime, then we can see that this model represents a universe in which there is just one point particle moving inertially for all time.

Of course I have not given a complete catalog of the principles of interpretation for these models, and there is room to disagree about just what those principles are. (Relationalists, for example, may deny that \mathbb{R}^4 represents spacetime, and may instead maintain that it is a fictional device for encoding spatiotemporal relations among point particles.) But there is one set of principles which, together with facts about what models there are, entails (L1). Suppose that in addition to the principles I gave in the last paragraph we also accept that each point of \mathbb{R}^4 represents some actual point of spacetime and that it represents the same point of spacetime in every model.⁵ With these interpretive principles in place, consider (again) the model in which a particle sits at $(t, 0, 0, 0)$ for all $t \in \mathbb{R}$. It follows from facts about the theory that if there is such a model, then there is also another model in which a particle sits at $(t, 1, 0, 0)$ for all $t \in \mathbb{R}$. And since in these two models a particle is located at different points in \mathbb{R}^4 it follows from what I've just said that these models represent a particle as located at different points of spacetime. Since each model represents a lone particle moving inertially for all time, these models correspond to distinct possibilities that differ only with regard to which points of spacetime the particle

⁵There is an obvious problem with interpreting the models this way, which I will set aside. On this way, if relationalism is true, (and so if there are no actual points of spacetime), then the substantivalist interpretation of the theory is necessarily false: not only is every model a false representation, no model could have been a correct representation. (Even if there had been points of spacetime, they would not be the actual points of spacetime, and no model can be correct unless the actual points exist.)

occupies. So these possibilities differ merely non-qualitatively. (If we suppose that the direction from $(0, 0, 0, 0)$ to $(0, 1, 0, 0)$ is the left-ward direction in space, then these models differ only in that according to one of them, the lone particle is shifted one foot to the left of its location according to the other model, at each time.)

Similar reasoning applies to more complicated models. Assuming that the theory is true, then, there is a model which represents the world as it is, and another model that represents each material object at each time occupying a region one foot to the left of the region it actually occupies at that time. Add the premise that whatever a model represents is (physically) possibly true, and (L1) follows.

But *by itself* this is not a good way to argue for (L1). For the mathematical models of our theories don't come with their representational properties built in. They get their representational properties from us. And substantialists are not required to give them the representational properties needed to make this argument work. A substantialist could, for example, interpret the models so that the models yield only qualitative information about the world. On this second way of interpreting the theory, no claims about *de re* possibility—no claims about which particular points of space a particular particle might be located at—follow from inspection of the models along with the rules for interpreting them. So, since the first premise of the Leibniz shift argument is a claim about *de re* possibility, it will not follow from inspection of the models along with the rules of interpreting them.

To clarify this second way of interpreting models, return to our two models, one in which a particle sits at $(t, 0, 0, 0)$ for all $t \in \mathbb{R}$, the other in which a particle sits at $(t, 1, 0, 0)$ for all $t \in \mathbb{R}$. On the second way of interpreting models both of these models say the same thing about the world: they both say that there is just one particle moving inertially for all time. Neither says anything about just which (actual) points of spacetime that particle occupies. (Indeed, we cannot even say whether it is one and the same particle that both models represent.)

So the above argument for (L1) has no force without some argument that substantialists must give the models one set of representational properties rather than another. What might such an argument look like? The most obvious way to argue that substantialists should assign models the first set of representational properties goes like this: substantialists believe that it is possible that everything

be one foot to the left of where it actually is at each time; so they should interpret the models in such a way that there is a model for each of these possibilities. But such an argument is question-begging in the current context: we are trying to argue for (L1), so we cannot use (L1) as a premise.

How else might one argue for (L1)? The only real ‘argument’ for (L1) is an appeal to one’s modal intuitions. It seems intuitive to many people that some restricted combinatorial principle governs possibility.⁶ In particular, the following seems intuitive: *any* way of distributing particles across spacetime is a *possible* way of distributing particles across spacetime. That is, as far as (metaphysical) possibility goes, the spatiotemporal location of a given particle is independent of the location of any and all other particles. But this principle (together with facts about what the laws are) entails (L1).

(L1) also seems intuitive when we consider what happens when a substantialist denies it. Denying (L1) is objectionable for the same reason that the more general denial of the existence of possible worlds that differ merely non-qualitatively is objectionable: it entails implausible essentialist claims.⁷ (L1) concerns a world in which everything is shifted one foot to the left at each time. But presumably anyone who denies (L1) will also deny similar premises asserting that substantialists must believe in worlds where everything is shifted *two* feet to the left, or one foot to the right, and so on. For short, let’s say that such a person asserts that there are no shifted worlds. Now, for simplicity, let’s consider the case where there is no time, only (three-dimensional Euclidean) space. Suppose that there are actually only two point-particles, Joe and Moe. Joe and Moe are two feet apart, and they occupy points *p* and *q*. Then if there are no shifted worlds it is *necessary* that if Joe and Moe are two feet apart (and nothing else exists), they occupy *p* and *q*. And we can say something stronger. According to (L1) substantialists must accept possible worlds that are shifted relative to the actual world. But any reason to believe that

⁶The *unrestricted* combinatorial principle is much more controversial. It entails that anything could have had any combination of fundamental properties at all. So (given plausible assumptions about which properties are fundamental) it entails that I could have been a positron, or a point of spacetime, or the successor of four.

⁷(Adams 1979), (O’Leary-Hawthorne and Cover 1996).

would be a reason to believe that substantivalists must accept possible worlds that are shifted relative to *any* physically possible world. The negation of this claim is this: for any specification of the distances between n point particles there is exactly one possible world in which only those n point particles exist, and they have those inter-particle distances.⁸ But how could this be? In the case of Joe and Moe, what is so special about p and q that makes them the only possible locations of Joe and Moe, when Joe and Moe are two feet apart? Of course there is nothing special about them; if you think that there are no shifted worlds, you must believe that this is just a brute modal fact.⁹

That's all I'm going to say to motivate (L1). What further debate there may be about (L1) can easily be translated into debate about the first premise of the hole argument; since that argument is my focus, I'll save the further debate for my discussion of it.

Since I accept (L1), I will briefly explain where I think the Leibniz shift argument goes wrong. What is supposed to be wrong with recognizing the possibilities in (L1)? The possibilities at work in (L1) are qualitatively indiscernible; but what

⁸I'm continuing to assume that there is no time, only space, and that that space's existence is physically necessary.

⁹Admittedly, there is something special about p and q : Joe and Moe actually occupy them. I find it hard to believe that this explains why they have the modal property in question. And there are other brute modal facts that someone who denies (L1) must accept that cannot be explained in this way. There could have been three particles, instead of two, each two feet from the others; if (L1) is false then there are three points of space that are the only possible locations of these three particles, when they have those inter-particle distances. What makes them so special? Not, in this case, that they are actually occupied.

One might claim that there are no shifted worlds and try to avoid these consequences by asserting that it is indeterminate where each material object is located. (Presumably it will still be determinately true that there are points p and q such that necessarily, if only Joe and Moe exist and they are two feet apart, then they occupy p and q . But there are no points p and q such that it is determinately true *of them* that necessarily, if only Joe and Moe exist and are two feet apart, then they occupy p and q .) But facts about where things are located are supposed to be fundamental facts; and I don't believe that there could be indeterminacy where such fundamental facts are concerned.

is bad about qualitatively indiscernible possibilities? Relationalists mention several distinct problems. There is an epistemic problem: we can never know which of the indiscernible possibilities is actual. There is a theological problem: God could have no reason to actualize this world rather than a possibility qualitatively indiscernible from it. The theological problem worries no one these days, and the epistemic problem is a pseudo-problem.¹⁰ What remains is just the bare claim that

(4) There are no qualitatively indiscernible possibilities.

Some find (4) intuitively plausible. I do not, and I also think there are reasons to reject it. Above I argued that substantialists who accept (4) and deny (L1) must accept implausible essentialist claims. As I mentioned above, I believe more generally that anyone who accepts (4), substantialist or not, must accept implausible essentialist claims. But I won't argue for this more general thesis here.

2.2 Back to the Hole Argument

As in the Leibniz shift argument, the first premise of the hole argument asserts that substantialists must believe in possible worlds qualitatively indiscernible from the actual world. But just what do these alternative possibilities look like? In the context of the Leibniz argument, it was easy to say: those were possibilities which differed with regard to where material objects were located. In the context of the hole argument things are not so easy. Now general relativity is our background physical theory. And general relativity (when given a substantialist interpretation) differs in ontology from Newtonian mechanics. According to the latter theory (or the version of it that I had in mind) there was spacetime, on the one hand, with its geometrical properties; there were material objects, on the other hand, with their intrinsic properties (properties like mass and charge); and material objects and points of spacetime 'interacted' by the former being located at the latter. But general relativity is a field theory. In the mathematical models of the theory, there are vector

¹⁰Earman (1989) finds both problems in Leibniz's letters to Clarke, though Earman puts the epistemic problem in verificationist terms. Maudlin (1993) dissolves the epistemic problem: it is a priori that each particle is where it actually is at each time, rather than shifted one foot to the left.

and tensor fields (like the metric and the stress-energy tensor) defined on the four-dimensional manifold which represents spacetime. But what is this theory telling us about the world? Of course, if we're substantivalists, we think it is telling us that spacetime exists. But what else is there? Are there in addition a bunch of mathematical objects related to the points of spacetime? That is a strange view. Or are there in addition a bunch of very large material objects—fields—the properties of which vary from point to point? Or is there instead just spacetime, which in addition to its geometric properties, has non-geometric intrinsic properties? (Or is some fourth interpretation correct?)

We can't settle these questions here. But we need to have some way to characterize the possibilities mentioned in the first premise, if we're going to look at arguments that substantivalists need to believe in them. We can of course characterize the possibilities in very general terms—we can say, they are possibilities which are qualitatively indiscernible, but which differ non-qualitatively because some spacetime points 'play different roles.' But without further information, we don't really know what this means.

Let's suppose we give general relativity the last ontology I mentioned—the one according to which there is just spacetime, with geometric and non-geometric properties. (Nothing will turn on this choice.) Then the possibility we are asked to consider is this one: in some future region of spacetime (the 'hole'), the geometric and non-geometric properties of points of spacetime in that region are 'pushed around' smoothly (so that nearby points stay nearby¹¹), leaving things qualitatively as they actually are. Of course, putting it this way is slightly misleading, because we are 'pushing around' geometrical properties. Since the geometry of spacetime is changing, there is no sense in which the points of spacetime are 'staying put' while the properties are 'moving.' We could just as truly characterize the possibility this way: in some future region of spacetime, the geometric and non-geometric properties are left where they are, but the points of spacetime are 'pushed around' underneath them in a suitably smooth manner. And indeed this characterization is

¹¹The ensures that the function we're using to move the points around is continuous, though actually the function must meet stricter requirements: it must be a diffeomorphism.

somewhat more accurate. For consider two points inside the hole, p and q , such that q ‘gets pushed to where p used to be.’ Now take any point r outside of the hole. r is actually some distance d from p .¹² In the possibility we’re contemplating, the distance from r to p is (probably) not d ; instead the distance from r to q is d . Similar facts hold for other points. So in this non-actual possibility q is playing the ‘geometric role’ that p actually plays. (Here is a picture of what is happening: Imagine God looking down on the spacetime manifold with the properties distributed across it as one might look down at the island of Manhattan; imagine Him ‘lifting up’ these properties as one might lift up the buildings of Manhattan; imagine Him then focusing on some future region of spacetime underneath the properties, and pushing around the spacetime points in that region as one might then push around the dirt underneath some region of the upper west side (say, the region under Columbia University); imagine him finally putting the properties back down, ‘just the way they were.’ After God’s activity, things are just as they were before, qualitatively speaking.)¹³

2.3 Confusions About The First Premise

What reason is there to believe that the first premise is true? In section 2.1 above I discussed a bad argument for the first premise of the Leibniz shift argument. An analogous bad argument often appears in discussions of the hole argument. As

¹²I’m speaking loosely here. Take r to be a point such that there is a unique (either timelike or spacelike) geodesic connecting r and p , and let d be the ‘length’ of that geodesic.

¹³These possibilities are somewhat complicated, because continuum-many points are being ‘moved around.’ Are there simpler possibilities that work just as well? Why not take just *two* spacetime points in the (absolute) future and have *them* ‘switch roles?’ That is, why not just take two spacetime points in the future and have them exchange all of their geometric and non-geometric properties (including relational properties like the property of being ten meters from point q)? This possibility is certainly easier to imagine: we simply imagine God focusing his attention on two spacetime points in the future, reaching down and removing them from the spacetime manifold, and then placing each in the hole left by the other. Melia (Melia 1999, section 2.1) argues that these simpler possibilities do work just as well.

before, this argument appeals to facts about the models (in this case, models of general relativity) together with contentious principles for interpreting them. A model of general relativity is a four-dimensional manifold together with vector and tensor fields defined on it. It is a fact about the theory that if M is a manifold in some model then there is another model that also contains M but in which the vector and tensor fields have been pushed around, so that they make the same pattern on M but the points in M play different roles. (The mathematical device that does the pushing around is a function from the manifold to itself called a ‘diffeomorphism’ and the two models are said to be ‘related by a diffeomorphism.’) We are asked to accept the first premise because the theory contains diffeomorphically related models.

So for example Carl Hoefer writes, ‘given the identification [of spacetime with the manifold; that is, given that the four-dimensional manifolds that appear in the models represent spacetime], *and a straightforward interpretation of the mathematical apparatus*, the substantialist is committed to an infinity of qualitatively indistinguishable possible worlds’ (Hoefer 1996, page 7; italics mine). Hoefer supposes that an interpretation of the models according to which they yield non-qualitative information (an interpretation according to which each point of a manifold represents a particular spacetime point, and represent the same spacetime point in each model in which it occurs) is more straightforward than an interpretation according to which models yield only qualitative information. I don’t see how one interpretation is any more straightforward than the other. It may be that one fits better with the antecedently established modal commitments substantialists make; but (as I remarked above) this justification is question begging in the current context. Again, arguing for the first premise in this way by looking at the models of the theory does not work.

Part of the trouble is that the first premise is often *stated* as a claim about the representational properties of the mathematical models. So Carolyn Brighthouse writes, ‘The central claim [of the hole argument] is that the substantialist has to view diffeomorphically related models as representing distinct situations’ (Brighthouse 1994). Similarly, the central question that Jeremy Butterfield thinks the hole argument raises is, Do models related by a hole diffeomorphism represent the same possible world (1989, page 12)? Stating the first premise in these terms leads to

confusion. Brighouse's claim is stronger than the first premise of the hole argument as I have written it down; it is the conjunction of my premise and some claim about the representational properties of models. So her claim could be false while my first premise is true. Arguments against the first premise as Brighouse states it, then, are not sufficient to block the hole argument.

Although Earman in his book (1989) also presents the first premise as a claim about models and their representational properties, Earman and Norton (1987) do not seem to make this mistake. They *first* assert (in the context of the Leibniz shift argument) that substantivalists must accept the possibility of shifted worlds, and *then* 'translate' this claim into a claim about the representational properties of models. Claims about the representational properties of models do not appear as premises in an argument for the first premise.¹⁴

2.4 Defending the First Premise

We still have no argument that substantivalists need to accept the first premise. When I discussed the first premise of the Leibniz shift argument, I gave an argument for it that relied on certain combinatorial intuitions. Can we appeal to those intuitions here?

It's not clear that we have the combinatorial intuitions needed in this case. It is obvious (to me anyway) that (assuming points of spacetime exist) the spatiotemporal locations of material objects are independent of each other. But the indiscernible possibilities at work in the Leibniz shift argument do not differ with regard to the *geometrical* role that each spacetime points plays, as do the possibilities at work in the hole argument; and it is not nearly so obvious that a point of spacetime could have played a different geometric role from the one it actually plays.

