

Fifty Million Elvis Fans Can't be Wrong

GORDON BELOT
University of Michigan

Abstract

This essay revisits some classic problems in the philosophy of space and time concerning the counting of possibilities. I argue that we should think that two Newtonian worlds can differ only as to when or where things happen and that general relativistic worlds can differ in something like the same way—the first of these theses being quaintly heterodox, the second baldly heretical, according to the mores of contemporary philosophy of physics.

1. Introduction

It's a great title (for an album).¹ And: it is true. About some things—such as the fact that Elvis was peachy-keen—his legions of fans could not be wrong, just because there were so many of them.² Similarly, might makes right whenever a large group of people uses a word or a concept in a certain way—if you are interested in providing an account of how that word or concept works, then you had better be able to accommodate their use (not necessarily *only* their use, of course).

With this in mind, I take another look here at some classic questions about the counting of possibilities from the philosophy of space and time, paying special attention to how physicists count possibilities. Section 2 is concerned with preliminary groundwork. Section 3 argues that philosophers should follow physicists in accepting that Newtonian worlds can differ only as to where or when things happen. Section 4 makes a similar case concerning general relativistic worlds. Section 5 floats some speculative suggestions concerning wider questions, touching on the metaphysics of laws of nature and the most perspicuous way of thinking of physical theories.

First, a warning. In what follows, conclusions are often somewhat overstated and the tone is often too preachy. An irenic formulation of the point behind the bluster: among the starting points of philosophical work on the problems discussed here should be the project of trying to make sense of how physicists seem to think about these problems—but because this project has not received the attention that

Earlier versions of this paper were presented in 2014 and 2015 at the Université Paris Diderot, UC Irvine, the University of Rochester, the London School of Economics, a meeting of the Philosophy of Science Association, the White Stallion Ranch, and MIT. For helpful comments and discussion, thanks especially to Dave Baker, Caro Brighthouse, Jeremy Butterfield, Shamik Dasgupta, Cian Dorr, Ned Hall, Kathrin Koslicki, Alyssa Ney, Laura Ruetsche, Brad Skow, David Wallace, and Jim Weatherall.

it deserves from philosophers, there is something unbalanced about the current literature.

Next, another warning. My topics lie at or near the intersection of metaphysics and philosophy of physics. I have aimed to write so that the main thread of argument will be accessible to readers from either audience. On occasion I enter into somewhat technical discussions, liable to be regarded as essential by one constituency and as eminently skippable by the other. These are for the most part relegated to footnotes or to the subsections headed 'Worries.' Very likely the result will not be entirely to anyone's taste.

2. Preliminaries

At the level of the mathematical objects that we use to represent possibilities, isomorphism is cheap: except in special cases, if we have a mathematical object with a given structure, there will be others in which the same cast of characters instantiates the same structure, but with the roles permuted. If I take an equilateral triangle and paint one vertex green, you can make a different but isomorphic structure by erasing my work and painting one of the other vertices. If we are given three points and told to arrange them to form a scalene triangle with sides of lengths of three, four, and five, there are six structures we could end up building (differing as to which points lie opposite the longest and shortest sides).³

A question naturally arises: Is this multiplicity of isomorphic structures reproduced at the level of possibilities? Isomorphisms generate new representations—but do they also generate new possibilities? Two special forms of this question have played a large role in the philosophy of space and time over the years.

Shifts in Newtonian Spacetime Picture two-dimensional Newtonian spacetime as an infinitely extended sheet of graph paper. Horizontal lines correspond to instants of time, vertical lines to points of absolute space. A possible history of a system of n particles moving with finite velocity is represented by n curves, each never tangent to a horizontal line. Given a pattern of curves of this kind, one can always generate another by rigidly translating the given pattern horizontally and vertically: lay a sheet of glass over the graph paper; trace the given curves onto the glass, then erase them from the graph paper; slide the glass horizontally, then vertically along the paper; then trace the curves back onto the paper. Two patterns of curves differing from one another in this way are isomorphic and, intuitively, represent the histories in which the same patterns of matter in motion obtain (worldlines approach or recede from one another in certain ways)—but differ as to “where” and “when” events occur (e.g., as to at which point of the underlying mathematical spacetime two worldlines intersect). At the mathematical level, it is clear that spatial and temporal shifts generate new representations. But do shifts also generate new possibilities? Are there possibilities that differ only as to when and where things happen?⁴

Generalized Shifts in General Relativity The preceding case was analogous to the problem of taking an equilateral triangle and privileging one vertex—we were given

a geometrical structure (Newtonian spacetime) and looked for a way to paint a pattern of a certain kind on it (determined by some given pattern of curves representing particles). In general relativity, spatiotemporal geometry varies from one solution of the theory's equations to another, rather than being fixed. But if we consider one of the many possible geometries that arise in this way, we face a problem like that of constructing a scalene triangle of given dimensions out of three given points—there will be many ways of instantiating the given geometry using the materials at hand. To get a feeling for how this works, consider the following toy problem. We are given a string that can be stretched and compressed.⁵ And we are asked to break its length up into one thousand parts, which we will then consider to be of equal size. One way to proceed is to stretch and or compress the string until it is matched up against a metre stick, then to transfer the millimetre markings from the stick to the string. But there is in general no canonical way to do this: if we restretch and recompress our string, then repeat the procedure, the marks we end up with again give a good segmentation of the string; but, in general, we expect to arrive at a different segmentation whenever we repeat the procedure.⁶ Similarly, in general relativity, it is easy at the level of representation to find structures that instantiate the same geometry, but differ as to which point of their shared mathematical space plays which role in this geometry—we will say that *generalized shifts* generate new representations in general relativity.⁷ Does this also happen at the level of possibilities—are there possible general relativistic histories that differ only as to when and where things happen? Do generalized shifts generate new possibilities?

The philosophical literature on these questions is dominated by two positions.

THE OLD WAY. Isomorphisms sometimes generate new possibilities—shifts always do, generalized shifts never do.

THE NEW WAY. Isomorphisms never generate new possibilities—so, in particular, shifts and generalized shifts never do so.

My goal here is to argue that neither of these gets it right—even the old way is too stinting. In the first section below, I will first argue that we should take shifts to generate new possibilities in theories set in Newtonian spacetime. In the following section, that we should take *some* generalized shifts to generate new possibilities in general relativity. Throughout I will concentrate on temporal (generalized) shifts—but there is no fundamental difference between the spatial and temporal cases.

3. The Shifty and the Shiftless

Among the very few things that Leibniz and Clarke agreed about in their correspondence was that shifts generate new possibilities on Newton's absolutist view of space and time (on which the instants of time and parts of space are individuals as real as any material object) but not on Leibniz's relationalist view of space and time (on which space and time and their parts are not themselves objects on a par

with material objects but are rather a sort of abstract framework for talking about actual and possible relations between bodies).⁸

These days, a growing number of philosophers think that Leibniz and Clarke were wrong about this—they hold that no matter what one's account of the ontology of space, time, or spacetime one ought to count possibilities in the relationalist manner, and deny that two possible histories can differ only as to what happens when or where.⁹

I maintain that this new shiftlessness is wrong-headed: if your ontology includes regions of spacetime, then you should cleave to the old shifty way and admit that shifts generate new possibilities in the context of theories set in Newtonian spacetime.¹⁰ Henceforth, in discussing shifts I restrict attention to views that include regions of spacetime in their fundamental ontology, since this is where the action is. I will give two arguments. The first, in §3.1, is more or less internal to the usual dialectic on this issue. The second, in §3.3, is of a more methodological cast and turns on claims about how physicists count degrees of freedom. (This discussion will focus on temporal shifts—but a similar story could be told about spatial shifts or rotations.)

3.1 *The swerve*

Here is a *very* simple physical theory that one might have:

Spacetime is two-dimensional and Newtonian. There is a single particle. It is at rest up until some moment at which it swerves into motion, after which it moves uniformly to the right with unit speed.¹¹

Prima facie, this theory is indeterministic. It tells us *that* a swerve will occur but does not say *when* it will occur, even if we supplement the laws with initial conditions (i.e., the initial location of the particle). And, *prima facie*, this theory is one in which time is homogeneous—neither the spacetime geometry nor the laws of the theory draw distinctions between the various instants of time.

The shifty—those who hold that shifts generate new possibilities—can vindicate these intuitive judgments. Let w_0 be a world falling under our theory. And for each real number t , let w_t be the world that results from temporally translating the particle worldline in w_0 by t units.