One can try to make the two cases look similar, so that the same intuitions that support the Leibniz shift also support the hole construction, as follows:

In the context of the Leibniz shift argument, you accept that material objects could be located in regions other than the regions at which they

¹⁴Though it is true that their paper contains no arguments for the first premise; they seem to think its denial is inconsistent with substantivalism.

are actually located. But as we are understanding general relativity there are no things other than points of spacetime. (This is not to say that according to general relativity people do not exist, or that according to general relativity there are no tables; they are just not what we thought they were.) So in the context of general relativity the possibility according to which you are located in some region other than the one in which you are actually located is not correctly described, at a fundamental level, as one in which certain particles bear the location relation to certain points of spacetime. Instead, it is correctly described as one in which certain non-geometric properties that certain spacetime points instantiate have been changed around in an appropriate way. So what your combinatorial intuitions are telling you, in this context, is that there is no impossibility in the idea of changing which non-geometric properties certain points of spacetime instantiate. But there is no principled distinction to be made, in the context of this theory, between the geometric and the non-geometric properties. So you ought to admit that it is possible to shuffle the geometric properties as well.

Why is there no principled distinction to be made? There are several reasons that could be offered here. For one thing, just as there are many physically possible ways to distribute non-geometric properties in spacetime, in general relativity there are many physically possible geometries for spacetime. For another, in general relativity different distributions of non-geometric properties in spacetime *require* as a matter of physical law different distributions of geometric properties.¹⁵

This line of reasoning has some plausibility. But it is not nearly as strong as the combinatorial support for the first premise of the Leibniz shift argument. Here's one reason. The combinatorial intuition I cited was not an intuition that the spatiotemporal locations of material objects can be changed around, whatever the correct physical and metaphysical theory is. Rather, my intuition was this: if *this* is the correct way of describing the world—there are material objects, there is

¹⁵Something like this line of thought occurs on page 519 of (Earman and Norton 1987).

spacetime, material objects are located in spacetime—*then* it seems intuitive that the locations of material objects can be changed around.¹⁶ So the support for the first premise of the Leibniz shift argument does not transfer automatically to the first premise of the hole argument.

Still, I accept the first premise. I think it becomes clear that substantialists should accept it when we see what is involved in denying it. Denying the first premise involves some appeal to essentialism. Since the first premise is about futures in which spacetime points switch roles in the geometry, first-premise deniers often appeal to some form of geometrical essentialism.¹⁷ There are various ways to formulate geometrical essentialism: one could claim that (actual and possible) points of spacetime have their qualitative geometric properties (including their curvature properties) essentially; or one could claim that points of spacetime have their qualitative properties and their non-qualitative relational geometric properties (like *being ten feet from point p*) essentially.

¹⁶I think it clear that one's intuitions about possibility can depend on one's beliefs about fundamental metaphysical matters. Here's an example. Forget about time for a minute; and suppose space is Euclidean. And suppose it is Euclidean in virtue of the points of space (along with numbers) instantiating a three-place relation, *the distance from x to y is r*, in a certain pattern. Then it seems perfectly possible that every pair of points of space be twice as far apart as they actually are. For this simply involves each pair *m* and *n* instantiating *the distance from x to y is 2r* iff they actually instantiate *the distance from x to y is r*, where *r* is some real number. But now suppose instead that space is Euclidean in virtue of the points of space instantiating two relations, betweenness and congruence, in a certain pattern. Then it seems impossible that every pair of points of space be twice as far apart as they actually are—because there are no absolute facts about how far apart two things are, there are only comparative facts about whether two things are the same distance apart as two other things. (One might assert that 'the congruence relation can hold between points of space in different possible worlds,' and that this allows one to make sense of the possibility in question; but this assertion only makes sense if one is a modal realist, which I am not. And even modal realists will hesitate to say that the congruence relation can hold between points of space in different possible worlds: Lewis defines 'possible world' as 'maximally spatiotemporally related concrete object.' It follows from this definition that parts of distinct possible worlds bear no spatial relations (and congruence is a spatial relation) to each other.)

¹⁷(Maudlin 1990).

We should reject these versions of geometrical essentialism. In general relativity, the geometry of spacetime depends on the distribution of mass-energy.¹⁸ So, if general relativity is true, then if I had raised my hand a moment ago, the geometry of the region of spacetime around me would have been different. So, if these versions of geometrical essentialism are true, then if I had raised my hand a moment ago, (part of) the region of spacetime I actually occupy would not have existed. But certainly it was up to me whether I raised my hand a moment ago; so if these versions of geometrical essentialism is true, it was up to me whether a certain region of spacetime exists. But that is absurd.¹⁹

There is a weaker version of essentialism that does not entail that what points of spacetime exist depends on what I do. It is an instance of the more general doctrine that there are no possible worlds that differ merely non-qualitatively. On this view it is impossible that (a) all the actual points of spacetime exist, (b) the geometry of and distribution of matter in spacetime is just as it actually is, and (c) some points of spacetime play different roles than they actually play. But this view does not entail that a given spacetime point have any particular curvature property essentially, or that it must be any particular distance from some other point. So this view evades the objection above: it is consistent with this view that even if I had raised my hand, all the actual points of spacetime would still have existed.

What motivates those who appeal to geometrical essentialism is the idea that a point of spacetime cannot be separated from its geometrical properties. This idea

¹⁸The argument to follow depends on special features of general relativity; analogues of the hole argument in the context of other physical theories are immune to it.

¹⁹Maudlin replies to an argument like this one by ‘appealing to counterpart theory’ (1990, page 550). That is, after arguing that ‘spacetime (the actual spacetime) could have had a different geometry’ is false, he proposes a counterpart-theoretic semantics for our modal vocabulary on which ‘spacetime (the actual spacetime) could have had a different geometry’ is true. But for this response to work, he must claim that my objection seems plausible when we understand the modal vocabulary that occurs in it in the new way, but does not seem plausible when we understand it in the old way. But I don’t think this is so; when I wrote the objection down, I meant to be using the modal vocabulary in just the way Maudlin does when he defends essentialism; and the argument seems plausible when read that way.

does not motivate the weaker essentialism I am now discussing. I gave my reasons for rejecting its motivation, the doctrine that there are no merely non-qualitative differences, above on page 32: it leads to implausible essentialist claims. So I look elsewhere for a response to the hole argument.

2.5 Counterpart Theory

Butterfield asserts that one can deny the first premise of the hole argument without being an essentialist by appealing to counterpart theory (1989, page 22). How might this work? Essentialism is a collection of *de re* modal claims. Counterpart theory is not; it is (part of) an *analysis* of *de re* modal claims. (Roughly speaking, *de re* modal claims like ‘Nader could have won’ are analyzed not as ‘There is a possible world in which Nader wins,’ but as ‘There is a possible world in which a counterpart of Nader (someone sufficiently similar to him) wins.’) And counterpart theory is not incompatible with essentialism. There are counterpart relations that make certain essentialist claims true: there is a counterpart relation, for example, which makes true the *de re* modal sentence, ‘I could not have been a poached egg.’

Of course there are other counterpart relations that make this and many other essentialist claims false. To use counterpart theory to deny the first premise without embracing essentialism, then, one must find one of these anti-essentialist counterpart relations that makes the first premise false. But this cannot be done. There are no such counterpart relations because the denial of the hole argument’s first premise is *equivalent* to a version of essentialism. The first premise just is the claim that certain spacetime points can switch their geometrical roles; for this premise to be false is for the points to have (some aspect of) their geometrical roles essentially. This is so whether you analyze *de re* modal claims in terms of counterparts or not.²⁰

²⁰In Lewis’s original version of counterpart theory, as presented in ‘Counterpart Theory and Quantified Modal Logic’ (1968), he made it an axiom that each thing has at its own world only one counterpart: itself. This seems to lead immediately to the result people like Butterfield want: for (reverting for the moment to the Leibniz shift argument) you might think that the truth of ‘Each thing could have been one foot to the left of where it actually is’ requires each point of space to have a counterpart at its own world other than itself. This is still essentialism: but one might think that it is made more palatable by following immediately from Lewis’s theory.

This argument shows more than that appealing to counterpart theory is not a way to avoid essentialism while denying the first premise. It shows that the *only* way to deny the first premise is to embrace some form of essentialism.

3 The Second Premise

I have said that substantivalists should accept the first premise of the hole argument, and so accept that (if general relativity is true) there is more than one physically possible future consistent with the way things are now: the actual future, and a future that is qualitatively indiscernible in which some future spacetime points have switched roles. Now, the second premise of the hole argument is:

- (2) If there is more than one physically possible future consistent with the way things are now, then determinism is false.

To uphold the claim that general relativity is a deterministic theory, then, substantivalists should deny this second premise. I know this premise looks analytic; but this is merely an appearance. There are reasons to reject it which are independent of its role in the hole argument.

3.1 Determinism: Examples and Counterexamples

It is easy to convince yourself that not just *any* alternative physically possible future counts against determinism, as we ordinarily think of it. If one accepts that a given world has futures that differ merely non-qualitatively, then it should be intuitive that those futures do not count against determinism.²¹

But, in the first place, it does *not* follow immediately: we need an extra assumption—namely that there are no qualitatively indiscernible worlds—which Lewis does not make. And, in the second place, Lewis himself came to think this axiom too restrictive and abandoned it in *On The Plurality of Worlds* (Lewis 1986b).

²¹David Lewis's analysis of determinism in (Lewis 1983) entails that alternative futures that differ merely non-qualitatively do not count against determinism. He offers no explanation of why this should be, though. Butterfield's analysis in (Butterfield 1989) and Brighouse's in (Brighouse 1997) are equivalent to Lewis's. I look at Lewis's analysis in more detail below.

Joseph Melia (Melia 1999) gives an example of a world with possible futures that differ merely non-qualitatively, but which seems, intuitively, deterministic. Consider these simple laws: nothing ever moves; there are two kinds of particles, $P+$ and $P-$ particles; each $P+$ particle decays into a $P-$ particle five minutes after coming into existence; the $P-$ particles occupy the same places as the $P+$ particles that give birth to them. $P-$ particles never decay. (We are also to assume that worlds governed by these laws are relationalist: there is no space or time, only spatiotemporal relations between particles).²² Suppose that the history of the world up to a certain time looks like this: in the beginning there was nothing, and then three minutes ago two $P+$ particles were created three meters apart. It appears that the future is determined: in two minutes both $P+$ particles decay into $P-$ particles. There is no other possibility. But if we are working with the conception of determinism at work in the second premise of the hole argument, then this world comes out indeterministic. Call one of the $P+$ particles $\alpha+$, the other $\beta+$. Then $\alpha+$ actually decays into $\alpha-$, $\beta+$ into $\beta-$. But $\alpha-$ and $\beta-$ could have switched roles. So there are two possible futures (differing merely non-qualitatively) for the world; they differ with regard to which $P-$ particles $\alpha+$ and $\beta+$ decay into. But it is absurd to allow these differences count, when asking whether this world is deterministic.

If this line of thought is right, then we have reason to deny the second premise. Unfortunately there is an example which seems to show that this line of thought is not right, that we do allow the existence of physically possible futures which differ merely non-qualitatively to count against determinism.²³ Exert enough downward force on a cylindrical column and it will usually collapse into an elbow shape. We are to imagine a non-actual set of laws of nature according to which cylindrical columns, when they collapse, *always* collapse into an elbow shape, even when the situation is perfectly symmetric around the axis of the cylinder—the force is applied straight down on the center of the top of the cylinder, the cylinder itself is not

²²But wait: then what do we mean when we say ‘the $P-$ particles occupy the same places as the $P+$ particles that give birth to them’? Perhaps Melia means to allow for cross-time distance comparisons, so we can say that some $P-$ particle is zero meters from the $P+$ particle from which it decayed.

²³The following example first appeared in (Wilson 1993).

weaker on one side, and so on. (We are also to imagine that worlds governed by these laws contain Absolute (Newtonian) space and time.) Now imagine a world containing such a cylinder in which a sufficient force is applied straight down on the center of the top of the cylinder, causing it to collapse, but which is perfectly symmetric before the cylinder collapses. (The world contains only the cylinder sitting on a perfectly spherical planet and a sphere falling directly on top of the cylinder.) Intuitively, this world is not deterministic: the laws do not dictate in which direction the cylinder will collapse. But the different possible futures—the cylinder collapses *this* way, or it collapses *that* way—are qualitatively indiscernible.

What does this example show? Even if we agree that the column world is indeterministic, it does not show that the second premise is true. For I, at least, still have the intuitions that in some cases qualitatively indiscernible futures do not count against determinism: unless more is said, the collapsing column example does not suggest that these intuitions were misleading.

So we are left with the task of finding an analysis of ‘determinism’ which sometimes allows merely non-qualitative differences to count against determinism—as in the column world—and sometimes does not—as in the worlds in Melia’s example. And we want an analysis that is appealing for independent reasons, not just because it allows us to defend substantivalism from the threat of the hole argument.

3.2 Analyzing Determinism

There is such an analysis available. Consider the Laplacian picture of determinism: in a deterministic world, an extraordinarily intelligent demon who knew all the information about the present as well as the laws of nature would be able to deduce all the information about other times.²⁴ But what do we mean, *all* the information? We could mean, all the *qualitative* information; we could mean all the qualitative *and non-qualitative* information; or we could mean something in between. So this

²⁴This picture speaks of agreement at a time forcing agreement at all times. There are other varieties of determinism, differing with regard to which regions of space-time worlds must agree on to force agreement through spacetime (Earman (1986) contains a survey). But this is the one in play in discussions of the hole argument, so it is this one I shall focus on.

gives us several ways to make the Laplacian picture more precise: we can be more precise about what information we give the demon to start with, and more precise about what information we demand that he produce. Here's one way to make these more precise: we could give the demon all the qualitative information about the present, as well as the laws of nature, and demand that he produce all the qualitative information about other times. Here's another: we could give the demon all the qualitative *and* non-qualitative information about the present, as well as the laws of nature, and demand that he produce all the qualitative and non-qualitative information about other times. (This is the conception of determinism relationalists rely on in the hole argument.) The analysis I endorse captures the only natural middle position. We give the demon all the qualitative and non-qualitative information about the present, as well as the laws of nature. (If we were to write down the information he is given in a language the predicates of which express only qualitative properties, some of the sentences we would write down would contain proper names of things that exist in the present.) And we demand that he produce all the qualitative information about other times, *and also* all the non-qualitative information about other times which (if we were to write it down) can be expressed using only the proper names we have already given him. (So he need not produce any non-qualitative information that can only be expressed using proper names of things that exist only in the future.) If he cannot do this, then the world is not deterministic.

This picture of determinism meshes with our intuitions about determinism. Consider Melia's particle decay world. No Laplacian demon is needed to write down the relevant information; we can do it ourselves: given that three minutes ago two $P+$ particles, $\alpha+$ and $\beta+$, were created three meters apart, we know that in two minutes there will be two $P-$ particles, one having decayed from $\alpha+$, the other from $\beta+$, and they will be three meters apart, for the rest of time. This world is deterministic, even if we cannot deduce which $P-$ particle decays from which $P+$ particle; even if an extraordinarily intelligent demon cannot, either. For in order to express the required proposition we must use names for the $P-$ particles; and we are not required to deduce information that can only be expressed using such names.²⁵

²⁵By 'names' I mean names that do not have their reference fixed using descriptions. We can, of course, deduce that $\alpha+$ decays into $\alpha-$, because I introduced ' $\alpha-$ '

The column world is indeterministic, though, because (since the points of space endure through time and exist at the initial time) we must be able to deduce all the non-qualitative information expressible using the names of the points of space. So for any point p we must be able to deduce whether the column will collapse toward point p or not; and this we cannot do.²⁶

This picture of determinism also seems intuitively correct, when considered on its own. I have said that we should demand that the demon produce only that non-qualitative information about other times which can be expressed using only the proper names we have already given him. And if we are going to demand that the demon produce *some* non-qualitative information, how could we reasonably demand more information than this? Think of proofs from some given set of premises in some formulation of the first-order predicate calculus. It is impossible to derive any sentence containing names which do not occur in the premises you are given.²⁷ How can we ask of the demon something that, as a matter of logic, cannot be done?²⁸

by saying it names the particle into which $\alpha+$ decays.

²⁶The use of Absolute space is not necessary; even if the history of the world unfolds in neo-Newtonian spacetime, we have names for currently existing points of spacetime, so we must be able to deduce whether the column will collapse toward the unique spacetime point that exists at the time of collapse and is co-located with p relative to such-and-such frame of reference. This, too, we cannot do.

²⁷Well, not exactly. But it is true that for any derivable sentence $\ulcorner \varphi(a) \urcorner$ containing a constant a not occurring in the premises, you can also derive $\ulcorner \forall x \varphi(x) \urcorner$; so the non-qualitative information expressible using new names that you can derive is not substantive.

²⁸In this argument, I assume that the laws of nature are purely qualitative; they do not, as it were, mention any particular individuals by name. There are theories of laws of nature which may allow non-qualitative laws. The best system theory of laws (Lewis 1983), in particular, may do this: in simple enough or strange enough worlds, the strongest and simplest system may be one containing theorems which mention some particular individual by name. I think my analysis could be amended to accommodate such laws.

3.3 A Formal Definition

I began with the Laplacian picture of determinism and distinguished three ways of making it more precise. The three ways agree that a deterministic world is one in which a powerful demon, when given complete qualitative information about one time and the laws of nature, can produce complete qualitative information about all other times. They disagreed over how much non-qualitative information we give the demon about the initial time, and how much non-qualitative information we demand he produce about other times. I favored the precisification according to which we give the demon all non-qualitative information about things that exist at the initial time, and demand that he produce all non-qualitative information *about just those things* at all other times.

But this talk of non-qualitative information and of what the demon can produce is still vague. I will now precisify it further. To see how, first consider how to precisify the version of Laplacian determinism that deals only with qualitative information. To give the demon all the qualitative information about a time is to tell him how many things exist at that time, which qualitative properties each of them instantiate, and for each n (where n might be infinite), which qualitative n -place relations any n -tuples of them instantiate. A world that is deterministic in this (purely qualitative) sense, then, is one in which the demon, given all qualitative information about a time and knowing the laws of nature, can tell us how many things there are in total, what qualitative properties each of them instantiate, and so on. And there will be some time slice of the world he describes that will match the initial time we describe to him, where ‘match’ here means ‘be a qualitative duplicate of.’ Although the demon can tell us how many things exist in this world, though, he cannot tell us *which* things exist, and cannot distinguish between worlds in which two things have switched roles. Dispensing with the demon, this means that a world w is deterministic (in this purely qualitative sense) just in case any other world that is physically possible relative to w and is a qualitative duplicate of w at a time is a qualitative duplicate of w full-stop.

We have almost arrived at David Lewis’s analysis of determinism. Lewis’s analysis makes use of an analysis of qualitative duplication; and this analysis itself

makes use of the distinction between properties that are perfectly natural and those that are not. The perfectly natural properties (and relations) are the fundamental properties (and relations); wherever any other property or relation is instantiated, it is instantiated *in virtue of* the global pattern of instantiation of the fundamental properties and relations. The set of perfectly natural properties and relations, then, forms a supervenience base for the set of all properties. (I will also assume that the perfectly natural properties are purely qualitative.) Two things are duplicates, then, iff they share all the same perfectly natural properties, and their parts can be put into correspondence so that corresponding parts share the same perfectly natural properties, and corresponding pairs of parts stand in the same perfectly natural (two-place) relations (and so on). Lewis's analysis of determinism, then, is

(D0) A possible world w is deterministic iff every world that is a duplicate of w at a time and is physically possible relative to w is also a global duplicate of w .

The column world is a counterexample to Lewis's analysis. We want a Lewis-style analysis of determinism that corresponds to my modified version of the Laplacian picture in the way that Lewis's corresponds to the version that takes only qualitative information into account.

To give the demon all the qualitative information about a time *and* all the non-qualitative information about things that exist at that time is to tell him not just how many things exist at that time, but also which things exist at that time, and what qualitative properties each instantiates at that time (and so on). So the demon can distinguish between times that are qualitative duplicates in which some things have switched roles.

A deterministic world, then, is one in which the demon, given the relevant information about a time and knowing the laws of nature, can tell us not just how many things there are in total and what qualitative properties each instantiates (and so on), but can also tell us which of them are the things that exist at the initial time. So the demon can tell us what roles the things that exist at the initial time are playing in the global structure of the world. And there will be some time slice of the world he describes that will match the initial time we described to him; and match not just because it is a qualitative duplicate, but also because the same things exist and play just the same roles.

To express this as a Lewis-style definition I need the concept of a duplication function. Recall that two things are duplicates iff their parts can be put into a correspondence meeting certain conditions. Call such a correspondence a ‘duplication function.’ (Two things can be duplicates according to more than one duplication function: think of two congruent equilateral triangles.) The analysis is:

(D1) A possible world w is deterministic iff every world that is a duplicate of w at a time and is physically possible relative to w is also a global duplicate of w , *under a duplication function that is the identity function on the initial times.*²⁹

Worlds at which general relativity is true are deterministic on this analysis, as they were on Lewis’s. (The diffeomorphism that generates the hole is a duplication function; since it only changes which spacetime points play which roles in the future, it is the identity on the initial times). But the collapsing column world is indeterministic. Again, let ‘ p ’ and ‘ q ’ name points of space that lie in different directions from the column. There is a world w in which the column collapses toward p and

²⁹I’m assuming that we can speak of two things existing in more than one possible world. Counterpart theorists might be nervous about such talk of transworld identity appearing in an analysis of ‘determinism,’ but the analysis can be easily re-written in terms of counterpart relations: say that

w is deterministic iff for any world w' that is physically possible relative to w and duplication function d , if a time slice of w and a time slice of w' are duplicates under d , then w and w' are duplicates under some duplication function d^* that agrees with d on the initial time slices.

Here the function d is a counterpart relation between the two time slices. It provides the standard for identifying things that exist on the time slice in w with things that exist on the time slice in w' .

(D1) first appeared in (Belot 1995), though he formulates it in counterpart-theoretic terms, as does Melia (1999). Belot does not endorse (D1), but Melia does. Melia does (and Belot does not) attempt to show that there is something natural and intuitive behind this formal analysis, that it is not just an *ad hoc* device for avoiding the conclusion of the hole argument. But Melia’s intuitive characterization is different from mine. He appeals to branching possible worlds. But at its best this will only allow us to explain senses of determinism in which the past and the laws fix the future. My characterization can be generalized to other senses of determinism.

a world w^* in which it collapses toward q .³⁰ So a duplication function between w and w^* must map p to q . But p and q are distinct and exist at t , before the tower collapses. So no duplication function from w to w^* can act as the identity on t .

4 Earman's and Stein's Argument

In Chapter 3 of *World Enough and Space-Time* Earman (following Stein (1977)) argues that relationalism about motion, together with the possibility of determinism, entails relationalism about ontology.

Pause to consider how implausible this is. Suppose relationalism about motion is true, and so that the world contains a Machian spacetime. (The discussion to follow is unchanged if we consider Leibnizian spacetime instead.) Suppose that the one law governing particle behavior is this: inter-particle distances never change. The law looks deterministic: if I know there are just two particles, and that they are two feet apart right now, I know all there is to know about the future: there will continue to be two particles, two feet apart. Of course I won't know if they're in the same place later as they are now, or whether they're in the same place relative to some frame of reference as they are now. But I couldn't know these things, because in Machian spacetime there is no sense to be made of 'same place across time,' even relative to some frame of reference.

Earman's argument appeals to

(SP2) Any spacetime symmetry of theory T is a dynamical symmetry of T .

Rotations are symmetries of Euclidean space; by performing a (possibly different) arbitrary rotation on each instantaneous Euclidean space in Machian spacetime we obtain a symmetry Φ of that spacetime. If we suppose that spacetime is populated

³⁰There is a debate in modal metaphysics over whether distinct but qualitatively indiscernible possibilities really require distinct possible *worlds*. (Some defenders of counterpart theory, like David Lewis (1986b), say no.) This makes no difference to my argument; my argument works even if w and w^* are the same world, though in some cases I would have to phrase the argument in terms of counterpart theory (see footnote 29 above).

by particles, then this symmetry is a dynamical symmetry of theory T just in case for any model M there is another model M_Φ with the same spacetime manifold such that a point x in the spacetime of M is occupied by a particle iff $\Phi(x)$ is occupied in the spacetime of M_Φ . In his argument Earman presupposes (what I have complained about earlier) that substantivalists must say that distinct models of T correspond to distinct possible worlds.

The argument goes like this. Suppose relationalism about motion and substantivalism are both true. Then spacetime has a Machian structure. Suppose there are some particles obeying some laws. The function Ψ on Machian spacetime which is the identity for all times before and including t but rotates each of the later instantaneous Euclidean spaces through some angle about some axis after t is a symmetry of Machian spacetime. By (SP2), Ψ is also a dynamical symmetry. So the state of the world at t plus the laws fails to fix the state of the world after t : they fail to determine whether the particles are where they actually are after t , or are where they would be if rotated through some angle about some axis. That is, if relationalism about motion and substantivalism are both true, then no theory of particle motion is deterministic. If relationalists about motion want to allow for the possibility of determinism, they must also be relationalists about ontology.

These two futures at work in Earman's argument differ merely non-qualitatively: they differ merely over which spacetime points in the future the particles occupy. Earman's argument presupposes that the existence of futures that differ in this way is enough to render a theory indeterministic. As I have argued, this is not so. Ψ is a global duplication function that acts as the identity on the original times; so the distinct futures generated by it do not render the theory indeterministic, according to the analysis of determinism I have defended. Relationalism about motion does not require relationalism about ontology; and good thing, too, since I argued the opposite in chapter 1.

5 Conclusion

I have argued that the hole argument fails because its second premise is false. That premise may seem true at first, but only when read with an incorrect analysis of

‘determinism.’ I have articulated and defended a better analysis of ‘determinism,’ one that is a natural precisification of the Laplacian conception of determinism, that has intuitive appeal considered on its own, and that fits our intuitions about several cases. And I have shown how this modified Laplacian conception can be given formal expression in terms of possible worlds.

Chapter 3

Supersubstantivalism

The best current thinking does not claim that particles are *not* built out of spacetime....For the time being, as a means to get on with the world's work, and to deal with particles on a practical working basis, it makes sense to *treat* particles as if they are foreign objects. This working procedure does not exclude any longer-term possibility to account for a particle in terms of geometry—as one today accounts for the eye of a hurricane in terms of aerodynamics, and the throat of a whirlpool in terms of hydrodynamics. (Wheeler and Taylor 1963, page 193)

1 Introduction

Substantivalists believe, and relationalists deny, that spacetime exists. Set relationalism aside; substantivalists still have plenty to disagree about.

There is room for them to disagree about the nature of material objects. Some substantivalists—call them ‘dualists’—hold that material objects are distinct from spacetime. Others—‘supersubstantivalists’, as Sklar calls them¹—hold that material objects (if there are any) are identical with regions of spacetime. Supersub-

¹In *Space, Time, and Spacetime* (Sklar 1974). I believe that Sklar coined this term; it first appears on page 214.

stantivalists identify material objects with the regions of spacetime that dualists say those material objects occupy. Dualists, then, like their namesakes in the philosophy of mind, accept additional ‘ontological categories’: just as mind-body dualists hold that minds differ in kind from physical things, so dualists in this debate hold that physical objects differ in kind from points and regions of spacetime. As John Wheeler puts it, dualists regard material objects as ‘strange and nongeometrical objects immersed in spacetime’ (Wheeler and Taylor 1963, page 191).

Supersubstantivalism may initially seem incredible, but that has not stopped some very smart people from believing it. Descartes identified body with extension, and so looks like a supersubstantivalist. Bennett (1984) attributes supersubstantivalism to Spinoza. More recently, Ted Sider (2001) briefly defends supersubstantivalism in the course of arguing for the doctrine of temporal parts.² And supersubstantivalism is not just a metaphysician’s fantasy. Respected physicists have held views close to or identical with it: Isaac Newton was a supersubstantivalist, and so was Einstein during one period of his career. And the many physicists who believe that there are only fields (like the 19th century physicists who identified charged particles with parts of the electromagnetic field³) are close to being supersubstantivalists. This is an impressive catalog.

There are several varieties of supersubstantivalism. In the next section I’ll survey these varieties, and then I’ll look at arguments for and against the view.

2 Varieties of Supersubstantivalism

Dualism is straightforward: there is a special fundamental relation, the occupation relation. Every concrete material object bears the occupation relation to some region of spacetime.

Generic supersubstantivalism is also straightforward: there is no such fundamental relation as the occupation relation. Every concrete material object is identical with some region of spacetime. But there are more and less radical varieties of supersubstantivalism.

²Other philosophers who flirt with supersubstantivalism include David Lewis (1986b), Hartry Field (1989), and W. V. Quine (1981).

³The ‘electromagnetic view of nature’ is discussed in (McCormmach 1970).

Different versions of supersubstantivalism give different accounts of what fundamental properties and relations points and regions of spacetime instantiate. At one extreme lies the view that the only fundamental properties and relations spacetime instantiates are geometrical ones. Hard-headed scientifically-minded philosophers are likely to respect this version most, because it is linked to a scientific research program in a way that the others are not. (Sklar considers this version a worthy topic of discussion, while dismissing a more modest version as a ‘linguistic trick’ (1974, pages 166, 233).) Inspired in part by the general theory of relativity, John Wheeler advocated this variety of supersubstantivalism. In general relativity Einstein eliminated the gravitational force: the motions that this force explained are in Einstein’s theory explained instead by the curvature of spacetime. Wheeler hoped that *all* forces could be eliminated, and *all* motion explained by the curvature of spacetime. (It’s this hope that he expresses in the quotation that heads this paper.) Of course, the sense in which such a theory ‘explains motion’ is now a little funny: if we’re supersubstantivalists, we don’t think there are particles distinct from spacetime, particles that move by being located in different places at different times, whose motions we can attempt to explain. When an extreme supersubstantivalist claims to explain the motion of bodies he means something like this: the whole story of the history of the universe is contained in the geometrical structure of spacetime; there are laws governing that structure which allow us to explain why a given instantaneous space has the geometry it does by citing the geometry of earlier instantaneous spaces together with the laws; and ordinary talk of the motion of bodies is in *some* sense talk of the curvature of spacetime.⁴ The Grand Unified Theory that Einstein sought was to be a theory that provided explanations like this: the theory would not postulate (dualistically-conceived) particles, but certain regions of spacetime in possible worlds permitted by the theory would look ‘particle-like’ (see (Earman 1995, page 16)). Producing such explanations is a real challenge, one

⁴Not all spacetimes permitted by general relativity are well-behaved; some cannot be divided up into a series of instantaneous spaces ordered in time. In such spacetimes we may not be able to explain why a given spacelike region of spacetime has the structure it does by appealing to the structure of some earlier spacelike region and the laws.

that has not yet been met: while Einstein showed how to ‘reduce’ gravity to curvature, no one has succeeded in reducing electromagnetism, or the other forces, to curvature in a plausible way.

Although the success of general relativity helped motivate this research program, its roots go back farther. Back in 1870 W. K. Clifford, inspired by Riemann’s theory of curved surfaces, hypothesized that what we ordinarily call the motion of matter is nothing but ‘variation of the curvature of space,’ and that ‘in the physical world nothing else takes place but this variation.’ Like Wheeler, Clifford hoped to produce scientific explanations of phenomena that appeal only to the geometry of space: he hints that ‘I am endeavoring in a general way to explain the laws of double refraction on this hypothesis’ (1882, page 22). Even earlier Newton secretly toyed with a version of supersubstantivalism according to which the only fundamental properties space instantiates are geometrical, though he makes no attempt to use the curvature of space to explain motion. (Newton thought space was flat, and in his time the geometry needed for dealing with surfaces with variable curvature—Riemannian geometry—had not been invented. I’ll look at Newton’s theory below.)

I’ve been discussing extreme versions of supersubstantivalism. Other versions are more modest. More modest versions are more liberal about what fundamental properties they allow spacetime to instantiate. One modest version accepts as fundamental the (non-geometrical) fundamental properties we ordinarily think that fundamental particles instantiate—properties like mass and charge—and allows spacetime to instantiate them.

This more modest version does not face the daunting challenge that the extreme version does: that of rewriting the laws of electromagnetism and the other forces as laws governing the curvature of spacetime. In fact, existing spacetime theories (like Newtonian gravitational theory and electromagnetism) can easily be reinterpreted to make them compatible with the more modest form of supersubstantivalism.

3 Is Dualism the Default View?

Strictly speaking, dualism and supersubstantivalism are not incompatible. Both are true in a world with spacetime but no material objects. But our world is not one of

these worlds, so we must choose.

Dualism seems to most substantialists like the default position. They want powerful arguments showing that dualism is incoherent, or that it conflicts with some other strongly held beliefs, before they take supersubstantialism seriously. I don't know of any arguments like that; so if that is what is required, supersubstantialism is in bad shape.