- (i) At any time before $t = 0$, the histories of w_0 and w_1 match perfectly (the particle is at rest at the same point of space in both worlds). But the particle swerves at $t = 0$ in w_0 but not in w_1 . So fixing the laws and an initial segment of history isn't sufficient to fix the global history. The theory is indeterministic.¹²
- (ii) Since the swerve theory is set in Newtonian spacetime, the question of time's homogeneity reduces to the question whether the laws draw any distinctions between instants. In w_0 , the swerve occurs at time $t = 0$. But for any other instant of time $t^* \neq 0$, there is a world w_{t^*} just like w_0 except that the swerve occurs at t^* . Both w_{t^*} and w_0 fall under the same laws. So time is homogenous in this theory: for any time and any event permitted by the laws, there is a world where that event happens at that time.

But if you are shiftless—if, that is, you deny that shifts generate new possibilities in this sort of context—then you must deny that the theory is indeterministic and deny that the theory represents time as being homogeneous.

- (a) Those who deny that shifts generate new possibilities will hold that there is only a single world w at which the swerve theory is true. They will then have to regard the theory as being deterministic: if you have only one world to work with, you will certainly not be able to find two worlds whose initial histories match but whose global histories diverge.¹³
- (b) Consider the instant t_0 at which the swerve occurs in a swerve-world w . If you maintain that w is the only world at which the swerve theory is true, then you maintain that it is nomologically necessary that the swerve occurs at t_0 . So you deny that time is homogeneous in the swerve theory, since you think that the laws draw a distinction between t_0 and the other instants of time.¹⁴

Intuitively, our swerve theory is at least a little bit different from a theory that says: spacetime is Newtonian except that one instant is singled out as being the special instant at which the swerve occurs. But if you deny that shifts generate new possibilities, these theories collapse into one. How embarrassing! Better, then, to be shifty than shiftless.

3.2 Worries

But wait! What about cheap haecceitism? A not too tendentious characterization of the main point of the above discussion is that if one denies that shifts generate new possibilities, then one is driven to say that the swerve *has* to occur whenever it *does* in fact occur. But that is too quick, surely! After all, on Lewis's account (= qualitative counterpart theory with a reliance on sameworldly counterparts) we can have it both ways.¹⁵ Lewis originally characterized his account as a form of cheap quasi-haecceitism—but I will follow what is now standard usage and call it simply *cheap haecceitism*.¹⁶

For the cheap haecceitist, shifts don't generate new possibilities: no two possible worlds differ only as to how they distribute roles over individuals—so, *a fortiori*, no two worlds differ only as to when and where things happen.¹⁷ On the other hand, there is a sense in which, relative to a swerve world w , the swerve could have occurred at times other than the instant t at which it did occur (in virtue of the fact that t has counterparts at w other than itself). So it may appear that the swerve discussion above has little force in the end—one can be shiftless and yet uphold the indeterminism of the swerve theory.

Well—there may be some way to weasel out of the idea that shiftlessness commits one to thinking that the swerve theory is deterministic. But I don't think that cheap haecceitism does the trick.

- (i) On the one hand, while cheap haecceitists certainly are entitled to say that there is a sense in which the swerve at world w could have occurred at some time other than the instant t at which it in fact occurs, this sense is frustratingly

thin. After all, this sense is underwritten by counterpart relations that count arbitrary instants at w as counterparts of t . But, of course, arbitrary instants at w are not qualitative duplicates of t , since t is the only instant at which a swerve occurs.¹⁸ Here we have “an extraordinarily generous counterpart relation, one which demands nothing more of counterparts than that they be things of the same kind.”¹⁹ On a standard qualitative counterpart theory, one expects that such counterpart relations will be ruled out in tolerably strict conversational contexts.²⁰ So according to factory-model cheap haecceitism, it is in many (most?) contexts *false* to say that the swerve could have occurred at some time other than t —and hence often (usually?) false to say that the swerve theory is indeterministic.²¹

- (ii) On other hand, there is something very strange about the package the cheap haecceitist presents us with. The swerve theory is indeterministic—i.e., there are possible histories that share an initial history but then diverge. But shifts don't generate new possibilities—so there is in a sense only one swerve world. A lot of weight is being put here on a distinction between possible histories (=sameworldly counterparts of a possible world considered as a mereological sum) and possible worlds (up to duplication). Here I think it is natural to follow a line taken by Kment: (a) in philosophical accounts of modality, we should regard ‘possible world’ as a technical term and there should be no presumption that it corresponds directly to informal notions, such as ways the world could be; (b) indeed, given the work that the two notions do in this theory, it seems like the world-plus-counterpart relation is in fact a better correlate of our informal possible-worlds notions than is the cheap haecceitist's official category of possible worlds.²² But then it is natural to say that this is an account on which shifts *do* generate new possibilities, in precisely the same sense in which it is an account on which the swerve theory is indeterministic.

But wait! I have good philosophical reasons to be a cheap haecceitist—doesn't that mean that I am committed to shiftlessness in tolerably strict contexts? In a word, no.

Let it be settled that we are going to be qualitative counterpart theorists who make free use of sameworldly counterparts. The next question facing us is what sort of counterpart relations we are going to countenance—different choices will correspond to different coherent families of modal concepts.

For consider how things look if we simply decide to adopt counterpart theory. We can agree that no particular exists at more than one world, and that the truth conditions for *de re* modal claims about a given object concern objects related to it by admissible counterpart relations. But this much agreement leaves us a great deal of latitude in the modal concepts we can capture.

- (a) It would, for instance, remain open to us to be neo-Spinozists: we could maintain that identity is the only admissible counterpart relation—and thereby collapse the distinction between the contingently true and the necessarily true.

Interesting—but no one would ever consider this package to be adequate to modelling our modal concepts.

- (b) Or we could, as Lewis originally did, require that under admissible counterpart relations no thing have any counterpart other than itself at its own world.²³ It would then follow that no matter how loosely we were speaking we could never truly say that Lewis could have been poor Fred there, or that a person living in one epoch of a world of eternal recurrence could have lived in another.²⁴ Very likely there is a coherent notion of possibility relative to which that is the right verdict. But it is not ours. Lewis drew the moral that we ought to allow same-worldly counterpart relations other than the identity relation if we want to be able to model our actual modal concepts.²⁵

By varying the constraints that they require admissible counterpart relations to satisfy, qualitative counterpart theorists too can reconstruct a variety of modal concepts within their framework. So far in considering qualitative counterpart theories, we have always considered the standard version, according to which one can say truly that the swerve could have occurred at a different time than that at which it actually did occur only in contexts that license relatively generous counterpart relations. To the extent that it looks like our modal concepts are such that even in strict contexts we can truly say that the swerve could have occurred other than when it did, that is a strike against the standard approach (in the same way that the prevalence of haecceitistic intuitions was a strike against Lewis's original approach, with its ban on same-worldly counterparts).

But it is easy to cook up an alternative to the standard approach that lacks this problem: stipulate that if x is a part of spacetime, then all (same- or otherworldly) counterparts of x are likewise parts of spacetime and further that the only sort of qualitative similarity that counts towards counterparthood for parts of spacetime is similarity in geometric respects. One can think of these restrictions on counterparthood for parts of spacetime as embodying the traditional substantialist claim that space and time are in some sense independent of, and prior to, matter.²⁶

Under this approach: at a swerve world, every instant will be a counterpart of every other instant, even in strict contexts (since geometrically, there is nothing to choose between them). So the swerve theory will be indeterministic (and represent time as being homogeneous), since for each instant at a swerve world, the swerve *could have* occurred at that instant. Of course: for reasons given above in the discussion of cheap haecceitism, I think that it is natural to think of the resulting account as being one on which shifts generate new possibilities—but now the sense in which the swerve theory is indeterministic and the sense in which shifts generate new possibilities are thoroughly robust. One can be a shifty cheap haecceitist.

3.3 *Elvis fans*

What are we doing when we ask questions like whether shifts generate new possibilities, then look at the consequences of affirmative and negative answers? We are trying to get clear on the contours of our modal concepts as part of a larger

project of constructing some sort of analytic framework for making sense of modal discourse.

Back and forth about intuitions and cases (like that engaged in above) can play a role in helping us to understand how our modal concepts work—when there happens to be agreement about how the concepts apply to cases. But when it comes to the question of shifts (and the related question of haecceitism), such agreement is hard to come by.²⁷ Philosophical opinion is divided in part because philosophers on either side of the issue find that they just don't have the intuitions about cases that are so widespread on the other side of the divide. Locked in a dispute of this kind, one naturally begins to suspect that there is some indeterminacy in "our" concepts—so that the sort of cases we cannot agree about can safely be consigned to that thick and convenient file, Spoils to the Victor.