How does dualism get this default position? Not because we can *just see* that material objects aren't regions of spacetime. Maybe substantialists give dualism default status because it is entrenched: maybe substantialists have believed it for a long time, are used to picturing the world in its terms, and have never really considered the alternative. But psychological entrenchment does not justify giving dualism default status.⁵ If preferring supersubstantialism requires strong arguments against dualism, that must be because there are good reasons in favor of dualism in the first place. But what are these reasons?

I can think of two reasons we give dualism default status, and substantialists should not think either are any good.

First, we might oppose supersubstantialism because we feel some motivation for relationalism. I'm sitting in my chair and my bamboo plant is across the room, and (it seems) there is absolutely nothing between us. But if supersubstantialism is correct then we are both swimming in a sea of regions of spacetime, our borders contiguous with other things made of the same stuff that we are. So supersubstantialism conflicts with the way things ordinarily seem.

Even if this is a correct account of the way things ordinarily seem—and surely sometimes it is—substantialism already conflicts with it. And I do not see that dualism comes closer to matching the ordinary appearances than supersubstantialism.

Second, we might oppose supersubstantialism because it sounds odd to say that I bought my coffee from a region of spacetime. But, in general, the demise of ordinary language philosophy has taught us to be suspicious of arguments that start with 'it would sound odd to say....' And in particular, history should make us

⁵Maybe those who accept 'conservative' accounts of belief revision (like Harman (1986)) will disagree.

suspicious: there was a time when it sounded odd to say that I bought my coffee from a swarm of elementary particles.⁶

There are, in addition, more sophisticated philosophical reasons substantialists might oppose supersubstantivalism. Maybe supersubstantivalism is false because material objects move but regions of spacetime do not; maybe it is false because material objects lack but regions of spacetime have temporal parts; maybe it is false because my bamboo plant could have been located elsewhere, but no region of spacetime could have been located elsewhere. These are important objections, and I address them below. But I don't think they capture the reasons why we initially recoil from supersubstantivalism and give dualism default status.

I conclude that substantialists should not give dualism default status. We should consider dualism and supersubstantivalism as competitors on an equal footing, and see where argument leads.

So what reasons are there to be supersubstantialists? In sections 4 through 7 I discuss instances of two strategies for defending supersubstantivalism, one due to Descartes, and the other due to Newton.

4 Cartesian Supersubstantivalism

Descartes is typically classified as a relationalist. But while Descartes was certainly a relationalist about motion, it is less clear that he was a relationalist about ontology, rather than a supersubstantivalist. (I argued in previous chapters that relationalism about motion is not only consistent with substantivalism, it requires it.) In the *Principles of Philosophy* Descartes writes, 'in reality the extension in length, breadth and depth which constitutes a space is exactly the same as that which constitutes a body' (2:10).⁷ That sounds like a strong (and not particularly plausible) version of supersubstantivalism: not only is every material object (or 'body') identical with some region of space, but every region of space is a material object. (Since Descartes lived in a pre-relativistic age, this version of supersubstantivalism identifies material objects with regions of space, not with regions of spacetime.)

⁶(Sider 2001, sect. 4.8).

⁷Unless otherwise noted, all citations of Descartes are to the *Principles* (part:section).

Is Descartes really a supersubstantialist? Field (1989) claims that supersubstantialists can be distinguished from relationalists by asking the following question: is it possible that there be empty regions of spacetime? Only supersubstantialists will answer ‘yes.’ Descartes says ‘no,’ and so counts as a relationalist by this criterion. But I’m not sure this criterion is a good one. A dualist might think that it is necessary that every region of spacetime is occupied by a material object. Surely a supersubstantialist can accept the analogous claim, that it is necessary that every region of spacetime is a material object. Still, it may be that there is a merely semantic difference between some relationalists and supersubstantialists—they just differ over how they are willing to use the words ‘region of space(time).’ I will set this issue aside and speak as if Descartes is properly classified as a supersubstantialist.

Descartes’s argument for Cartesian supersubstantialism rests on the claim that ‘the nature of body consists in...its being something which is extended in length, breadth and depth’ (2:4). If we say that talk of something’s nature is talk of the conjunction of all of its essential properties (and this does seem to be what Descartes has in mind in 2:11), then Descartes’s claim amounts to this:

- (1) Being spatially extended is the only essential property of material objects.

Descartes also accepts a partial converse of (1), namely

- (2) Necessarily, every extended thing is a material body.

But the arguments he gives in 2:4 and 2:11 (which I examine below) are clearly arguments for (1), and without extra premises (1) does not entail (2). Nor is (2) sufficient to establish supersubstantialism: it entails that every region of space is a material body, but not that every material body is a region of space.

In 2:11 Descartes says ‘Yet this [namely, the idea of something that is extended in length, breadth, and depth] is just what is comprised in the idea of space.’ I take this to mean

- (3) Being spatially extended is the only essential property of regions of space.

This looks plausible. And we can get from (1) and (3) to Cartesian supersubstantialism if we add another premise:

(4) Distinct kinds of things cannot share all of their essential properties.

So there is only one kind of thing that has spatial extension as its only essential property; and we have two names that apply to (all and only) things of this kind: ‘region of space’ and ‘material object.’⁸

This argument aims to establish a particular version of supersubstantivalism. So we can evaluate it from two points of view. First, does it give us reason to accept supersubstantivalism in the first place? Second, if we are already supersubstantialists, does it give us reason to accept Cartesian supersubstantivalism instead of some other version? I will take these questions in reverse order.

Cartesian supersubstantivalism is hardly the most plausible version of supersubstantivalism. Supersubstantialists should allow for the possibility of regions of space that are not also material objects. So which premises do they reject? Not the premises about essential properties: presumably supersubstantialists agree that it is essential to regions of space (and so to material objects) that they be spatially extended. They should deny the (suppressed) premise that material objects are a ‘kind of thing,’ things that can be characterized in terms of their essential properties. Instead, say that material objects differ from other regions of spacetime in their accidental properties—like, for example, mass and charge.

Of course this reply is not open to dualists; dualists do not think material objects differ from regions of space only accidentally. So (turning now to my first question), how do dualists resist this argument?

Dualists should reject (1). Descartes’s argument for (1), as it occurs in his discussion of the stone in *Principles* 2:11, is to cast about for a counterexample and fail to find one.⁹ He thinks that for any other property some material object has,

⁸One quick objection to this argument is to say that material objects can be point-sized, and so have no spatial extent at all. But in that case the same is presumably true of regions of space—if there can be point-sized material objects there can be point-sized regions—and the argument could be modified to take this into account. Descartes would not accept this new argument, though, since he thought it impossible for a material object to lack all spatial extent. And this plays a crucial role in his argument against atoms in 2:20.

⁹In 2:4 Descartes argues that no perceptible properties, like hardness and color, are essential to material objects. Something similar goes on in Descartes’s discus-

there can be a material object that lacks it; and so no other property can be essential to material objects.¹⁰

A dualist might say that mass is an essential property of material objects. (Descartes, of course, lived before any such property was recognized.) But this can't be right, because there are massless particles. Bennett (2001, pp.30-32), following Locke and others, objects that Descartes fails to rule out solidity, or impenetrability, as an essential property of bodies. But I'm not sure that impenetrability is essential to material objects. Neutrinos may have some power to exclude other particles from occupying just the region they do, but it is very weak (that is why they are so hard to detect); do we really think it is impossible for there to be a kind of particle that is completely uncoupled from other kinds, and does not interact with them at all? In any case, dualists have a better reason to reject (1). They should say that it is essential to material objects that they occupy regions of space.¹¹ But it is not essential to regions of space that they occupy regions of space. So these two kinds of spatially extended thing do not share all their essential properties.

Descartes's argument for supersubstantialism is not very good, but I think there is a better argument for supersubstantialism that is Cartesian in spirit.

5 Spatiotemporality

Descartes thought there couldn't be two kinds of spatial thing. I think that reflection on what it is to be a spatial (or, better, spatiotemporal) thing can still give us some reason to be supersubstantialists, even if not the reason Descartes thought.

A virtue of supersubstantialism, to the eyes of a metaphysician, is that it provides the materials for a better account of spatiotemporality than does dualism. This gives us some reason to be supersubstantialists.

sion of the wax in the second meditation; but there his primary aim is to argue that the nature of the wax 'is in no way revealed by my imagination, but is perceived by the mind alone' (page 85).

¹⁰For most of the passage he argues that for any property a stone has, that very stone could lack it. But when he argues that heaviness is not essential to the stone, he cites the fact that fire (which he regards as a material body) is not heavy.

¹¹Markosian (2000) defends something like this view.

What is it to be a spatiotemporal thing? What is it to be extended in space and time? Supersubstantialists have a simple answer: to be spatiotemporal is to be a region of spacetime, to be part of spacetime.

What do dualists say? They will say that to be spatiotemporal is either to be a region of spacetime, or to bear the occupation relation to some region of spacetime. But this view has several disadvantages.

First, there is a problem about essences. It is essential to material objects that they be spatiotemporal. For dualists this means: for anything at all, if it occupies some region of spacetime, then it is necessary that it occupies some region of spacetime. But this is to accept a necessary connection between distinct existences. Supersubstantialists do not need to accept this necessary connection.

Second, there is a problem of inclusion. Which things are spatiotemporal? Material objects certainly are. But some material objects are parts of other material objects; and composite material objects are located where their parts are. That is, it is necessary that if x is part of y and x occupies region L then y also occupies L .¹² Why is this necessary? Dualists can give no answer. They cannot explain the necessary connection between parthood and occupation. Supersubstantialists do not believe in an occupation relation, and so have no necessary connection to explain. (Or, better, supersubstantialists do believe in an occupation relation, but they deny that it is a fundamental relation, and instead analyze it in terms parthood: a region of spacetime occupies just those regions that are parts of it. On this account, there is no mystery in the necessity that material objects occupy those regions that their parts occupy.¹³)

¹² L need not be the largest region y occupies. Some might accept a weaker condition than the one in the text. They might say that it is necessarily only that if x is part of y and x occupies L then y occupies a region that contains L . The difference between these two does not affect my argument.

¹³I'm oversimplifying the way we use 'occupies' and our talk of locations and places generally. Supersubstantialists will say that sometimes we use 'occupies' to express the converse of the parthood relation; sometimes we use it the way relationalists think we always use it: to express some complicated relation to other material objects, as when I stand on the subway platform and say, 'every morning I stand in the same place and wait for the train' to mean something like 'every morning I stand next to this wall and wait for the train.' This will come up again below.

Third, there is a problem of exclusion. Some philosophers want to expand the realm of the spatiotemporal. Penelope Maddy, for example, claims that sets are spatiotemporal. She uses this claim to defend the existence of sets by arguing that we can see sets, and so that we have empirical evidence for their existence.¹⁴ It is easy for supersubstantialists to reject Maddy's claim. They say that only regions of spacetime are spatiotemporal, and no set is a region of spacetime. It is not so easy for dualists. They must argue that sets cannot bear the occupation relation to regions of spacetime. But it is open to someone like Maddy to say: dualists admit a necessary connection between occupation and parthood. I admit an additional necessary connection between occupation and membership: sets occupy just those regions their members occupy. I cannot explain why there should be this connection; but neither can dualists explain why there should be a connection between occupation and parthood.

6 Newtonian Supersubstantialism

Descartes hoped to derive supersubstantialism from premises we should all accept. That's an ambitious project. There is a less ambitious way to argue for supersubstantialism. Instead of arguing that material objects *must* be regions of spacetime, we argue that they *can* be. Regions of spacetime can do all the work we need material objects to do. For the sake of simplicity, then, we should identify material objects with regions of spacetime.¹⁵

Newton took this approach. His manuscript 'De Gravitatione' contains an attempt to describe the nature of material objects.¹⁶ He ends up describing a possible world in which, he claims, it would be appropriate to say that certain regions of space are (also) material objects. Newton proceeds in stages.

First he asks us to imagine the following scenario. Suppose the dualism is true, that material objects really are distinct from spacetime. God chooses some region of empty space—call it 'Joe.' Whenever a material object is about to occupy

¹⁴(Maddy 1990), chapter 2, especially pp.58-59.

¹⁵(Sider 2001) and (Bennett 1984) also give this motivation for supersubstantialism.

¹⁶(Newton 1962). Howard Stein discusses this manuscript in his (1970).

Joe, God exerts a force on that object in the direction away from Joe sufficient to prevent it from occupying Joe. So when they get close enough to Joe, superballs bounce right back; photons are reflected; and beer bottles shatter. Over time God may change which region he watches over in this way; we may assume that he changes it continuously and that all the regions have the same shape. (We may be tempted to say, as God varies which region of space enjoys his favor, that Joe is *moving*. But this is just a misleading appearance; Joe is a region of space, and Newton denies that regions of space can move.)

Of course even in Newton's physics material objects act on each other in ways other than by contact; but we can easily imagine God also causing each material object to be accelerated towards Joe, that acceleration being indirectly proportional to the square of the distance between that object and Joe (or the distance between that object's center of mass and some point that falls inside Joe). Material objects in this world would move, then, just as if there were some material object occupying Joe and attracting them gravitationally.

Of this scenario Newton asserts, 'it seems impossible that we should not consider this space [that is, Joe] to be truly body from the evidence of our senses....for it will be tangible on account of its impenetrability, and visible, opaque and colored on account of the reflection of light, and it will resonate when struck because the adjacent air will be moved by the blow' (139).

Now it is doubtless true that were we in such a world, confined to exploring it with our senses, we would conclude that Joe is 'truly body'; but since we are not so confined, but know more directly the metaphysical truth of the situation, it is clear that Joe is no material object. A world in which God behaves in the way described is a world in which God deceives some of his creatures. (Descartes would be appalled.)

The problem is that while some regions of space—Joe, in particular—are masquerading as material objects, there are also, in the world, *real* material objects—that is, material-objects-according-to-dualists. Our use of 'material object' is not so permissive that it could apply to both regions of space and things distinct from regions of space in the same world (considered as actual). There's an obvious way to avoid this problem: Newton next asks us to imagine a world that contains none of

what dualists call ‘material objects,’ but instead contains only space, certain regions of which God has chosen to watch over as he watched over Joe in the previous scenario. (Of course, now watching over a given region involves ensuring that no other chosen region overlaps it, rather than ensuring that no dualist material object occupies it.) In such a world, Newton (perhaps) asserts—he declines to wholeheartedly endorse the thesis—supersubstantialism is true.¹⁷

Newton’s assertion is implausible, even to someone sympathetic to supersubstantialism. In the first scenario, the one in which there were also dualist material objects, there was *something there* for God to repel away from his chosen empty region of space. Without those things, though, in what way do the chosen regions of space differ from the unchosen regions? They do not differ in their intrinsic properties, since it is no part of God’s activity (as Newton has described it) to endow the chosen regions with special intrinsic properties. Nor does God endow these regions of space with the disposition to repel each other. For I take it that two things that differ in their dispositional properties must also differ intrinsically. Intuitively, a bar of lead does not become as fragile as glass just because a magician resolves to cast a spell causing it to shatter whenever anyone strikes it with a light blow; God, though more powerful, no more alters the dispositions of the regions of space than the magician alters the dispositions of the bar.

So the special regions differ from the ordinary ones only in their relational properties: some regions have, and others lack, the property of being chosen. Surely there must be more of a difference than this between regions which are also material objects and regions which are not. It can seem that all we have here is a world containing no material objects, only empty space, in which God has chosen to play a game with himself whereby he ‘outlines’ certain regions of space and changes which regions are outlined according to certain definite rules.¹⁸

I am not sure why Newton wanted God to do so little for the regions of space he has chosen to be material objects. Why have God merely watch over them, when God could give them the power to watch over themselves, by endowing them with certain intrinsic properties? Perhaps the following idea motivated Newton: all that

¹⁷That is, supersubstantialism is true *and* there are material objects.

¹⁸(Bennett 1984) makes a similar complaint.

really matters, when we describe a supersubstantialist world, is how the regions interact with other things; as long as they interact in the right sort of way, it doesn't matter what they are like 'in themselves.'¹⁹

7 Moderate Supersubstantialism

Newton's strategy, again, is to argue that regions of space can do all the work we need material objects to do. I object to the way he implemented this strategy: material objects do more than just prevent other material objects from occupying the regions they occupy. Let's keep the strategy and change the implementation.

The physical theories that have been considered fundamental since Newton postulate two kinds of material objects: particles and fields. What work do particles and fields do in these theories? I will start with the particles. Particles are the primary bearers of intrinsic properties like mass and charge; and they move around.

A moderate version of supersubstantialism can allow regions of spacetime to instantiate mass and charge. So that bit of work can be done by regions of spacetime as easily as by material objects. But what about motion? Can a region of spacetime move?

Versions of supersubstantialism that treat space and time as separate things—like Newton's—have trouble here. According to these versions, material particles are regions of space. Normally substantialists say that a particle moves if it occupies different regions of space at different times. But this won't work if particles just are regions of space; whatever sense we give 'occupies' in this context, a region of space cannot occupy different regions of space at different times.

This kind of supersubstantialist might do better by adopting a relationalist account of motion: a region moves just in case its distance from some other region changes in time. But it seems impossible for regions of space to move in this sense, either: it seems necessary that the distance between any two points of space be the same at all times. There is some conceptual room to maneuver around this

¹⁹Something like this line of thought occurs on page 140. Others have shared Newton's motivating idea that dispositions don't require categorical bases—Faraday (1965) inclined toward supersubstantialism for similar reasons, and Blackburn (1993) gives the idea a more general defense.

impossibility, but the moves aren't attractive. I'll mention them, just to set them aside.

We might give up mereological essentialism for regions of space, and say that one and the same region can include different points at different times. Then even if the distance between any two points of space is the same at all times, the distances between regions can change. Or we might accept mereological essentialism, deny that particles really move, and try to explain away appearances to the contrary. Different regions are particle-like (by instantiating mass properties and so on) at different times; since the pattern of mass-instantiation changes continuously, it looks like one persisting thing is moving around, but this is an illusion.²⁰

Luckily, we can avoid having to say either of these things by denying that space and time are separately existing things. Believe instead in a four-dimensional spacetime that unifies them. In this context dualists explain what it is for a particle to move by citing features of its worldline. So, for example, a particle accelerates just in case its worldline is not straight. Supersubstantialists who identify a particle with its worldline can say the same thing.

So much for particles. What about fields? The case for identifying the parts of physical fields (like the electromagnetic field) with regions of spacetime is even better than the case for identifying particles with regions of spacetime. In this case we don't have to worry about motion: whether and how the electromagnetic field moves is of no importance to the theory of electromagnetism.

Of course this was not always so. According to classical electromagnetism, light is a wave in the electromagnetic field. And many in the 19th century assumed that if there is a wave then something must be waving: something must be changing its shape, as the water in the pond changes shape when a rock falls in it. So there had to be some physical stuff that carried electromagnetic waves: the ether.

I think that even this theory can be given a supersubstantialist interpretation: identify each part of the ether with its worldline in four-dimensional spacetime, and

²⁰This is analogous to what mereological essentialists who aren't supersubstantialists say: strictly speaking, there are no persisting tables, but it looks like there are because different things are momentarily table-like in the same place at different times.

explain its motion in the same way the motions of particles are explained. But the ether interpretation of electromagnetism has been discredited. We no longer think that there can be a wave only if something physical is changing shape, and so we no longer need a moving ether.

8 The Modal Argument Against Supersubstantivalism

As I said in the last section, supersubstantivalism is best combined with a belief in four-dimensional spacetime, rather than in space and time separately. This version of supersubstantivalism entails the doctrine of temporal parts: every persisting thing is composed of instantaneous temporal parts. Arguments against the doctrine of temporal parts, then, are also arguments against supersubstantivalism. But I have nothing to add here to the debates over the doctrine of temporal parts.

The most serious remaining argument against supersubstantivalism is the modal argument:

- (1) I could have been three feet to the left of where I actually am right now.
- (2) No region of spacetime could have been three feet to the left of where it actually is right now.
- (3) So I am not a region of spacetime.

This argument should look familiar: it is structurally similar to the problem of motion for supersubstantivalists that I discussed above. There the problem is accounting for the apparent fact that one and the same material object can be in different places at different times; here the problem is accounting for the apparent fact that one and the same material object can be in different places in different possible worlds. For this reason, parts of my reply to the modal argument will resemble parts of my discussion of the problem of motion.

Supersubstantivalists could deny the first premise and try to explain away our tendency to accept it. They could say: a region three feet to the left of me could have been intrinsically just like I am right now. Maybe we are confusing this possibility with the possibility according to which *I* am three feet to the left.

This move doesn't look very plausible. What, then, about denying the second premise? Before looking at how to deny it, we need to clarify its meaning. What does talk of where a region of spacetime actually is mean? Supersubstantialists should insist it means something relational: for it to be possible that a region of spacetime be three feet from where it actually is at some time, there must be a possible world in which its distances from other regions at that time are different from what they actually are. So if region *A* is actually one foot to the left of region *B* at that time, then in the other possible world *A* must be four feet to the left of *B*, and so on.

Is this really possible? There are two ways in which this possibility might be realized. First, the points in *A* might be different distances from the points in *B* in the two possible worlds. Second, while all points of spacetime stand in the same geometrical relations to each other in both possible worlds, the points that compose *A* and *B* differ in the two worlds. So to deny (2) supersubstantialists must give up one of two doctrines. They must either reject (some version of) geometrical essentialism, the doctrine that points of spacetime have their geometrical properties essentially; or reject compositional essentialism for spacetime regions, the doctrine that regions of spacetime cannot contain different points in different possible worlds.²¹

I am happy to give up geometrical essentialism, for reasons I discussed in 2: it does not sit well with general relativity. To supersubstantialists who do accept both of the above doctrines I offer the following reply to the argument. They should say that our modal talk is context dependent. In some context the essentialist doctrines are true; in some context they are false. In the first kind of context, then, (1) is false and (2) is true. In the second kind, (1) is true and (2) is false. In no context are both premises true: when we are led to think both are true, it is because the context is shifting back and forth.

One way to explain this context dependence is to give a counterpart-theoretic analysis of these *de re* modal claims. In some contexts (contexts in which we are

²¹This is the doctrine that deserves the name 'mereological essentialism for spacetime regions;' but this name is already widely used for the claim that regions cannot contain different points at different times.

thinking about regions of spacetime as regions of spacetime) we use a counterpart relation that values geometrical similarity and makes these essentialist claims true. In other contexts (contexts in which we are thinking about regions of spacetime as material objects) we use a counterpart relation that values other kinds of similarity—like similarity with regard to mass and charge distribution—over geometrical similarity, and makes the essentialist claims false.²²

²²By ‘values geometrical similarity’ I mean more than that counterparts are geometrical duplicates; I mean also that counterparts of two regions stand in the same geometrical relations as the regions do.

Chapter 4

What Makes Time Different From Space?

1 Introduction

No one denies that time and space are different; and it is easy to catalog differences between them. I can point my finger toward the west, but I can't point my finger toward the future. If I choose, I can now move to the left, but I cannot now choose to move toward the past. And (as D. C. Williams points out) for many of us, our attitudes toward time differ from our attitudes toward space. We want to maximize our temporal extent and minimize our spatial extent: we want to live as long as possible but we want to be thin.¹ But these differences are not very deep, and don't get at the essence of the difference between time and space. That's what I want to understand: I want to know what *makes* time different from space. I want to know which difference is the fundamental difference between them.

I will argue for the claim that (roughly) time is that dimension that plays a certain role in the geometry of spacetime and the laws of nature. (In this paper, then, I focus on what is distinctive about time, and say little about what is distinctive about space.) But before giving the argument I want to put my question in slightly different terms. Instead of asking, 'what makes time different from space?'

¹(Williams 1951, page 468). Williams actually says that we care how long we live but do not care how fat we are.

I want to ask, ‘what makes temporal directions in spacetime temporal, rather than spatial?’² After rejecting some bad answers to this question I’ll present my view.

2 Spacetime Diagrams and Directions in Spacetime


It is often helpful, when approaching problems in physics and in metaphysics, to draw a spacetime diagram. Spacetime diagrams represent the careers in space and time of some material objects. Traditionally in a two-dimensional spacetime diagram (the easiest kind to draw on paper) the horizontal axis represents space and the vertical axis represents time. So suppose I’m confined to one dimension in space: I can only move to the left or to the right. Then the diagram in figure 2 might represent part of my career. The zig-zag line represents me; it’s my worldline. (I’m incorrectly represented as point-sized, but that’s not important.) According to the diagram I stand still for a while; then I walk to the left, stop, stand still for a little while longer, and then walk back to the right.

I said I wanted to ask what makes temporal directions in spacetime temporal, rather than spatial. So what is a direction in spacetime? We can use spacetime diagrams to get a sense for what directions in spacetime are. To represent a direction in spacetime (at some spacetime point) on the diagram we can draw an arrow, or vector, on the diagram at the point that represents that point of spacetime. So in figure 2 the arrow labeled ‘A’ points in the leftward direction in space and the arrow labeled ‘B’ points in the future direction in time. There are in this diagram, then, at least two temporal directions: toward the future and toward the past; and two spatial directions: toward the left and toward the right.

Two arrows may point in the same direction while being of different lengths. A direction then is an equivalence class of vectors—the set of all vectors that point

²It is important to distinguish this question from another commonly discussed question. Many philosophers want to know what makes the future different from the past. But that is not what I am asking. Toward the future and toward the past are both temporal directions, and I am not asking what makes one temporal direction the direction toward the future and the other, the direction toward the past. Instead I’m asking, what makes *either* of them a temporal, rather than spatial, direction in the first place?


Figure 4.1:



in the same direction and differ only in their length. Following standard usage, I will sometimes call a vector that points in a temporal direction a ‘timelike vector,’ and a vector that points in a spatial direction a ‘spacelike vector.’

What about the arrows labeled ‘C’ and ‘D’? They don’t seem to point in either a temporal or a spatial direction. What to say about arrows like C and D really depends on what geometrical structure the spacetime represented by the diagram has. In (two-dimensional) neo-Newtonian spacetime every arrow that does not point either to the left or the right points in a temporal direction, while in Minkowski spacetime (the spacetime of the special theory of relativity) arrows that are less

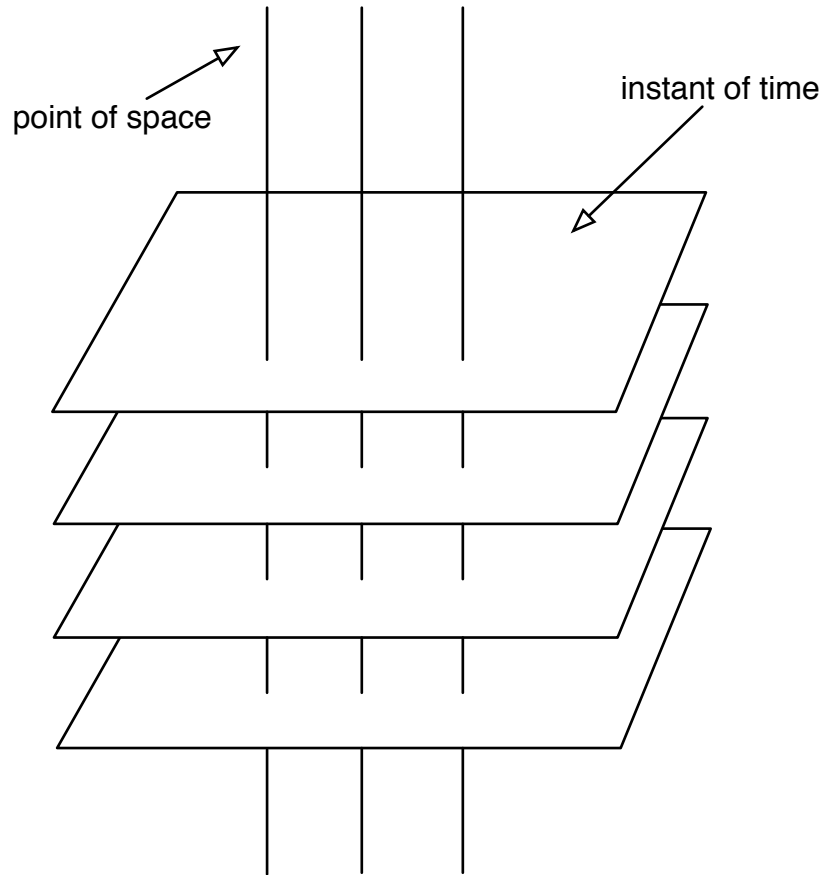
Figure 4.2:



than 45° from the vertical (like C) point in a temporal direction, while arrows that are more than 45° (like D) from the vertical point in a spacelike direction.

Why frame the discussion in terms of spatial and temporal directions, rather than space and time? Modern physical theories are formulated in terms of a four-dimensional spacetime, instead of in terms of three-dimensional space and one-dimensional time separately. In some older theories (Newtonian mechanics, in particular) there is a way to identify certain regions of spacetime as points of space and other regions as instants of time. But in more recently theories, especially the

Figure 4.3: Newtonian Spacetime



general theory of relativity, this cannot always be done.

We can identify points of space and instants of time with certain regions of Newtonian spacetime because Newtonian spacetime has certain special geometrical features (see figure 4.3). There is a unique and geometrically preferred way to divide up this four-dimensional spacetime manifold into a sequence of three-dimensional Euclidean submanifolds. Each of these Euclidean submanifolds is well-suited to play the role we want instants of time to play: events that occur on the same submanifold occur simultaneously. So these submanifolds *are* instants of time. A point of space, then, is a line in spacetime perpendicular to each time. (Events located on the same line, whether simultaneous or not, occur in the same

place.) Facts about which regions are points of space, and which are instants of time, are absolute: not relative to any observer or frame of reference.

In Newtonian spacetime the distinction between time and space and between temporal and spatial directions coincide. A vector that points in a spatial direction is one that points along (is tangent to) a time, and so points toward other points of space. A vector that points in a temporal direction is one that points at an angle to a time, and so points in the direction of future or past times.

But the special geometrical features that allow us to identify points of space and instants of time with certain regions of Newtonian spacetime are missing in other spacetimes. In neo-Newtonian spacetime, though there is a non-relative way to identify regions of spacetime with times, there is not a non-relative way to identify regions of spacetime with points of space. Different inertial observers will regard different events as occurring in the same place, and so will regard different lines in spacetime as points of space, without the geometry of spacetime privileging one of them over the others.³ This happens to time as well in Minkowski spacetime, the spacetime of special relativity.

In general relativity it gets worse. According to this theory the geometry of spacetime varies from world to world, depending on the distribution of matter in each world. In *some* of those worlds spacetime can be divided up into instants of time and points of space. (As before the times are three-dimensional submanifolds having certain geometrical properties, but in general relativity they need not have a Euclidean geometry. Points of spaces are curves in spacetime (they need not be straight lines) that run oblique to each time and meet some further geometrical conditions. Typically there are many ways to divide spacetime up, just as in Minkowski spacetime, none more preferred by the geometry than the others.) But in other worlds spacetime cannot be divided in this way at all. (Gödel's solution is an example.) In these worlds no regions of spacetime count as points of space or instants of time, not even relative to some inertial observer. But even in these worlds there is a distinction between the temporal and the spatial aspects of space-

³I said above that in Newtonian spacetime a point of space is a line in spacetime that intersects each time at a right angle. In neo-Newtonian spacetime we can no longer say which lines intersect each space at a right angle and which do not.

time, because we can still distinguish spacelike from timelike directions. (Given a point on the worldline of a conscious observer in one of these spacetimes, for example, we can still tell in which directions from that point his future mental episodes lie.) Since the distinction between directions is more general than the distinction between space and time, it is the distinction on which I want to focus.

3 Does Geometry Distinguish Temporal from Spatial Directions?

The view that the distinction between temporal and spatial directions is a geometrical one is a natural one to take when one studies spacetime theories. Those theories attribute one or another geometrical structure to spacetime, and when a given theory is explained, the distinction between temporal and spatial directions is usually explained in geometrical terms.

For example, according to the special theory of relativity, the geometry of spacetime satisfies the axioms of Minkowski geometry. That geometry allows us to assign lengths to vectors in spacetime. And the lengths of vectors that point in temporal directions have a different sign than the lengths of vectors that point in spatial directions. One kind has negative lengths; the other, positive.

We might hope to explain what makes temporal directions temporal by appealing to the signs of the lengths of vectors that point in those directions; but I don't see how this approach could work, for three reasons.

First, while looking at the signs of the lengths of the vectors may allow us to pick out two disjoint classes of vectors (namely, the class of vectors with negative lengths, and the class of vectors with positive lengths), it doesn't allow us to figure out which of these the class of vectors that point in temporal directions. For it is a matter of convention whether vectors that point in temporal directions get assigned positive or negative lengths; different textbooks adopt different conventions. And the distinction between temporal and spatial directions is not a conventional one.