But sometimes the questions we face aren't as subtle and fiddly they appear. Philosophical opinion on the question whether shifts generate new possibilities may be divided, but there is a large group of people for whom the issue is clear: physicists. The vast majority of them recognize a distinction between two sorts of symmetry (see below) that presupposes that shifts *do* generate new possibilities. And there are a lot more physicists out there counting degrees of freedom than there are philosophers debating questions about shifts and haecceitism. So if we are serious about wanting to get clear on *our* concepts, then the situation is straightforward: we need to make room within our analytic framework for a cluster of modal concepts that ground the judgement that shifts generate new possibilities, since such a cluster is in wide use.²⁸

What is the physicists' distinction that I have in mind? It is one that is customarily made whenever an author has occasion to explain the notion of a gauge symmetry (or related notions, such as that of a gauge theory). Here are a couple of representative quotations.²⁹

General covariance is a *gauge symmetry*. As with other gauge symmetries, the term "symmetry" is a misnomer. Gauge symmetries are not symmetries of the Hilbert space; the Hilbert space is invariant under the entire gauge group. Instead, gauge symmetries represent a redundancy in our description of the theory. (Seiberg 2007, 169)

[G]auge theories are . . . deeply disturbing and unsatisfying in some sense: They are built on a redundancy of description. The electromagnetic gauge transformation $A_\mu \rightarrow A_\mu - \partial_\mu \Lambda$ is not truly a symmetry stating that two physical states have the same properties. Rather, it tells us that the two gauge potentials A_μ and $A_\mu - \partial_\mu \Lambda$ describe the same physical state. (Zee 2010, §III.4)

These authors are drawing attention to a completely standard distinction between two types of symmetries that relate solutions of the equations of theories—the mathematical objects that are used to represent possible histories.

GAUGE SYMMETRIES: symmetries that relate solutions that *cannot* be taken to represent distinct situations.

PHYSICAL SYMMETRIES: symmetries that relate solutions that *can* be taken to represent distinct situations.

Corresponding to a class of solutions related to one another by physical symmetries, there will be a multiplicity of physical situations represented by them.³⁰ The shift symmetries of classical physics are typically supposed to be the paradigm examples of physical symmetries.³¹

So there doesn't seem to be much room to deny that there is a large community of people according to whose modal concepts shifts generate new possibilities—large, that is, relative to the number of philosophers interested in these matters. So shiftless philosophers are engaged in the revisionary project of trying to construct new modal concepts to replace ones in common use. That may well be worthwhile. But unless they are imagining converting physicists to the use of these new concepts (a tall order, to say the least), they will still end up facing the problem the rest of us face: trying to construct an analytic framework to make sense of a collection of modal concepts according to which shifts generate new possibilities—since they will, presumably, still want to make sense of the modal concepts actually in use.

3.4 *Worries*

But wait! There is a difference between making a discovery and changing the subject! The discussion above makes it sound like the movement in recent years among philosophers of space and time away from shiftiness and towards shiftlessness amounts to the development and spread of a new package of modal concepts. But there is another possibility: this movement reflects a new clarity about how our good old modal concepts really work.

One can of course find examples of both sorts of process in the history of philosophical reflections on modality. In the case at hand, I think what we are seeing is something more like a change in the subject of discussion than a discovery. I will not drag the reader through a detailed examination of the relevant literature. Let me just say that it seems to me that defences of shiftlessness fall into one or more of three broad camps.³²

First, there are those that work backwards from the hole argument. The idea is that in thinking about general relativity, we learn that in that theory, isomorphic models never represent different possibilities—and this may seem to teach us that the received view about shifts in the Newtonian context was mistaken. A lot remains to be said about how we are supposed to get from the lesson of the hole argument to a thesis about shifts. But in any case, the starting point here is wrong-headed—the lesson of the hole argument has been misunderstood (or so I will argue in §4 below).

Second, there are those who take as their starting point a cheap haecceitist account of modality and argue that if this account is accepted, then we must deny that shifts and generalized shifts generate new possibilities. But, as we saw above, cheap haecceitism is compatible with shiftiness.

Third, we have approaches that attempt to parlay reflections on the notion of symmetry into views about which quantities are physically real—and that have as a

corollary that the location of the centre of mass of the universe is not a physically real quantity (and hence that shifts do not generate new possibilities). For such a procedure to be successful, it is crucial that we be able to determine the relevant symmetries of a physical theory without already knowing which quantities are physically real. Unfortunately, the standard mathematical notions of symmetries of physical theories are unsuitable: they count solutions as related by symmetries that no reasonable interpreter would want to count as representing the same physical situation.³³ It seems plausible, in fact, that no mathematical notion is suited to do the sort of work that philosophers expect here—if one wants a notion of symmetry that relates all and only solutions capable of representing the same physical situation, then one needs to start with a notion along the lines of “transformation under which physically real quantities transform in a suitable way.”³⁴

But wait! Physicists aren't philosophers! In particular, you can't expect them to be careful about the counting of possibilities—so any route from their casual remarks on this subject to claims about how their concepts work is going to be arduous.

About some issues that is clearly right. Consider a form of super-haecceitism according to which for any possible individual x , possible world w , and possible individual y in w , there is a possible world w^* that is qualitatively identical to w and in which x plays the role that y plays in w . The question whether our modal concepts accord with such a view turns on fine-grained questions that are of interest to some philosophers but to virtually no physicists—and so it would indeed be surprising if one could get very far in investigating this question by looking at remarks of physicists. But our question about shifts is not obviously so far removed from questions physicists care about. So let's try, and see how things work out.

Physicists have a systematic way of counting physical states. They say that a system of n particles has $3n$ degrees of freedom because it takes three numbers to specify the location of each particle relative to a chosen set of coordinates.

Shifty philosophers will think: when physicists say that n particles have $3n$ degrees of freedom, they mean that the possible ways that a (static) system of particles can be arranged in space form a $3n$ -dimensional space of possible worlds.

But shiftless philosophers can account for this practice via a revisionary semantics. On their view, it takes $3n$ numbers to specify the configuration of an n -particle subsystem of the universe—the ambient system provides a reference frame, and each way of specifying the coordinates of a particle relative to this frame corresponds to a possible way that a particle of the subsystem could be located. So when speaking of subsystems, counting degrees of freedom is the same thing as counting possible configurations. But when it comes to possible configurations of the universe as a whole, the story changes. Now it just takes just $3n - 6$ numbers to fix the configuration of an n particle system (since now we need to identify configurations that differ from one another by a translation or rotation in Euclidean space). So it takes a little care to properly interpret what physicists are up to when they count degrees of freedom. Degrees of freedom

directly correspond to ways the world could be when we are talking about states of subsystems, but not when we are talking about global states. But this distinction is invisible when physicists talk about counting degrees of freedom because they can't be bothered to fuss about the distinction between subsystem states and global states.

We have focussed on just one sort of example. But a structurally similar story will play out whenever we look at how physicists count states in a context in which there are no gauge symmetries in play. Shifty philosophers will see a direct correspondence between the instantaneous states recognized by physicists and (partial) time slices of possible worlds, while shiftless philosophers will recognize such a correspondence only when the system under consideration forms a subsystem of the universe—when universal states are in question, the shiftless will maintain that the physicists' space of states over-counts possibilities. Again, one might reasonably expect this distinction to be invisible in physicists' practice: when they are talking about universal states, they don't quite mean what they seem to say—but this is just because they can't be bothered to make the distinction between global states and subsystem states.

The subtler story preferred by the shiftless turns out to have impressive resources. According to a standard view, relativity principles say that there is no detectable difference between possible global situations that differ by a velocity boost.³⁵ This standard reading is unavailable on the subtle view under consideration—but it turns out to be possible (and interesting) to reconstrue relativity principles as being claims about the relations between subsystems of the universe.³⁶ One might worry that this subtle approach would efface the distinction between physical symmetries and gauge symmetries (since, roughly speaking, it tells you to identify global states related by either sort of symmetry)—but, again, it turns out to be possible to reconstruct this distinction in terms of the relations between states of subsystems of the universe.³⁷

So now we can put the worry at hand in sharp terms: it is natural to wonder whether the argument from Elvis fans above really has much force—one should not necessarily expect to be able to determine from the casual and pedagogical mutterings of physicists whether their position is best reconstructed in terms of the subtle account preferred by the shiftless or the straightforward account preferred by the shifty.

That is right: one shouldn't *expect* that. But in fact we can find a trickle of physicists explicitly advocating aspects of the subtle approach.³⁸ And one can also find a much broader stream in which the subtle approach is fairly explicitly rejected: an examination of the way in which physical states are individuated in the literature on monopoles shows that physicists working in that field are hip to the distinction between subsystem states and global states, but stubbornly persist in counting global states related by shifts as being physically distinct. So the revisionary semantics offered by fans of shiftlessness looks implausible—and the physicists' practice of counting degrees of freedom stands as a reason to think that our modal concepts count shifts as generating new possibilities.

Remark Counting Monopoles: A Few Details. At the classical level, a monopole is a magnetically-charged particle-like solution to the Yang–Mills–Higgs equations in which the energy of the system is well-localized in space.

Relative to a given frame of reference, it makes sense to ask: How many solutions are there in which a monopole of unit charge is at rest? There is a huge family of solutions of this kind—but every member in the family is related to every other one by a symmetry. Today it is standard to count most of these symmetries as gauge symmetries, but to count those that shift the location of the monopole or change its global phase as physical—according to this standard way of counting, there is a four-dimensional family of static monopoles of unit charge (it takes three real numbers to specify a location for the monopole, one to specify its global phase). Further, it turns out to be possible to have static multi-monopole configurations. According to the standard way of counting, the family of situations involving k static unit-charge monopoles is $4k$ -dimensional—it takes $3k$ parameters to specify the position in space of each of the k monopoles, and another k parameters to specify the global phase of each monopole.³⁹

This way of counting can be taken at face value by shifty philosophers. The shiftless will instead want to maintain: (i) that physicists are *really* thinking that while it takes $4k$ parameters to specify the state of a static k -monopole subsystem of the world, it only takes $4k - 4$ parameters to specify the state of a world containing nothing but k static monopoles (because the centre of mass of the system and its global phase are unphysical, while the relative distances and relative phases are physical); but (ii) physicists state their view in a misleading way, because they cannot be bothered fussing about the distinction between what it takes to specify the state of a subsystem and what it takes to specify the a global state.

But it turns out that in the present case this revisionary semantics is implausible. Not too long ago, physicists adhered to a *different* method of counting monopoles, according to which it takes $4k - 1$ parameters to fix the state of a static assemblage of k monopoles: $3k$ parameters to fix the locations of each of the monopoles; and $k - 1$ parameters to fix the *relative* phase relations between the monopoles.⁴⁰ Under this older scheme, the location of the centre of mass of an assemblage of monopoles was counted as physical but its global phase was counted as unphysical.

This older way of counting monopoles is itself difficult to square with the revisionary semantics. On this way of counting, in order to specify the state of each single-monopole subsystem, one has to specify both its position and its phase; but to specify the global state, one has to specify less information than is required to specify the state of each single-monopole subsystem. That may sound like good news for the shiftless philosopher. But the details are wrong: the shiftless philosopher says that there is in truth a $4k - 4$ parameter family of k -monopole situations (because neither shifts nor phase transformations generate new possibilities at the global level), which is misdescribed as a $4k$ parameter family by those not bothering about the distinction between subsystem states and global states. There is no room in this scheme for people who count monopoles in the old way—and hence who recognize the distinction in question—but who nonetheless maintain that global shifts generate new possibilities while global phase transformations do not.

Further: the fact that the transition from the older way of counting monopoles to the newer way took place in living memory undermines the plausibility of the revisionary semantics strategy for explaining away the newer way of counting (requiring as it does the attribution to physicists of ignorance of, or disinterest in, the distinction between the sort of information required to specify global states and the sort required to specify states of a subsystem).⁴¹

A further challenge for anyone hoping to defend the shiftless philosopher's revisionary semantics: make sense of the fact that Witten, within a single set of lectures, takes global phase to be unphysical in the vacuum sector of a Yang–Mills–Higgs theory, but physical in the monopole sector of the same theory.⁴²

4. General Relativity and Generalized Shifts

A solution (V, g) to the equations of general relativity consists of a four-dimensional manifold V (think of a smooth space that has topology but no geometry) together with a metric tensor g that endows V with a nomologically possible spatiotemporal geometry.⁴³ Let (V, g_1) and (V, g_2) be distinct solutions defined on the same manifold V . Because the metrics g_1 and g_2 are defined on the same spacetime manifold, the two solutions agree about a number of facts (e.g., concerning the topology of spacetime). Because they are distinct, g_1 and g_2 disagree about the assignment of geometric roles to points of the spacetime manifold V . There are two ways they might disagree. On the one hand, g_1 and g_2 might disagree about which geometric roles are instantiated—e.g., g_1 might say that there are regions of very intense curvature, while g_2 might say that there are no such regions. In this case, we say that g_1 and g_2 disagree about the geometry of spacetime. On the other hand, g_1 and g_2 might agree about which roles are instantiated and about their pattern of instantiation, but disagree about which points in V instantiate which roles—e.g., g_1 and g_2 might agree that there is a region of intense curvature (and about much else besides), but disagree about whether a given point $x \in V$ lies in that region or elsewhere.⁴⁴ In this case, we say that g_1 and g_2 differ by a generalized shift. This means that there is well-behaved map from V to itself, with the feature that the image of a given point under this map is assigned the same geometric role by g_2 that the original point was assigned by g_1 .⁴⁵

How should we think of generalized shifts? In the Newtonian case, most (almost all?) physicists and, until very recently, almost all philosophers took shifts to generate new possibilities. But there is a powerful reason—the hole argument—to think that it would be a mistake to take generalized shifts to generate new general relativistic possibilities.⁴⁶

THE HOLE ARGUMENT. Let V be a four-dimensional manifold, g a metric tensor that endows it with a spatiotemporal geometry, and Σ a three-dimensional subset of V that represents an instant of time in the spacetime (V, g) . Then there will be generalized shifts that act nontrivially only to the future of Σ . A spacetime metric g^* on V that is related to g by such a generalized shift will agree with g about the spacetime geometry and also about the way that geometric roles are distributed among points in V to the past of Σ , but will

disagree with g about which roles are assigned to (at least some) points to the future of Σ .⁴⁷ If such generalized shifts generate new possibilities, then the theory is indeterministic—since there will be distinct histories that match perfectly up until the instant represented by Σ . But: no physicist would agree that general relativity is grossly indeterministic in this way.⁴⁸ So: we need an interpretation of the theory on which generalized shifts do not generate new possibilities.

Philosophers of physics have by and large bought this.⁴⁹ So they have sought interpretations of general relativity on which generalized shifts don't generate new possibilities, typically following one or the other of two main paths.

STRAIGHT-UP CHEAP HAECCEITISM. Adopt a qualitative counterpart theory and deny the existence of worlds that are qualitative duplicates of one another. Then one can maintain that there is only one possible world corresponding to a whole family of mathematical spacetimes related to one another by generalized shifts. The threat of indeterminism vanishes: no two worlds are related to one another in the same way as solutions related by a generalized shift.⁵⁰

METRIC ESSENTIALISM. Maintain that points of spacetime possess their geometric properties essentially. Then it is impossible that a point of spacetime could exist in, but play different geometric roles at, multiple worlds. So generalized shifts don't generate new possibilities. So general relativity is not grossly indeterministic in the way threatened by the hole argument.⁵¹

Most philosophers seem to think that one or the other of these approaches provides a satisfactory response to the hole argument. Others think that there is an important open problem here.⁵² But very nearly everyone seems to agree that on an acceptable interpretation of general relativity, no generalized shifts generate new possibilities.

My aim here is to move the goal posts. Look again at the hole argument as characterized above. There is an obvious gap. At best, the argument shows that we should deny that *certain* generalized shifts generate new possibilities—namely, those that involve us in indeterminism if they are not handled that way. But it does not obviously follow that we should deny that *any* generalized shifts generate new possibilities. I am going to argue that there are indeed cases in which we should hold that some generalized shifts generate new possibilities. It follows that we should not be satisfied with responses to the hole argument that imply that no such cases exist (as do the two dominant approaches mentioned above).