Second, a difference in the signs of their lengths is too formal and abstract to be the fundamental difference between vectors that point in temporal and those that point in spatial directions, and so between time and space.

And third, distinguishing between the two kinds of directions in terms of the

signs of vectors' lengths only works in spacetimes with Minkowski geometry. No vectors have negative lengths in neo-Newtonian spacetime, but we do not want to say that no directions in that spacetime are temporal. Of course in that spacetime the geometrical distinction between spatial and temporal directions is explained differently. Roughly speaking, in neo-Newtonian spacetime we have two ways to assign (positive) lengths to vectors. We are told that vectors that point in timelike directions are those that have non-zero length according to a particular one of those ways. (The same is true in Newtonian spacetime.) But it won't do to say that what makes a timelike vector timelike is that it satisfies the following condition: either it is a vector in Minkowski spacetime and it has (say) negative length, or it is a vector in neo-Newtonian spacetime and it has positive length according to one particular metric (and so on with clauses for each different spacetime). For temporal directions in a world with one spacetime geometry have *something in common* with temporal directions in a world with some other spacetime geometry. And an answer to the question, 'in virtue of what are temporal directions temporal?', must tell us something about what they have in common. Even if the current proposal correctly distinguishes temporal from spatial directions, it doesn't say what temporal directions have in common. So it is not the answer we're looking for.

While we have not yet found a geometrical way to distinguish timelike from spacelike directions, we do have a geometrical way to distinguish directions that are *either* spacelike or timelike from those that are neither. It is a matter of convention whether we say that timelike vectors in Minkowski spacetime have positive or negative lengths. But it is not a matter of convention that they (along with spacelike vectors) have *non-zero* length. And we can establish that vectors with zero length are neither spacelike or timelike on geometrical grounds alone. The zero vector, which points in no direction at all, has zero length, but in some spacetimes other vectors do as well. In particular, 'lightlike' vectors in Minkowski spacetime have zero length. (These vectors point along possible paths of light rays.) Whichever spacetime geometry we look at, that geometry assigns lengths to vectors as a way to assign distances, either spatial or temporal, between spacetime points. (This length is determined by adding up (really, integrating) the lengths of vectors tangent to a certain path between the two points.) But adding up a bunch of zeros just

gives zero, so adding up the lengths of vectors with zero length couldn't be a way to determine the temporal or spatial distance between two points.

We have, then, a geometrical way to divide the vectors in any spacetime into the class of vectors that are either spacelike or timelike, and the class of vectors that are neither. And we also have a geometrical way to divide the class of vectors that are either spacelike or timelike into two subclasses. (In Minkowski spacetime (and in the spacetimes of general relativity as well) we divide them into the subclass with negative lengths and the subclass with positive lengths. In Newtonian and neo-Newtonian spacetime we divide them into the subclass with positive length according to one way of assigning lengths, and the subclass with positive lengths according to the other way of assigning lengths. (Let's say that two vectors that belong to the same subclass are 'of the same kind.' Talk of vectors that are of the same kind, then, is reserved for vectors that are either spatial or temporal.)) But we don't yet have a way to designate one of those subclasses as the class of vectors that point in temporal directions.

4 Dimensionality

If we're looking for a geometrical way to distinguish temporal from spatial directions, dimensionality considerations are probably our best bet. In four-dimensional Minkowski geometry, whichever convention about the signs of lengths one uses, time ends up being one-dimensional and space ends up being three-dimensional. Perhaps it is *because* it is one-dimensional that time is time.

Before examining this thesis I'll say something about what it means to say that time is one-dimensional.

Intuitively speaking, to say that time is one-dimensional is to say that we can represent time as a line, and that all events that occur in time can be assigned a position on that line.⁴

As I said above in section 2, in Newtonian and neo-Newtonian spacetime there is a unique geometrically preferred way to slice up the four-dimensional

⁴Circles are one-dimensional too, so strictly speaking time could be one-dimensional even if we had to represent time as a circle.

spacetime into a one-dimensional sequence of three-dimensional Euclidean submanifolds. Each three-dimensional submanifold is a time, and the sequence gives their temporal ordering. Since every event occurs somewhere in spacetime, every event occurs somewhere in this one-dimensional sequence. So time is one-dimensional in the intuitive sense in these spacetimes because it divides up in this way.

In Minkowski spacetime there are many geometrically preferred ways to slice up the four-dimensional spacetime into a sequence of three-dimensional Euclidean submanifolds. We can still, then, temporally order all events on a line, even though there is no unique way to do so. (There are pairs of events x and y such that x occurs before y according to one slicing but y occurs before x according to another slicing.)

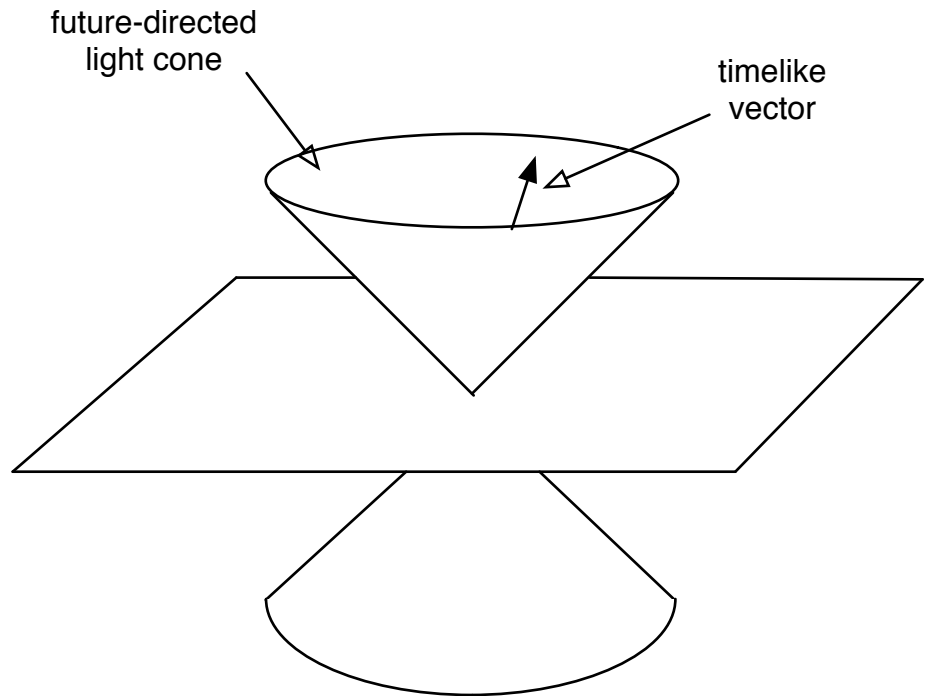
But we do not want to say in general that time is one-dimensional just in case spacetime divides up naturally into a one-dimensional sequence of three-dimensional submanifolds. For, as I mentioned, some general relativistic spacetimes cannot be divided up into a sequence of three-dimensional submanifolds. And there is still a (somewhat technical) sense in which time is one-dimensional in worlds containing those spacetimes.

The vectors in spacetime at any given spacetime point form a four-dimensional vector space; the maximum dimension of a subspace containing only timelike vectors (and the zero vector) is one, while the maximum dimension of a subspace containing only spacelike vectors (and the zero vector—I'll leave this implicit from now on) is three.⁵ Time is one-dimensional in this more technical sense not just in general relativistic spacetimes, but also in the other spacetimes I've mentioned; so this more technical sense is a generalization of the one I gave earlier.

To see why this is so in Minkowski spacetime, consider (for ease of visualization) three-dimensional Minkowski spacetime, depicted in figure 4. From a given point in that spacetime the set of points that can be reached by light rays forms a

⁵Here's a brief, intuitive explanation of vector spaces. Think of a (three-dimensional) vector space as the set of all arrows that can be drawn (in ordinary space) from a point. The arrow of zero length counts: it's the zero vector. A subspace of that vector space is a subset of the arrows such that either all of the arrows lie in the same plane (that's a two-dimensional subspace) or all of the arrows lie on the same line (that's a one-dimensional subspace).

Figure 4.4: Minkowski Spacetime



double cone: the future and past light cones at that point. The set of vectors that lie inside either light cone are the timelike vectors. If you pick any line through that cone and consider the set of vectors that point along that line, then (as I explained in footnote 5) those timelike vectors form a one-dimensional subspace of the space of all vectors at that point. But there couldn't be a *two*-dimensional subspace of timelike vectors: if there were, then there would have to be some plane through that point in spacetime such that all the vectors at that point that lie in that plane also fall inside either the past or future light cone. But no plane lies entirely inside the two light cones.

It's not controversial that time is one-dimensional in the familiar spacetime theories. The controversial claim is that it is dimensionality that *makes* timelike vectors timelike. To be explicit, the controversial claim is:

DIMENSION: A timelike vector \vec{v} is one that satisfies the following: the maximum

dimension of a subspace containing only vectors of the same kind as \vec{v} is one.⁶

I objected above to the first way of using geometry to distinguish spatial from temporal directions that it didn't tell us what temporal directions in different spacetimes had in common. DIMENSION does better, because it does say something about what they have in common.

There are, however, several problems for DIMENSION. First of all, it entails that it is necessary that time is one-dimensional. Now, I do not want to deny that time is necessarily one-dimensional. (Some philosophers have tried to describe worlds in which time has two dimensions, but I don't find those descriptions convincing.⁷) But I am not sure I want to affirm it, either. It is better to have a view that does not rule it out.

Second, the dimensionality difference between time and space doesn't seem deep enough to be the fundamental difference between the two. I take it that the fundamental difference between space and time will illuminate the other, less fun-

⁶In Minkowski spacetime the maximum dimension of a subspace containing only null vectors (lightlike vectors and the zero vector) is one. But DIMENSION does not entail that these vectors are timelike. In section 3 I argued that null vectors are not timelike on geometrical grounds alone; DIMENSION, like the other principles I will later advance, is only meant to determine which of the remaining vectors are timelike. (Recall I said in section 3 that talk of two vectors being of the same kind is reserved for vectors that (unlike null vectors) are either spatial or temporal.)

⁷Judith Thompson tries to describe such a world in her (1965). (This example is discussed in more detail in (MacBeath 1993).) It's a world in which two people disagree about the temporal order of certain pairs of events. Each person thinks he outlives the other. The argument that time is two-dimensional in this world goes, in outline, like this. (1) There is no way to temporally order events on a line so that both people are right. (2) But both are equally well-placed to determine the temporal order of events, so we do not want to say that one is right and the other wrong about the temporal order. So (3) the events *cannot* be temporally ordered on a line; time has more than one dimension.

But this conclusion does not follow. Time can be one-dimensional even if there is no uniquely correct way to temporally order events on a line. In Minkowski spacetime, for example, two observers can disagree about the temporal order of events. But time is still one-dimensional in that spacetime.

damental, differences between them, and help us explain those other differences. But the difference in dimensionality doesn't do this.

Third, and most important, it seems possible that both time *and* space be one-dimensional. But in a possible world with two-dimensional Minkowski spacetime, *all* the (non-null) vectors meet the condition in DIMENSION. So DIMENSION entails that all (non-null) directions in that world are temporal, and so (since no direction is both spatial and temporal) that that world contains no spatial directions at all. But that can't be right: surely it's possible that special relativity be true and that time and space both be one-dimensional.

The problem is that while vectors of one kind satisfy the condition in DIMENSION, vectors of the other kind do as well; while we have already established that all timelike vectors are of the same kind. Someone who maintains that dimensionality is the only factor that does any work to distinguish timelike from spacelike directions, then, must admit that the condition in DIMENSION is necessary but not sufficient for a direction to be timelike; and that no condition is sufficient in every world. He might then revise his view as follows:

DIMENSION*: If a vector is timelike then the maximum dimension of a subspace containing only vectors of the same kind as it is one; all timelike vectors are of the same kind; and to the extent that these conditions fail to determine which vectors are timelike, it is to that extent indeterminate which vectors are timelike.

In worlds where just one kind of vector satisfies the condition in DIMENSION, then, DIMENSION and DIMENSION* agree that vectors of that kind are timelike. But in worlds with two-dimensional spacetimes DIMENSION* entails that it is indeterminate which kind of vector is timelike.

In the end, though, this move to indeterminacy fails. It fails not because I insist that it must be perfectly determinate in every world which directions are timelike. But surely in *some* two-dimensional worlds, complicated worlds in which plenty is happening, there is a fact of the matter about which directions are temporal. So we should reject DIMENSION* as well as DIMENSION.

I have looked at two ways to distinguish temporal from spatial directions in geometrical terms, and found reasons for rejecting both. Might some other attempt to distinguish them in geometrical terms succeed where these have failed? The answer is ‘no.’ The two-dimensional spacetimes that make trouble for DIMENSION also make trouble for any attempt to use geometry alone to distinguish temporal from spatial directions. For the roles that timelike and spacelike directions play in the geometry of two-dimensional Minkowski spacetime (and two-dimensional Newtonian spacetime) are symmetric. Since timelike and spacelike directions play symmetric roles in the two-dimensional spacetime geometries, any attempt to distinguish temporal from spatial directions by isolating a geometrical role that one but not the other plays is bound to fail.

So what else other than or in addition to the geometry of spacetime makes the difference between spacelike and timelike directions?

5 Laws of Nature

Timelike and spacelike directions play different roles in the laws of physics that we have taken seriously as the fundamental laws governing our world. Those laws govern the evolution of the world in timelike directions, but not in spacelike directions.

This claim might look analytic (‘of course *evolution* happens in *time*’), but I’m using ‘govern the evolution of the world’ in a stipulated sense that does not build time into its meaning. The laws govern the evolution of the world in some direction just in case the laws, together with information about what is going on in some region of spacetime, yield information about what is going on in regions of spacetime that lie in that direction from the initial region.

My meaning can be made more precise using an example. Earlier I said that the spacetimes of Newtonian mechanics and special relativity, as well as some of the spacetimes of general relativity, can be partitioned into a sequence of times—a sequence of three-dimensional submanifolds. Now in Newtonian gravitational theory, given information about what is going on on one time, the laws determine what is going on on the rest of the times.⁸ These laws govern the evolution of the

⁸I am pretending here (for purposes of illustration) that Newtonian gravitational

world from one time to the others. And a similar claim is true for other laws we've considered fundamental.

Now, timelike vectors are not tangent to any time, on any way of partitioning any given spacetime into times. Rather, no matter which partitioning of spacetime into times you use, timelike vectors point from one time toward others. So timelike vectors point in the directions in which the laws govern the evolution of the world.⁹

The same is not true of points of space. If I know what is going on right here (at this location in space) for all time, the laws tell me nothing about what is going on anywhere else at any time. The laws do not govern the evolution of the world in spacelike directions.

I used the laws of Newtonian gravitational theory as an example, and these laws are deterministic. If some laws are deterministic, then given information about the state of the world on some time, those laws yield complete information about the state of the world on other times. But not all possible laws are deterministic; on some interpretations the laws of quantum mechanics are an example of indeterministic laws. But even here there is a difference between the roles timelike and spacelike directions play in the laws: given information about the state of the world at a time, these laws assign probabilities to possible states at other times; but they do not do so for points of space.

(There is another role that timelike directions play in some familiar laws that spacelike directions do not: quantities like mass, charge, and energy are conserved in timelike directions, but not in spacelike directions. But when there are conservation laws like this, they are usually derived from the dynamical laws (as, for example, the law of conservation of charge follows from Maxwell's equations—the dynamical laws for electromagnetism). So I need not explicitly mention this as a theory is deterministic, even though there are good arguments that it is not. Earman discusses these arguments in his (1986).

⁹As I've said, not every general relativistic spacetime can be partitioned into times. But the laws of general relativity still govern the evolution of worlds with those spacetimes in timelike directions at a local level: some four-dimensional regions of those spacetimes can be partitioned into times, and (if the region is the right shape) the laws determine what is happening at all times given information about what is happening at one time.

second role in the laws that distinguishes timelike from spacelike directions.)

It is not controversial that timelike directions play these roles in the laws with which we are most familiar. I propose that we take these roles as constitutive: *what it is* to be a timelike direction is to play these roles in the laws. To be explicit, the proposal is this:

LAWs: Any direction (that is either spatial or temporal according to the geometry) in which the laws govern the evolution of the world is a timelike direction.