4.1 Gauge equivalence

Fix a four-dimensional manifold V and consider the set \mathcal{S} of metrics on V that obey the equations of general relativity—each element \mathcal{S} corresponds to an admissible spacetime metric g defined on V . We are interested in two equivalence relations on \mathcal{S} . (i) *Isometry* (the notion of isomorphism proper to spacetime geometries): distinct metrics are isometric just in case they encode the same four-dimensional geometry for V , differing only as to which point of V plays which geometric role (in other words, distinct metrics are isometric if and only if they are related by a generalized shift). (ii) *Gauge equivalence*: metrics on V count as gauge equivalent just in case they have to be viewed as representing the same possible situation (because they

jointly represent just one possibility). The first relation is a strictly mathematical one, but the second one has an interpretative dimension—interpretations can differ as to whether two given metrics in \mathcal{S} are gauge equivalent.

The moral that philosophers have generally drawn from the hole argument is that we should look for an interpretation of general relativity on which spacetime metrics on V are gauge equivalent if and only if they are isometric. In support of this, they often quote scripture: “diffeomorphisms comprise the gauge freedom of general relativity.”⁵³

But rather than rushing to that conclusion, let’s begin by looking at a more general question: How do physicists normally decide whether to count two solutions of the equations of some theory as gauge equivalent?

Physicists have a number of fancy machines that tell them which symmetries of a given (physically reasonable) classical theory are gauge and which are physical (including Dirac’s theory of constrained Hamiltonian systems and its modern elaborations and a distinct but related approach that exploits the technology of covariant phase spaces). But the basic idea is: *count a symmetry of a classical theory as physical, unless doing so spoils the determinism of the theory.*⁵⁴

GAUGE EQUIVALENCE. Call two isomorphic solutions defined on a manifold V *spoiler-related* if they differ only to the future (or past) of some region $\Sigma \subset V$ representing an instant of time. Take *gauge equivalence* to be the most discerning equivalence relation that extends spoiler-relatedness.

These machines are extremely general in their field of application. But they do come with one restriction: they are geared to theories whose equations can be derived from a certain sort of function called a Lagrangian (the derivation of equations from a Lagrangian involves a sort of maximum principle—think of the way that the laws of reflection and refraction for light rays can be derived from the principle that light takes the quickest path).⁵⁵ In fact, the family of theories that arise in this way includes just about any classical theory that a physicist would be likely to write down. In practice, such theories fall into two categories.

- (i) **Nondegenerate Lagrangians:** If the theory’s Lagrangian is well-behaved in a certain technical sense, then each solution is gauge equivalent only to itself—so every symmetry is considered physical.
- (ii) **Degenerate Lagrangians:** If the theory’s Lagrangian fails to be well-behaved in this sense, then gauge equivalence classes will be vast (indeed, infinite-dimensional)—and infinitely many symmetries will be classified as gauge symmetries.⁵⁶

We are going to consider what happens when we run general relativity through these machines. The answer is going to be: that depends. General relativity breaks into a number of sectors, corresponding to different sorts of universes one might be interested in modelling. In some sectors, the relations of gauge equivalence and isometry are coextensive relations; in others, gauge equivalence is strictly more

discerning than isometry. In order to get a feeling for how this works, it is helpful to begin with some toy examples.

4.2 Squiggles on graph paper

Here are three unexciting games that involve drawing curves on (possibly infinite) sheets of graph paper. In each game there are rules that determine which drawings count as admissible and two equivalence relations on the set of admissible drawings. At this point, the selection of equivalence relations for each game probably ought to seem fairly arbitrary.

GAME I. Our sheet of graph paper is infinite in extent. An admissible drawing consists of a single non-horizontal straight black line. Two drawings count as *isomorphic* if their lines have the same slope (i.e., make the same angles with the vertical and horizontal background lines on the graph paper). Two drawings count as *equivalent* if their respective lines occupy exactly the same region of the graph paper. If two drawings are equivalent then they are also of course isomorphic. But two isomorphic drawings can be inequivalent—changing the location of a drawing's line while leaving its slope unchanged will result in a drawing that is isomorphic to the original, but not equivalent to it.

GAME II. Beginning with a sheet of graph paper that is finite in horizontal extent, roll it up into a cylinder with with a vertical axis. An admissible drawing consists of any number of coloured curves, each of which is closed and wraps around the cylinder once. No curve is allowed to intersect itself or any other. Curves are allowed to be as straight or as squiggly as you like. Two drawings are *isomorphic* if and only if they involve the same colours, occurring in the same order. Two drawings are *equivalent* if and only if they are isomorphic.

In the first game, the rich structure of the graph paper plays a role in demarcating the admissible drawings and in defining isomorphism and equivalence—we pay attention to angles that the drawn line makes to the horizontal and vertical background lines and even to the precise location of the drawn line on the background. In the second game, we make use only of the topological structure of the graph paper—we need to distinguish between continuous curves and others, between closed curves and open curves, between curves that wrap once around the cylinder and those that don't, and between curves that intersect and curves that don't. But we don't pay any attention to geometric distinctions like horizontal vs. vertical, straight vs. curved, or short vs. long. It is also possible to cook up games that latch on to an amount of structure intermediate between that relevant to Game I and that relevant to Game II.

GAME III. Begin with an infinitely extended sheet of graph paper. We again draw a family of coloured curves. The colours are again arbitrary. Again, intersections are forbidden. But now the curves drawn are required to be infinitely long and *asymptotically horizontal*: for each drawn curve, there is a horizontal line on the page such that the drawn curve approaches the horizontal line closer and closer without limit as you go further and further out towards both the left and the right. In this game, as in the preceding one, two drawings count as *isomorphic*

just in case they feature the same colours, in the same order. Two drawings count as *equivalent* just in case they are isomorphic and, further, for each coloured curve in one drawing there is curve of the same colour in the other, with both curves asymptotic to the *same* horizontal line on the page.

In Game I, if we seek to make a drawing equivalent to a given one, there is only one thing we can do: make an exact reproduction of the given drawing. There is more latitude in creating a new drawing isomorphic to a given drawing: lay down a sheet of glass on the graph paper and trace the black line from the given drawing onto the glass; erase the line in the given drawing and slide the glass as far vertically and/or horizontally as you like; now trace the black line from the glass on to the graph paper.⁵⁷

In Game II, isomorphism and equivalence coincide. Given an admissible drawing, there are many ways to create a new one isomorphic/equivalent to it: trace the lines of the given drawing on to a glass cylinder and then erase them from the graph paper; now stretch and compress the graph paper cylinder (now thought of as being made of rubber . . .) as you like and then trace the lines from the glass cylinder back on to the graph paper (while it is stressed); now let the graph paper return to its resting state. In general the new lines will be in very different places from the old ones—and will probably be very different shapes. None of that matters for our purposes though: this process will leave the colours and their order invariant.

The situation in Game III is a sort of hybrid of the other cases. Given an admissible drawing, we create another isomorphic to the first as follows: (i) lay a glass plate on the original drawing; trace the coloured curves onto the glass plate and erase them from the graph paper; (ii) stretch and compress the graph paper as you like, but make sure that as you go out towards the left and right limits of each horizontal line, the effects of your stretching and compressing shrink to zero; (iii) now lay down the glass plate on the stressed graph paper however you like, so long as the coloured curves are asymptotic to horizontal lines; (iv) trace the coloured curves back on to the stressed graph paper and allow it to return to its relaxed state. The recipe for creating a drawing equivalent to the given one differs in just one way: at Step (iii) you have to assure that each coloured curve on the plate is asymptotic to the *same* horizontal line that it was asymptotic to in the original drawing. As in Game II, isomorphic drawings have to feature the same colours in the same order—but they will in general differ radically about the shapes and precise locations of their curves (although they will of course agree that every curve “becomes horizontal at infinity”). But the difference between isomorphism and equivalence is similar to that in Game I. In particular: in Game I, it makes sense to speak of two isomorphic but inequivalent drawings as being related by a vertical shift of some given amount; in Game II this is of course impossible; but in Game III it makes sense to speak of two isomorphic but inequivalent drawings as being related by a “vertical shift at infinity” of some given amount (corresponding to the distance along the vertical direction that the glass plate is slid between step (i) and step (iii)).

4.3 General relativity and temporal shifts at infinity

The solutions of the equations of motion of a classical theory can be thought of as being painted onto the mathematical spacetime V that is used in formulating those equations. But V and its structure play very different roles in different types of theory.