Let me make two remarks about this proposal.


First, I do not claim that it is necessary that the laws of nature govern the evolution of the world in some direction(s) in spacetime.¹⁰ I do not claim, that is, that it is necessary that there be some timelike role in the laws to be filled. Perhaps there are possible worlds with strange laws of nature in which there is no such role. But I *do* claim that in such worlds no direction is a timelike direction.

Second, my proposal presupposes that the laws of nature are more fundamental than the distinction between timelike and spacelike directions. It presupposes, roughly, that it is possible to state the laws of nature without using the words ‘time’ and ‘space.’ For if in order to state the laws we had to presuppose that time and space had already been distinguished, then it would be going in a circle to then appeal to the laws to distinguish time from space. This is not a problem, though. We standardly state laws without appealing to the distinction between time and space. In formal presentations of, say, Newtonian gravitational theory, the distinction between timelike and spacelike directions is always made in some informal remarks after the author has described the geometrical structure of spacetime and before he writes down the equations of the theory. But the equations can be understood perfectly well with the informal remarks removed.

Or, to take a more concrete case, consider Newton’s first law: a body not acted on by any forces moves with a constant velocity. ‘Velocity’ means the same as ‘change in spatial location with respect to change in temporal location.’ I claim that we can re-write this law to remove the reference to space and time. We can do this by using quantifiers: first, replace ‘velocity’ with ‘change in location along the

¹⁰I’m using ‘the laws of nature’ non-rigidly.

Figure 4.5:



x direction with respect to change in location along the y direction'; then, preface the laws with the quantificational phrase, 'there are two (distinct) directions, x and y , such that x and y play such-and-such geometrical roles, and...'

6 Testing Our Intuitions

Let's take a look at a particular world with a two-dimensional spacetime and see if we agree that the laws distinguish space from time even though the geometry of spacetime does not. Consider the spacetime diagram in figure 4.5.

This diagram depicts the distribution of matter in spacetime in some possible world. (Suppose that spacetime has the familiar Newtonian structure, so that talk of space and time makes sense.) Normally we read the vertical axis of spacetime diagrams as the time axis. I ask you to forget about that convention for now and

suppose you do not know which axis represents time. I contend that if you do not know what the laws are, you are unable to tell which axis is time. And that is evidence that it is the laws that are doing the work to make one axis time.

It certainly seems that there are possible worlds correctly represented by this diagram in which the vertical axis represents time; and also possible worlds correctly represented by this diagram in which the horizontal axis represents time. To make this plausible, let me describe for you one world of each kind. Call the world in which the vertical axis represents time ‘the vertical world,’ and the other, ‘the horizontal world.’ First, a description of the vertical world:

VERTICAL: there are two particles that accelerate toward each other, until they meet in an elastic collision and rebound back the way they came; then they slow down, turn around, and accelerate back toward each other, repeating this cycle for all time.

And here’s a description of the horizontal world:

HORIZONTAL: for a long time nothing happens. Then an infinite number of pairs of particles come into existence; one member of each pair moves off to the right, the other to the left. Each particle bounces off a particle coming from the other direction, then is annihilated in a collision with the particle with which it was created. Then nothing happens for the rest of time.

By itself the diagram doesn’t favor one of these descriptions over the other.

Things are different, I think, when I tell you what laws of nature govern the world the diagram depicts. Let the laws be Newton’s three laws of mechanics and a slightly amended version of Newton’s law of universal gravitation.¹¹ Of course Newton’s laws contain terms like ‘velocity’ and ‘acceleration’ which are defined in

¹¹According to Newton’s law the force between two bodies is inversely proportional to the square of the distance between them. As the distance between two bodies goes to zero, the force gets infinitely large. A natural way to extend Newton’s laws to deal with this case is to have the particles bounce off each other in a perfectly elastic collision.

terms of space and time; suppose them to be redefined in the way I mentioned above on page 87, so that velocity ends up being change in position along the horizontal axis with respect to change in position along the vertical axis.¹² So in these laws the vertical axis plays the time role in the definition of ‘velocity,’ and also plays the role that I described above: the laws govern the evolution of the world along the vertical axis. It seems to me that when we add these laws to the description of the world depicted by the diagram, we know enough to know that the vertical axis is the time axis, and we rule out HORIZONTAL as a correct description of that world.

Let me address two worries one might have about this example. First, one might worry that my strategy here—maintaining that knowing only how matter is distributed in spacetime, without knowing what the laws are, leaves us unable to tell which axis is time—is incoherent if a broadly empiricist theory of laws of nature is true. According to such theories, the laws of nature supervene on the occurrent (categorical, non-dispositional) facts. Since the occurrent facts surely include facts about the structure of spacetime and the distribution of matter in it, an empiricist will say that by fixing the facts about matter and spacetime I’ve already fixed what the laws are, and so already fixed which directions are timelike.

This may be so; I won’t take a stand here on whether some empiricist theory of laws is true. But even if an empiricist theory is true, it is coherent to ask you to remain ignorant about the laws even though you know about the distribution of matter in spacetime, for it is not always obvious what the laws are in a world with

¹²So where the original laws contained ‘...velocity...’ the new laws will say ‘there are two (distinct) directions, x and y , such that x and y play such-and-such geometrical roles, and....change in position in the x direction with respect to change in position along the y direction...’

Given that the horizontal and vertical axes play symmetric geometrical roles, you might wonder why it is the direction along the horizontal axis, rather than the direction along the vertical axis, that gets to play the time-role in the definition of velocity. But there is no answer to this question. The way the laws are stated guarantees that (in two-dimensional worlds with Newtonian spacetime) one kind of direction will play the time role in the definition of velocity and the other will not; but it is undetermined which—there are other worlds with the same spacetime geometry and the same laws in which the other plays that role.

a given set of occurrent facts. And the fact that we cannot tell which axis is time when ignorant of the laws but can when we know the laws is still evidence that it is the role they play in the laws that makes timelike directions timelike.¹³

Second, one might worry about the descriptions in VERTICAL and HORIZONTAL. Those descriptions seem to contain more information about the distribution of matter in spacetime than the spacetime diagram alone does. And the points at which the descriptions go beyond the information contained in the diagram are points at which the descriptions disagree. The spacetime diagram tells us only which spacetime points are occupied by material objects. The descriptions contain further information about *how many* material objects there are, and *which* points each object occupies. In VERTICAL I said that there were two particles, while in HORIZONTAL I said there were infinitely many. Now, I claimed that knowing only how matter was distributed in spacetime isn't enough to allow us to figure out which axis is time. But one might complain that by presenting the spacetime diagram I hadn't given complete information about how matter is distributed in spacetime. Complete information requires the kind of information contained in VERTICAL and HORIZONTAL. If I had said that there are only two particles, and each one occupies the points on just one of the curvy lines, then perhaps it would have seemed obvious that the vertical axis is time.

To allay this worry let me make some further stipulations about the world(s) represented by the diagram. I'll add information so VERTICAL and HORIZONTAL won't contain additional information.

Suppose that in worlds represented by the diagram there are uncountably many material objects, that some of them are mereologically simple (without proper parts), and that each simple object occupies exactly one of the occupied spacetime points. Suppose also that for any collection of the simple objects there is a mereologically complex object that they compose. (In slogans, then, I'm saying that

¹³An empiricist theory of laws entails that at most one of VERTICAL and HORIZONTAL describes a possible world. For the descriptions disagree about which axis is time, and so (on my view) disagree about the laws, but agree on the occurrent facts. Empiricists should read the possibilities at work in my argument as epistemic, rather than metaphysical, possibilities.

four-dimensionalism and the doctrine of unrestricted mereological composition are true in these worlds.) VERTICAL and HORIZONTAL seem to disagree about how many things there are and which regions those things occupy, but we can suppose that this is mere appearance. These descriptions are not complete; they don't tell the stories of the spatiotemporal careers of every material object. Instead they only tell the stories of a few salient ones. When we move from VERTICAL to HORIZONTAL we switch which axis we regard as the time axis; doing this brings about a shift in which of material objects are salient.

So far I've presented my proposal and given it some intuitive support. I will now discuss an objection to it.

7 Symmetric Laws and Indeterminacy

The claim that geometry distinguishes timelike from spacelike directions ran into problems with spacetimes in which timelike and spacelike directions play symmetric roles. I said that there is still a difference between timelike and spacelike directions in some of those worlds, and that it is their different roles in the laws which distinguishes them. But what about worlds in which timelike and spacelike directions play symmetric roles in both the geometry *and* the laws? There isn't any reason to deny that such laws are possible; there are even examples of such laws. The wave equation for a wave in one dimension, for example, is

$$\frac{1}{v^2} \frac{\partial^2 \phi}{\partial t^2} = \frac{\partial^2 \phi}{\partial x^2}$$

(ϕ is a function on spacetime; it tells you, intuitively speaking, 'how high' the wave is at each point of spacetime.) Now v here is the speed of the wave, and we're free to choose units in which it's 1. In that case, the equation becomes

$$\frac{\partial^2 \phi}{\partial t^2} = \frac{\partial^2 \phi}{\partial x^2}$$

It is clear that in this law time and space play symmetric roles: switching t and x leaves the equation the same. Moreover, the roles time and space play in these laws

both fit the description I gave above: these laws govern the evolution of the wave both along the time axis and along the space axis.

LAWs entails that all directions in this world are timelike, and that is not right. There are three ways to respond to this problem. First, we could conclude that we do not yet have a complete account of what makes time different from space, and we could search for some other feature of the world that is doing work to distinguish spacelike from timelike directions. Second, we could conclude that it is just a brute fact that the timelike directions in this world are timelike, that there is nothing informative to be said about what makes them timelike. Or third, (paralleling the move from DIMENSION to DIMENSION*) we could conclude that if the geometry and laws fail to single out one kind of vector as the timelike vectors, then it is indeterminate which directions are timelike.

We should choose the third alternative. It does not seem that it could be just a brute fact which directions are timelike; and in these highly symmetric worlds it is hard to see what else, other than the geometry and the laws, could distinguish timelike from spacelike directions. Worlds governed by the wave equation look the same no matter which axis we regard as the time axis.

To deal with these symmetric worlds, then, amend LAWs as follows:

LAWs*: If a direction is timelike then the laws govern the evolution of the world along it; all vectors that point in timelike directions are of the same kind; and to the extent that these conditions fail to determine which vectors are timelike, it is to that extent indeterminate which vectors are timelike.

LAWs and LAWs* agree on all worlds except worlds like the one governed by the one-dimensional wave equation. In such worlds LAWs* entails that it is indeterminate which kind of vector is timelike.

How bad is it to admit that it is sometimes indeterminate which directions are timelike? Certainly it's perfectly determinate in *our* world which directions are timelike. Earlier I claimed that there are some two-dimensional worlds—complicated ones in which there is a lot going on—in which it is perfectly determinate which directions are timelike. But I don't think that this must be perfectly determinate in every two-dimensional world, so I see no problem accepting that it is indeterminate

which directions are timelike in worlds in which time and space play symmetric roles in both the laws and the geometry of spacetime.

8 Laws That Govern in a Spacelike Direction

I turn now to a second objection. My view entails that it is not possible that there be laws that govern the evolution of the world in a spacelike direction. But (so the objection goes), surely this is possible. Surely we can produce examples of possible worlds with this feature.

I have yet to hear a convincing example. Let me go through some of the examples I have heard.¹⁴

Example 1: In this world, it is a law that to the left of every apple is an orange, and to the left of every orange is an apple. (Suppose we've fixed, once and for all, which direction in space is to the left.) This law governs in a spacelike direction: if I know that there is an apple here, I get information about what is going on to the left.

Example 2: In this world, there is a special plane dividing space in half; and it is a law that the contents of space on one side of the plane are perfectly mirrored on the other side. This law governs in a spacelike direction: if I know that there is a red sphere in a certain place on one side of the plane, I get information about what is going on at the very same time in the corresponding place on the other side of the plane.

I am enough of a metaphysician to take examples like these somewhat seriously. I have a three-part response. First, I don't think the laws in these examples are doing enough to govern the evolution of the world in a spacelike direction. Given information about what is going on at one place for all time, they do not give information about what is going on everywhere else at all times. These laws only give

¹⁴I heard many of these examples at the APA Eastern Division meeting, 2004. Among those who suggested examples are Eric Lormand, James Van Cleve, Philip Bricker, and Jonathan Schaffer. There were others, as well, and I can no longer remember who suggested which examples.

information about what is going on in *some* other places. (The law in example 1 gives information about places to the right of, and some of the places to the left of, the initial place. The law in example 2 gives information about just one other place for all time: the place that is the mirror-reflection of the initial place.) So they're not doing in a spacelike direction what laws like Newtonian mechanics do in a timelike direction.

Second, even if there is a world in which it is true that to the left of every apple is an orange (and so on), and a world with mirror symmetry, I'm not sure why I should believe that it would be a *law* that to the left of every apple is an orange, or that the world exhibits mirror symmetry. (Certainly the law about apples couldn't be a fundamental law, but maybe it could be derivative.) For I find it hard to have intuitions about what the laws of some world are, given descriptions of the goings-on in those worlds.

Of course, some philosophers can argue that if there is a world in which to the left of every apple is an orange (and so on), and the world is simple enough in other ways, then it is a law that to the left of every apple is an orange. These are philosophers who accept the best system theory of laws. According to the best system theory of laws, those true statements are laws that are theorems of the deductive system that best balances simplicity and strength.¹⁵ If the apple world is simple enough in other ways, then the statement that to the left of every apple is an orange might make it in to the best system, and so might be a law.

Earlier I tried to remain neutral between primitivist theories of laws and empiricist theories of laws (like the best system theory). Now it appears that I must take sides. I reject the argument in the previous paragraph because I reject the best system theory of laws.

This is not just an *ad hoc* move, though. I do not just reject the theory because it conflicts with my theory of the difference between space and time. My primary reason for rejecting it is that it fails to respect our modal intuitions about lawhood. Briefly, empiricism about laws (and so the best system theory in particular) entails that there cannot be worlds that differ merely in what laws govern them. But I accept

¹⁵(Lewis 1986b).

the counterexamples opponents of the best systems theory offer to this claim.¹⁶ For example, it seems that there is a possible world in which an empty spacetime has a Minkowski geometry, and in which special relativity is true; and another world in which an empty spacetime has a Minkowski geometry, and in which general relativity is true. (In the latter world, had there been any matter, spacetime would have been curved; this is not true in the former world.) Empiricists about laws must say that at least one of these worlds is not possible.

There is a third part of my response to examples like 1 and 2. It is part of the description of the possible worlds in these examples that certain directions are spacelike and certain others, timelike. But doesn't this beg the question? Why should we accept that the directions called 'spacelike' in the description really are spacelike?

One might refuse to answer this question. But then I am not sure why I should take seriously the claim that there are possible worlds like those described in the examples. Or, one might answer this question by proposing an alternative theory of what makes time different from space. I'll criticize a few such theories in the next section. Or, one might say: we can imagine the worlds described, and it is part of the content of our imagining that the directions called 'spacelike' are spacelike. And since imagination is a guide to possibility, that gives us reason to believe that these worlds are possible.

How does this act of imagination work? Maybe like this: we imagine watching the history of the world in question unfolding, as if we were watching a movie. We can tell which directions are timelike and which spacelike when watching a movie (without needing any theory of the difference to help us); in the same way, we can tell which directions are timelike and which spacelike in these worlds.

But imagining watching the history of the world unfold as if watching a movie is only a legitimate way to learn which directions in that world are timelike if those are the experiences an observer who *existed in the world* would have. But if the worlds described in these examples contain conscious observers, then there must be many more laws than the ones mentioned. Once all these laws are also taken into account, I doubt it will be so clear that the laws govern the evolution of the

¹⁶For example, in (Carroll 1994).

world in a spacelike, not a timelike, direction.

Example 3: In this world, the laws are those of Newtonian mechanics and Newton's law of universal gravitation. This law governs in a space-like direction: if I know that there is a particle here of a certain mass now, I know something about the value the gravitational field now has at every other point.

This example is not convincing. First, it is not clear that on a correct interpretation, Newtonian gravitational theory says that there really is any such thing as the gravitational field. And even if it does, knowing that there is now a particle here of a certain mass doesn't tell me the value of the field at any other point; to know that I'd have to know how many particles there were in total, and what their masses and positions are right now. If I know anything about what is going on elsewhere, it is only how the particle here contributes to the value of the gravitational field elsewhere; but since this is consistent with the field's actual value being anything at all, and since we get no information of the kind we're really looking for (namely, information about which other regions of space are occupied by material objects), these laws are not really governing the world in a spacelike direction.

Example 4: In this world, the laws of quantum mechanics govern the world. In an EPR-type experiment, there are two particles some distance apart, and if we measure the spin on one of them in some direction, we know with certainty the outcome of a measurement of the spin of the other particle in that same direction, even if the measurement events are spacelike-separated. So these laws govern the evolution of the world in a spacelike direction.

This example is not convincing, for two reasons. First, I want to complain that, as in examples 1 and 2, the determination by laws in a spacelike direction here is not the robust kind needed to refute my view. But there is a second, more important, problem. For laws to govern the evolution of the world in a spacelike direction, it must at least be the case that given information about what is going on in one place (say, right here), the laws give information about what is going on at other

places that are spacelike separated from it. But if all we know is that (after being measured) some particle right here has spin-up in some direction, the laws don't tell us anything about what is going on elsewhere. They only give us information if we also know that there is another particle somewhere else, and that the 'system' comprising the two particles is in an entangled state. But this information is not just information about what is going on right here.

9 Alternative Views

Finally, I will review some alternative views about what makes time different from space, and say why I reject them. My goal is not to give these theories detailed formulations and refutations; I aim only to suggest why I think they go in the wrong direction from the start.

9.1 The Causal Theory of Time

According to this view, a direction is timelike just in case it is a possible direction of causation. This theory was inspired, I think, by a certain way of thinking about Minkowski spacetime. There is a synthetic axiomatization of this spacetime's geometry using just one two-place predicate that can be taken to mean ' x and y are causally connectible.' Perhaps this axiomatization is getting the metaphysics right: the spatiotemporal relations in Minkowski spacetime, and so facts about which events occur before which other events, are derived from a more fundamental relation of causal connectibility.¹⁷

I reject this theory for two reasons. The first, and less important reason, is that it precludes the possibility of instantaneous causation. It looks like Newtonian mechanics involves instantaneous causation—according to that theory the sun's being a certain distance from the earth instantaneously causes the earth to experience a certain force—and I accept that Newtonian worlds are metaphysically possible. But I place more importance on a second reason. I just don't think that facts about causation are more fundamental than the difference between space and time. But they have to be, for this theory to be right.

¹⁷See (Sklar 1974).

(My theory, of course, requires that the laws be more fundamental than the difference between space and time. One might wonder why I am comfortable with this and uncomfortable with the causal theory of time. I won't give a detailed answer to this question; I will just point out that many contemporary philosophers share the feeling that causation is less fundamental than lawhood.¹⁸)

9.2 Three-Dimensionalism

Three-Dimensionalism is the view that material objects persist through time without having temporal parts. Since it is commonly admitted that material objects are extended in space by having spatial parts, there is an asymmetry here between space and time. One might try to distinguish space from time using this asymmetry: time is that dimension in which material objects are extended without being made up of parts.

Three-Dimensionalism is controversial, so insofar as I am not a three-dimensionalist I am not tempted by this proposal. But I do not think three-dimensionalists should be either. One standard argument that material objects that persist through time must have temporal parts is the argument from temporary intrinsics: if something is spherical at a time, then it must have a temporal part that is spherical *simpliciter*, on pain of making sphericity a relation to times, and not an intrinsic property at all.¹⁹ Three-dimensionalists think they can resist this argument. But if they can, then they can also resist the parallel argument that material objects that are spatially extended must have spatial parts, the argument from local intrinsics: if something is red in one place and green in another, then it must have a spatial part that is red *simpliciter*, on pain of making redness a relation to places. (Maybe redness and greenness are not the best examples, but there are others.) So three-dimensionalists ought to admit the possibility of material objects that are spatially extended without

¹⁸David Lewis (1986a) is one example: he analyzes causation in terms of counterfactuals, and his truth-conditions for counterfactuals appeal to laws. But even philosophers who reject counterfactual analyses of causation, like Maudlin (2004), believe that laws are more fundamental than causation.

¹⁹This argument is much discussed; it is presented, among other places, in (Lewis 1986b) and (Sider 2001).

having any spatial parts. But once that possibility is granted, there is no longer an asymmetry between the way material objects are (or can be) extended in space and the way they are (or can be) extended in time.

9.3 The Brute Fact View

According to the brute fact view, there is no need to appeal to geometry or the laws to distinguish spacelike from timelike directions. Instead, there is no way to distinguish spacetime from timelike directions in other terms. There is no informative answer to the question, ‘what makes timelike directions timelike?’

One way to put this view is to say that there is a fundamental property, the property of being a timelike direction. But it might seem odd to believe that things like directions could have fundamental properties. If we confine our attention to Newtonian spacetime we can avoid this oddness by avoiding talk of directions. Some regions of this spacetime are times, and others are not. The ones that are times have some fundamental intrinsic property that the others do not: the property of being a time. End of theory.

This theory is not plausible. Certainly the regions that have this special property must also play the appropriate role in the geometry. (Supposing we characterize Newtonian spacetime using two distance functions, the role is as follows: each region contains all and only the points that are zero distance from any point in that region, according to one of those distance functions.) But why is there this necessary connection between this special property and a certain geometrical role? The brute fact view gives no answer. My view does better: since it does not postulate the special property, it has no necessary connection to explain.

But maybe that is just a caricature of the brute fact view. Here is a closely related view one might have. Do not postulate a special non-geometrical fundamental property of being a time. Instead, pick out one of the geometrical relations that gives spacetime its structure, and make it special. Sticking to my focus on Newtonian spacetime, one way to put the view is as follows: this spacetime get its structure, let us suppose, from two distance functions. One of these is the *temporal*

distance function, and the other is the *spatial* distance function.²⁰ Times are regions containing all and only the points that are zero distance from any point in them, according to the temporal distance function. What makes one function the temporal distance function, rather than the spatial distance function? There is no answer. It is just a brute fact.

This view I take more seriously as a competitor to my own than the others I have considered. But I do think it is wrong. The reason appeals, again, to two-dimensional Newtonian spacetime.

Take a world w with two-dimensional Newtonian spacetime, and ‘rotate’ all the matter in that world ‘90 degrees’²¹ to produce a new world w^* . I can describe this world in a bit more detail, but I don’t want to beg any questions by calling certain regions of spacetime in w^* ‘times’ or ‘points of space.’ So I will have to pick out regions of spacetime in w^* using features those regions have in w . The description ‘regions of spacetime that are points of space in w ’ picks out a definite set of regions of spacetime in w^* , while leaving it open whether those regions are also points of space in w^* . In more detail, then, w^* looks like this: events that are simultaneous in w occur (in w^*) in regions that (in w) lie on different times but are in the same place. The laws of w^* are also obtained from the laws of w by ‘rotation’: the w^* -laws treat regions that are points of space in w as the w -laws treat the regions that are instants of time in w . I think w^* is qualitatively indiscernible from w . But the brute fact view cannot say this. According to the brute fact view, two-dimensional Newtonian spacetime is not symmetric in this way. So the brute fact view entails that either the ‘rotation’ operations cannot be performed (that is, there is no such ‘rotated’ world), or, if they can, they result in a world that is very different, qualitatively, from w . So I reject this view.

Both versions of the brute fact view are similar to the view that what makes the future different from the past is that the future direction in time has some special

²⁰There are analogous ways to formulate the view on other accounts of the fundamental spatiotemporal relations that characterize Newtonian spacetime.

²¹Of course, this doesn’t really make sense in the geometry of Newtonian spacetime. What I really mean is: consider this world represented on a two-dimensional Euclidean plane, like a piece of paper; then rotate everything 90 degrees on this Euclidean plane; now consider the world this new diagram represents.

intrinsic property that the past direction in time lacks.²² I reject the brute fact view for the same reason many reject this view about the difference between the past and the future.²³ In the later case, it seems that a world in which the distribution of matter were ‘mirror reversed’ around a given time (the laws, being time-reverse invariant (we may suppose), would be the same) would not be a world in which everything were ‘going backwards,’ but a world in which the direction that is actually the future direction is the past direction. Many who hold this view identify the future direction with the direction in which entropy increases;²⁴ so it is not intrinsic to the future direction that it *be* the future direction. I accept this view about the difference between the future and the past. And my view about the difference between space and time is analogous.

²²(Maudlin 2002) seems to defend this view, as does (Earman 1974).

²³(Price 1996) is an example.

²⁴(Reichenbach 1956) is one example.

Bibliography

- Adams, Robert (1979). 'Primitive Thisness and Primitive identity.' *Journal of Philosophy* 76(1): 5–26.
- Barbour, Julian (1999). *The End of Time*. New York: Oxford University Press.
- Belot, Gordon (1995). 'New Work for Counterpart Theorists: Determinism.' *British Journal for the Philosophy of Science* 46: 185–195.
- (2000). 'Geometry and Motion.' *British Journal for the Philosophy of Science* 51: 561–595.
- Bennett, Jonathan (1984). *A Study of Spinoza's Ethics*. Hackett.
- (2001). *Learning from Six Philosophers*, volume 1. New York: Oxford University Press.
- Blackburn, Simon (1993). 'Filling In Space.' In *Essays in Quasi-Realism*, chapter 14. New York: Oxford University Press.
- Brighouse, Carolyn (1994). 'Spacetime and Holes.' *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1: 117–125.
- (1997). 'Determinism and Modality.' *British Journal for the Philosophy of Science* 48: 465–481.
- Butterfield, Jeremy (1989). 'The Hole Truth.' *The British Journal for the Philosophy of Science* 40: 1–28.

- (2002). ‘Critical Notice: Julian Barbour, *The End of Time*.’ *British Journal for the Philosophy of Science* 53: 289–330.
- Carroll, John W. (1994). *Laws of Nature*. New York: Cambridge University Press.
- Clifford, William Kingdon (1882). ‘On the Space-Theory of matter.’ In Robert Tucker (ed.), *Mathematical Papers*, 21–22. London: Macmillan.
- Descartes, Rene (1988). *Selected Philosophical Writings*. New York: Cambridge University Press.
- Earman, John (1974). ‘An Attempt to Add a Little Direction to ‘The Problem of the Direction of Time’.’ *Philosophy of Science* 41(1): 15–47.
- (1986). *A Primer on Determinism*. Boston: D. Reidel Publishing Company.
- (1989). *World Enough and Space-Time*. Cambridge, MA: MIT Press.
- (1995). *Bangs, Crunches, Whimpers, and Shrieks*. New York: Oxford University Press.
- Earman, John and John Norton (1987). ‘What Price Substantivalism? The Hole Story.’ *The British Journal for the Philosophy of Science* 38: 515–525.
- Faraday, Michael (1965). ‘A speculation touching Electric Conduction and the Nature of Matter.’ In *Experimental Researches in Electricity*, volume 2, 284–293. New York: Dover.
- Field, Hartry (1980). *Science Without Numbers*. Princeton: Princeton University Press.
- (1989). ‘Can We Dispense With Space-time?’ In *Realism, Mathematics and Modality*, chapter 6. New York: Basil Blackwell.
- Forbes, Graeme (1987). ‘Places as Possibilities of Location.’ *Nous* 21: 295–318.
- Friedman, Michael (1983). *Foundations of Space-Time Theories*. Princeton: Princeton University Press.

- Harman, Gilbert (1986). *Change in View*. Cambridge, MA: M.I.T. Press.
- Hofer, Carl (1996). 'The Metaphysics of Space-Time Substantivalism.' *Journal of Philosophy* 93(1): 5–27.
- Kripke, Saul (1980). *Naming Necessity*. Cambridge, MA: Harvard University Press.
- Leibniz, G.W. and Samuel Clarke (2000). *Correspondence*. Ed. Roger Ariew. Indianapolis: Hackett.
- Lewis, David (1968). 'Counterpart Theory and Quantified Modal Logic.' *Journal of Philosophy* 65: 113–126.
- (1983). 'New Work for a Theory of Universals.' *Australasian Journal of Philosophy* 61: 343–377.
- (1986a). 'Causation.' In *Philosophical Papers*, volume 2, 159–171. New York: Oxford University Press.
- (1986b). *On The Plurality of Worlds*. New York: Blackwell.
- MacBeath, Murray (1993). 'Time's Square.' In Robin Le Poidevin and Murray MacBeath (eds.), *The Philosophy of Time*, 182–202. New York: Oxford University Press.
- Maddy, Penelope (1990). *Realism in Mathematics*. New York: Oxford University Press.
- Markosian, Ned (2000). 'What are Physical Objects?' *Philosophy and Phenomenological Research* 61: 375–395.
- Maudlin, Tim (1990). 'Substances and Space-Time: What Aristotle Would Have Said to Einstein.' *Studies in History and Philosophy of Science* 21: 531–561.
- (1993). 'Buckets of Water and Waves of Space: Why Spacetime Probably is a Substance.' *Philosophy of Science* 60: 183–203.
- (2002). 'Remarks on the Passing of Time.' *Proceedings of the Aristotelian Society* 102: 259–274.

- (2004). ‘Causation, Counterfactuals, and the Third Factor.’ In John Collins, Ned Hall and L. A. Paul (eds.), *Causation and Counterfactuals*, chapter 18, 419–444. Cambridge, MA: MIT Press.
- McCormmach, Russell (1970). ‘H. A. Lorentz and the Electromagnetic View of Nature.’ *Isis* 61(4): 459–497.
- Melia, Joseph (1998). ‘Field’s Programme: Some Interference.’ *Analysis* 58(2): 63–71.
- (1999). ‘Holes, Haecceitism and Two Conceptions of Determinism.’ *British Journal for the Philosophy of Science* 50: 639–664.
- Newton, Isaac (1962). ‘De Gravitatione.’ In Marie Hall and Alfred Hall (eds.), *Unpublished Scientific Papers of Isaac Newton*, 121–156. Cambridge: Cambridge University Press.
- O’Leary-Hawthorne, John and J. A. Cover (1996). ‘Haecceitism and Anti-Haecceitism in Leibniz’s Philosophy.’ *Nous* 30(1): 1–30.
- Pooley, Oliver and Harvey R. Brown (2002). ‘Relationalism Rehabilitated? I. Classical Mechanics.’ *British Journal for the Philosophy of Science* 53: 183–204.
- Price, Huw (1996). *Time’s Arrow and Archimedes’ Point*. New York: Oxford University Press.
- Quine, W. V. (1981). ‘Things and Their Place in Theories.’ In *Theories and Things*, chapter 1. Cambridge, MA: Harvard University Press.
- Reichenbach, Hans (1956). *The Direction of Time*. Los Angeles: University of California Press.
- Royden, H. L. (1959). ‘Remarks on Primitive Notions for Elementary Euclidean and non-Euclidean Plane Geometry.’ In Leon Henkin (ed.), *The Axiomatic Method with Special Reference to Geometry and Physics*, 86–96. Amsterdam: North-Holland.

- Sider, Theodore (2001). *Four-Dimensionalism*. New York: Oxford University Press.
- Sklar, Lawrence (1974). *Space, Time, and Spacetime*. Berkeley: University of California Press.
- Stein, Howard (1970). 'On The Notion of Field in Newton, Maxwell, and Beyond.' In Roger H. Stuewer (ed.), *Minnesota Studies in the Philosophy of Science*, volume 5, 264–287. Minneapolis: University of Minnesota Press.
- (1977). 'Some Philosophical Pre-History of General Relativity.' In J. Earman, C. Glymour and J. Stachel (eds.), *Minnesota Studies in the Philosophy of Science*, volume 8. Minneapolis: University of Minnesota Press.
- Tarski, Alfred (1959). 'What is Elementary Geometry?' In Leon Henkin (ed.), *The Axiomatic Method with Special Reference to Geometry and Physics*, 16–29. Amsterdam: North-Holland.
- Thomson, Judith Jarvis (1965). 'Time, Space, and Objects.' *Mind* 74(293): 1–27.
- Wheeler, John A. and Edwin F. Taylor (1963). *Spacetime Physics*. 1st edition. San Francisco: W. H. Freeman and Company.
- Williams, Donald C. (1951). 'The Myth of Passage.' *The Journal of Philosophy* 48(15): 457–472.
- Wilson, Mark (1993). 'There's a Hole in the Bucket, Dear Leibniz.' *Midwest Studies in Philosophy* 18: 202–241.