Consider a very simple (but, for present purposes, representative) example of a theory with a nondegenerate Lagrangian: a theory of a single free particle moving in two-dimensional Newtonian spacetime. A solution will take the form of a single straight line on Newtonian spacetime. If we use a sheet of graph paper to represent spacetime, with the convention that horizontal lines correspond to instants of time and vertical lines to points of space, then the line representing our particle will be non-horizontal. And the solutions to our theory's equations of motion will correspond to the admissible drawings of Game I above. As in Game I, solutions are isomorphic, by the natural standard for such objects, if and only if the lines they use to represent the particle's history have the same slope. And since the Lagrangian encoding the laws of the theory is nondegenerate, each solution is gauge equivalent only to itself. So there will be isomorphic solutions that are not gauge equivalent: isomorphic solutions that can be taken to correspond to distinct possibilities. Example: distinct solutions related by a temporal shift (these will be solutions that agree about that the particle has a certain non-zero absolute velocity, but which disagree about when it passes through a given point of space).

In theories with a degenerate Lagrangian, some solutions that differ strikingly count as gauge equivalent—so there are all kinds of ways that solutions can differ that do not correspond to ways that physical situations can differ. What sort of differences these will be vary from theory to theory: but our concern is with theories like general relativity, in which the relevant differences concern which points in the four-dimensional space V are assigned which geometric roles.⁵⁸ This means taking some aspects of the structure of V far less seriously than we do in the sort of case considered in the preceding paragraph.

For present purposes, the crucial point is that general relativity has sectors in which gauge equivalence and isometry coincide (in which generalized shifts do not generate new possibilities) and sectors in which gauge equivalence is a more discerning relation than isometry (in which there are isometric solutions that are capable of representing distinct possibilities). We will look at one example of each type.

A Cosmological Sector Various mathematical obstructions stand in the way of studying rich families of solution in which matter fills infinite space.⁵⁹ For this reason, in cosmological applications of general relativity one typically restricts attention to the case in which space is finite in extent. To this end, let us consider the sector of general relativity in which V is a sort of four-dimensional cylinder, each cross-section of which is a three-dimensional sphere. A solution of Einstein's field equations will be a metric tensor g defined on V that describes the geometry of spacetime (and of any matter whose configuration supervenes on geometry). A

typical solution g will assign each point of V a different geometrical role (e.g., possessing a set of curvature properties unique to it).⁶⁰

When we run this version of general relativity through the various machines that physicists use to identify gauge symmetries, the verdict is unanimous: two solutions are gauge equivalent if and only if isometric.⁶¹

This means, roughly speaking, that in this sector of general relativity we ignore any structure V has beyond its topological structure (which allows us to make a distinction between continuous curves and discontinuous curves, between curves that intersect and curves that don't intersect, and so on). So this version of general relativity is analogous to Game II above.

If we are willing to do a bit of violence to our formulation of general relativity, we can make the analogy more striking. Let us impose a slicing scheme: a rule that, given a solution (V, g) , breaks V up into instants of time in a unique way.⁶² So relative to a metric g on V , we get a family of three-dimensional subspaces of V , each with the topological structure of a three-sphere and with a spatial metric structure determined by g . So relative to a slicing scheme, each solution (V, g) corresponds to a history of the geometry of space. In Game II an admissible drawing consists of some coloured rings drawn on the cylinder, with drawings counted as isomorphic/equivalent just in case they featured the same colours in the same order. For our first version of general relativity, two solutions count as isometric/gauge equivalent if and only if, relative to our slicing scheme, they agree about the various geometries that space instantiates, as time passes, and about the order in which those spatial geometries occur.

Two isometric/gauge equivalent solutions that agree that an instant with certain spatial geometry occurs in V will not in general agree about which three-dimensional subset of V instantiates that spatial geometry—intuitively, such solutions agree about which events happen, but not “where” or “when” they happen within the mathematical space V . But since gauge equivalent solutions are incapable of representing distinct physical solutions, this sort of disagreement is effaced as we pass from our mathematical representations to the possible worlds that they represent—the sort of general relativistic worlds currently under discussion cannot differ only as to where or when things happen.

An Asymptotically Flat Sector Consider a second version of general relativity, this time with some boundary conditions at spatial infinity. The idea is to restrict attention to solutions with the feature that, for any given point and any given spatial direction at that point, in order to find a small patch of spacetime that looks as similar as you like to a given small patch of Minkowski spacetime, it suffices to proceed a certain distance outwards from the given point in the given direction. Solutions with this feature are said to be *asymptotically flat at spatial infinity*.⁶³ One way to make such a requirement precise is to choose as our mathematical spacetime V the space \mathbb{R}^4 of ordered quadruples, (t, x, y, z) of real numbers and to take V to be equipped with a non-physical background metric η endowing it with the flat geometry of Minkowski spacetime.⁶⁴ One then restricts attention to those solutions g of Einstein's field equations with the feature that if one goes far enough

towards spatial infinity in any direction, one can find sufficiently small regions of V such that the geometry of such a region according to g is arbitrarily similar to the geometry of that region according to η .⁶⁵

When we run this theory through the physicists' machines for identifying gauge symmetries, the verdict is that two solutions can be isometric without being gauge equivalent—in particular, a pair of isometric solutions can differ by a non-trivial time translation at spatial infinity.⁶⁶

The precise details vary somewhat from treatment to treatment.⁶⁷ For the sake of concreteness, let us consider the approach of Ashtekar *et al.* (1991). Here solutions are gauge equivalent if and only if they are related by a generalized shift that is asymptotic to the identity at spatial infinity (i.e., relative to the background metric, points related by this generalized shift become arbitrarily close together as one goes out towards infinity in any spatial direction). So this sector of general relativity is analogous to Game III above.

We can again make the analogy more striking by employing a standard slicing scheme that applies to solutions that are asymptotically flat at spatial infinity.⁶⁸ Under this scheme, each solution can be sliced into disjoint instants of time, each of which is asymptotic to an ordinary flat hypersurface of simultaneity of the background Minkowski metric—further, the decomposition is unique, except in the trivial case of Minkowski spacetime. So we can think of each solution g as telling us about the evolution in time of the geometry of space. Solutions g_1 and g_2 are isometric if and only if they correspond to the same sequence of spatial geometries. In order to be gauge equivalent, solutions g_1 and g_2 must satisfy a stronger condition:

For any instant Σ_1 in the special slicing of V relative to g_1 , there is a instant Σ_2 relative to the special slicing relative to g_2 , such that g_1 assigns the same geometry to Σ_1 that g_2 assigns to Σ_2 —and, further, the Σ_1 and Σ_2 are asymptotic at spatial infinity to the same flat surface of simultaneity of the non-physical background metric.

Isometric solutions agree about the sorts of spatial geometries that occur relative to the slicing scheme discussed above—but will in general disagree, for any such spatial geometry, about which set of points of V instantiates that geometry. Gauge equivalent solutions agree about more: subsets of V that they respectively view as instantiating a given spatial geometry “agree at infinity”—in the sense that they each asymptotically approaches the same instant of time of the background metric on V . So in this sector of general relativity, as in Game III above, asymptotic conditions play a role in determining which representations are legitimate and also in providing structure that can be used in defining an equivalence relation more discerning than isomorphism/isometry.

Further, in analogy with the behaviour we saw in Game III above, we find that in this sector of general relativity, two isometric solutions can fail to be gauge equivalent because they “differ from one another by a time translation at spatial infinity.”⁶⁹ Faced with such a pair of isometric solutions (= same pattern of events, instantiated differently on V) that are not gauge equivalent (= they are capable

of representing distinct physical possibilities), we conclude that there is a pair of physical possibilities that similarly differ by a temporal translation at spatial infinity. So, contrary to the common wisdom among philosophers, two general relativistic worlds *can* differ only as to when things happen, in the sense of differing by a time translation at infinity (and a parallel considerations show that two general relativistic worlds can differ by a spatial translation or rotation at spatial infinity, so that there is a sense in which two general relativistic worlds can differ only as to where things happen).

4.4 Worries

What is going on here? These two versions of general relativity correspond to different physical regimes. Solutions in the cosmological sector are used to provide descriptions of complete matter-filled universes (with a fixed closed spatial topology). Solutions in the asymptotically flat sector provide descriptions of finite isolated self-gravitating systems (in the case of vanishing cosmological constant). As Penrose (1982, 632) remarks:

Asymptotically flat spacetimes are interesting, not because they are thought to be realistic models for the entire universe, but because they describe the gravitational fields of isolated systems, and because it is only with asymptotic flatness that general relativity begins to relate in a clear way to many of the important aspects of the rest of physics, such as energy, momentum, radiation, etc.

And as Einstein says in his 1917 paper on cosmology, when we require that the spacetime metric approach the Minkowski metric as we go out towards infinity, we in effect “pre-suppose a definite choice of the system of reference, which is contrary to the spirit of the relativity principle” (i.e., we introduce structure that ruins full general covariance).⁷⁰ This extra “background structure at infinity” allows us to make finer-grained distinctions between isometric solutions than we can in the cosmological case—it becomes possible, e.g., to think of two such solutions as differing by a “time translation at infinity.” And the official physicists’ distinction between gauge and physical symmetries exploits this possibility. (The next two unpleasant paragraphs aim to flesh out the claims just made. They can easily be skipped.)

Today we say (without Einstein’s tone of disapproval): to impose asymptotic flatness at spatial infinity, is to single out a preferred class of inertial frames at infinity. Two inertial frames at infinity can be related to one another by a boost, a rotation, or a spatial or temporal translation—in short, the set of symmetries of the family of frames at infinity looks just like the set of symmetries of the background Minkowski geometry of our framework.⁷¹ A generalized shift relating two asymptotically flat solutions induces a map from the set of preferred frames at infinity to itself. So each generalized shift is associated with a symmetry of the background Minkowski spacetime structure.⁷² We call a generalized shift *asymptotically trivial* if it is associated in this way with the trivial identity symmetry of Minkowski spacetime.

Suppose now that asymptotically flat solutions g_1 and g_2 are spoiler-related. Then there is a generalized shift relating g_1 and g_2 that acts as the identity to the past (or future) of some instant of time. So the symmetry of the background Minkowski geometry associated with this generalized shift must also act trivially to the past (future) of that instant of time. But that means that this symmetry must just be the identity map. So the generalized shift in question must be asymptotically trivial. It follows that two metrics are gauge-equivalent, then they must be related by an asymptotically trivial generalized shift (since gauge equivalence is essentially just the transitive closure of spoiler-relatedness).

Okay. But Really—What is Going on Here? That isometric solutions of general relativity always count as gauge equivalent is all but universally accepted among philosophers of physics. And this article of faith can be (and, often, is) supported by appeal to explicit affirmations of this thesis in canonical textbook presentations of the theory (e.g., Hawking and Ellis 1973, 68 and Wald 1984, 260 and 438). So it certainly *looks* like the supposedly official notion of gauge equivalence just gives the wrong answer when applied to our asymptotically flat sector of general relativity, since its pronouncements apparently diverge from the judgements of relativists.

The truth of the matter is this, I believe: while relativists *do* often speak as if solutions of general relativity are gauge equivalent if and only if isometric, they drop this way of speaking when asymptotic boundary conditions (like asymptotic flatness at spatial infinity) are in view. Wald, for instance, employs a more restrictive notion of gauge equivalence when he discusses the Hamiltonian formulation of the asymptotically flat sector of general relativity in his textbook.⁷³

Why, though? What is the point of using a notion of gauge-equivalence more fine-grained than isometry in the presence of asymptotic boundary conditions? There is much to be said here. But here are two benefits of this approach.

- 1) Physicists think of the true degrees of freedom of a theory as being encapsulated in the space that results when one begins with the space of solutions, and then identifies gauge equivalent solutions. So one version of our question is: In the regime in which asymptotic flatness is imposed at spatial infinity, why think of the true degrees of freedom as being given by the space of solutions modulo the relation of being related by an asymptotically trivial generalized shift, rather than the relation of being related by any old generalized shift? One powerful consideration is that the Poincaré group (the family of symmetries of Minkowski spacetime) acts on the former space in an interesting way (its action on the latter space is trivial)—and that one finds in this case the usual rich and fruitful relation between symmetries of a physical theory and conserved quantities.⁷⁴ Thus, corresponding to any family of time translations at infinity will be a notion of the energy of an instantaneous state, which will be conserved over time in any asymptotically flat solution. With this notion in hand, we can, e.g., note that the physics of isolated systems is time-translation invariant in general relativity (with vanishing cosmological constant)—for any

solution, you can find another differing from it by a time translation at infinity, in which the same instantaneous states occur in the same order.

- 2) Under standard approaches to quantizing classical theories, classical physical quantities correspond to (self-adjoint) operators on the quantum Hilbert space, each of which is associated with a set of real numbers, its *spectrum*. In a quantum theory, the spectrum of a quantity is the set of possible values of measurement outcomes. Now: in typical physical theories, there is an association between symmetries and physical quantities—for any (continuous) symmetry, there will be certain physical quantities that “generate” it. Under standard procedures of quantization, quantities that generate gauge symmetries are quantized by operators with zero as the *only* member of their spectrum.⁷⁵ In the asymptotically flat sector of general relativity that we have been discussing, the physical quantity that generates asymptotic rotations is angular momentum. So if one decides to count all generalized shifts as gauge transformations, this will require one, in particular, to count asymptotic rotations as gauge symmetries—and this will be reflected at the quantum level as a prohibition on states with non-zero angular momentum.⁷⁶ One could conceivably have reasons for thinking that rotating systems should be impossible in quantum gravity—but it would be outrageous to impose this by fiat.

But wait! Why this setup? Some readers will be frustrated with talk of ‘sectors’—surely it is general relativity *as a whole* that we want to talk about. And if we think of the full family of solutions, surely the only reasonable criterion of gauge equivalence is isometry. Why not think of the theory as a whole, as just one big collection of mathematical spacetimes—say, all pairs (V, g) where V is a four-dimensional manifold and g is a suitably well-behaved metric obeying the Einstein field equations?

One thing to note is that for many purposes, it is essential to think of the solutions of the equations of general relativity as forming structured spaces, rather than a mere jumble.

CAUCHY STABILITY Specifying a well-behaved instantaneous state of the gravitational field suffices to determine a unique well-behaved spacetime geometry in which that state occurs. Further: there is an interesting sense in which the spacetime geometry so-determined depends continuously on the instantaneous state specified (if you specify a sequence of instantaneous states that converges to a given instantaneous state, then there is a sense in which the corresponding sequence of spacetime geometries converges to the spacetime geometry determined by the given instantaneous state).⁷⁷

NON-LINEAR STABILITY OF MINKOWSKI SPACETIME. The spacetime geometry determined by an instantaneous state sufficiently close to a (well-behaved) instantaneous state that arises in Minkowski spacetime is globally similar to Minkowski spacetime—in particular, the given instantaneous state has a future that looks as much like Minkowski spacetime as you like, as you proceed to infinity along any geodesic.⁷⁸

Stability results of these kinds are preconditions for the applicability of the theory. In the absence of Cauchy stability we would require precise knowledge of the current instantaneous state in order to be able to use the theory to make predictions even on short timescales. And most known solutions to the equations of general relativity are highly symmetric. Since a glance at the night sky shows that our universe is *not* endowed with perfect symmetry, in order for the global features of such a solution to be of physical interest, a symmetric solution must be nonlinearly stable in the above sense—for otherwise, even if we were to discover that some region of the universe was in a state very similar to an instantaneous state arising in the given solution, we could not conclude that the distant past or future evolution of that region would be relevantly similar to the what happens in the given solution.⁷⁹ Thus, it was a grave blow against the physical interest of the Einstein static universe when Eddington (1930, §5) showed that it was *not* stable in this sense—if you slightly decrease the density of matter in an instantaneous state, then the universe, rather than being approximately static, undergoes a process of eternal exponential expansion.

For our purposes, the point is that claims about stability only make sense once the relevant space of initial data has been equipped with (at least) topological structure. Standardly, one proceeds here by restricting attention to initial states in which space has some given, fixed, topology.⁸⁰

And when it comes to trying to understand the notion of mass/energy in general relativity (and related questions about the sense in which the dynamics of the theory is time-translation invariant), it becomes crucial not just to look at metrics on a manifold V that looks like \mathbb{R}^4 at infinity, but also to impose pretty strong asymptotic boundary conditions—if one wants mass to be a well-defined conserved quantity on well-behaved phase space for general relativity, then it appears that one has to restrict attention to metrics that are asymptotically flat at spatial infinity in a fairly tightly controlled sense.⁸¹ Further, as was alluded to above, the analytic techniques used to prove results in the cosmological and asymptotically flat sectors can't be used if one attempts to study generic spatially open universes (because certain crucial integrals that are finite in the tame cases diverge absent strong boundary conditions at spatial infinity).⁸²

It is for reasons like this that it is natural to investigate the asymptotically flat sector in isolation from other sectors of the theory. Still, one might think, couldn't we in principle characterize asymptotically flat sector not by looking at solutions asymptotically flat relative to some fixed background Minkowski metric on V , but rather by looking at the complete family of such solutions that arise for *all* ways of putting a background Minkowski metric on V ? If we were to do this, it would again become impossible to think of one solution as differing or as not differing from any other by a time translation at infinity—and so would be back to thinking of all generalized shifts as gauge transformations!

This is a coherent way to proceed. However, it is rarely adopted, in part because it breaks the prized connection between conserved quantities and physical symmetries (energy and angular momentum, for instance, would still be defined on this picture and would still be conserved within each solution—but one wouldn't

be able to think of them as generating time translations and rotations at infinity, because one would have thrown away the structure required to make sense of these notions). Here is a possibly too brief way of formulating what to me seems like the right way to think about the situation.⁸³ (i) The approach just sketched stands to the approach employed above as a generally covariant formulation of a field theory set in Minkowski spacetime stands to a standard, non-generally covariant formulation that uses a fixed physical background metric.⁸⁴ (ii) In each pair, the first member might fairly be accused of being merely fakely generally covariant—the fact is that some solution-independent geometric structure is being used to make sense of the equations of the theory, but this fact is being hidden. (iii) It is, in general, a good policy to make background structure manifest when it is being employed.⁸⁵

But wait! Surely here we are undeniably talking about subsystems of the universe! As Penrose was quoted saying above: asymptotically flat solutions provide idealized models of relatively isolated self-gravitating subsystems of our universe. So it may well seem that in this context it is safe to fall in with the sort of account canvassed in Section 3.4 above, on which symmetries are to be understood as generating new possibilities when they act on the state of a subsystem of the universe, but not when they act on global histories. And then we can set aside these funny asymptotic symmetries as irrelevant, and go back to flat-out denying that generalized shifts generate new possibilities.

Well—there is much that is controversial in this line of thought. In particular, its presupposition that a family of solutions is worth taking interpretatively seriously only if it offers realistic models of actual phenomena.

But there is no need to litigate that question here because there is a far more brutal and effective way of dealing with this sort of reason for denying that any generalized shifts generate new possibilities. These days, we believe that the universe is undergoing accelerated expansion. So we think that the physically correct version of the equations of general relativity includes a term with a positive cosmological constant Λ . This means that isolated systems should look asymptotically like de Sitter spacetime (the spacetime of maximal symmetry in the $\Lambda > 0$ regime) rather than Minkowski spacetime (the spacetime of maximal symmetry when $\Lambda = 0$). In the asymptotically de Sitter sector of general relativity, as in the asymptotically flat (=asymptotically Minkowski) sector, one finds that asymptotic symmetries are associated with conserved quantities in the usual way—and one again finds that every generalized shift can be factored into such an asymptotic symmetry and a generalized shift that acts trivially at infinity.⁸⁶ So the sort of reasons that we had above for thinking that only some generalized shifts should count as gauge rather than physical reappear in the context of asymptotically de Sitter spacetimes. But now there is a twist. When $\Lambda > 0$, it is not just isolated systems that tend to look asymptotically like de Sitter spacetime: the same is true for a wide variety of models of the universe as a whole.⁸⁷ In this context, asymptotic boundary conditions and the distinction between isometry and gauge equivalence that travels in their wake are not just for subsystems.

5. Conclusion

So there you have it. We have seen some reasons for thinking that we should (do what we can to) take shifts to generate new possibilities. Some of these were endogenous to the usual dialectic (intuitive judgements about costs and benefits of different ways of counting possibilities), some exogenous to it (driven by a concern to provide an account adequate to the modal concepts that appear to underlie the way that physicists count degrees of freedom). And we have seen some reasons (again arising out of attention to how physicists handle related questions) for thinking that we should (do what we can to) take *some* generalized shifts to generate new possibilities in the context of general relativity.

Let me end with a few remarks about how the considerations adduced above bear on a question about laws of nature.

QUESTION: How does one specify the laws that obtain at possible world w ?

THE STANDARD ANSWER: By singling out those true generalizations at w that hold of nomological necessity.

A NON-STANDARD ANSWER: By specifying a set of worlds that share their laws with w —and some structure on that set.

This standard answer should be given, e.g., by advocates of best-systems accounts and also by Armstrong-Dretske-Tooley. It is, I take it, the most common way of thinking about laws among philosophers (including those who are primitivists or necessitarians about laws). The non-standard idea is of a piece with some ideas that are in the air these days.⁸⁸ I think that it is worth a close look.⁸⁹

I maintain that there are interesting cases in which the standard conception of what it takes to specify laws leads to funny consequences—and where the non-standard idea fares better.

EXAMPLE: Consider a swerve world w . For the sake of argument, suppose that a single basic regularity holds of nomological necessity there: all material events lie on a single worldline, initially vertical, always straight except for a single elbow. But we can know that that regularity holds of nomological necessity at w without knowing whether physics at w is deterministic or time-translation invariant. Plausibly, we don't know what the laws are at a world unless we know whether or not they satisfy basic conditions of this sort. In order to fully know what the laws are at w , we need to know more than which true regularities are nomologically necessary at w —we need to know whether there are other worlds at which the same regularities are nomologically true and whether these worlds are related to one another by time translation.

EXAMPLE: At a general relativistic world, the nomologically true regularities are, presumably, given by Einstein's field equations and by any asymptotic boundary conditions that obtain. But in the asymptotically flat regime, fixing these leaves open the question of how to count possibilities—Do we treat all generalized shifts as gauge, or just those that are asymptotically trivial? And in leaving this question open, we again leave open the question whether physics is time-translation invariant. But, presumably, fully knowing the laws involves knowing

the answer to that question. In order to fully fix the laws, then, we need to go beyond saying which equations and boundary conditions obtain, and to say something about how many possibilities share these nomological regularities, and about how they are related to one another.

EXAMPLE: Certain simple classical mechanical systems (the free particle, the harmonic oscillator, the two-body problem) have equations of motion that can be derived from multiple Lagrangians, with distinct Lagrangians leading to distinct quantizations (disagreeing, e.g., about energy spectra and about which quantities are conjugate to one another).⁹⁰ A way to think of this: the set of classical free particle worlds (for example) breaks into sectors; distinct sectors contain the same Humean mosaics as one another; each sector is also associated with a different Lagrangian that induces on the set of worlds in that sector a distinctive geometry; each sector is associated with a particular quantum theory (the one that is its closest neighbour in logical space). Merely knowing which regularities/equations of motion hold of nomological necessity leaves something open—something that is settled either by saying how the relevant family of worlds is geometrically organized or by saying which quantum worlds are most closely related to these worlds.

The standard picture should go, I think. The non-standard idea clearly needs development. But it has some interesting implications for metaphysics and for philosophy of physics. For suppose that just saying which regularities obtain of nomological necessity is not enough to fix the laws at a world—you have to also say something about the geometry of the set of worlds co-legal with the given world. Then Humean supervenience of the nomological on the non-nomological fails dramatically—two worlds can share the same Humean mosaic and the same nomological necessities without having the same laws. And in philosophy of physics, we should spend more time thinking about spaces of solutions—with the structures and group actions that they carry. Then, perhaps, it will seem less strange to think that, even in general relativity, isomorphic solutions can sometimes be capable of representing distinct possibilities.

Notes

¹ Or, rather, it was a great title when it was used for a greatest hits compilation back in 1959, when Elvis had sold only fifty million records.

² Those who wonder what all the fuss was about might want to consult Bangs (1977).

³ Some philosophers think that reflection on mathematical practice teaches that it is a mistake to speak of different but isomorphic structures in cases like this. Recommended antidote: Burgess (2015, Ch. 3).

⁴ Puzzles about what sort of choice, if any, God faced in choosing a time and location for the creation of the material universe were traditional already in Augustine's time (CD XI.4 ff.).

⁵ Think of the string as being plastic rather than elastic, in the sense that it has no tendency to return to a default form.

⁶ So, for instance, we expect that each time we repeat the procedure, we will end up with a different point playing the role of the endpoint of the first one-thousandth of the string.

⁷ Here and below: two metrics defined on the same manifold differ by a generalized shift if and only if they are related to one another by a diffeomorphism of the underlying manifold.

⁸ See LC, esp. C.II.1, L.III.3–7, C.III.2–6, L.IV.15–18, C.IV.15–19, and L.V.55–70.

⁹ An early form of this view can be found in Field (1985). But the floodgates opened downstream from Earman and Norton (1987). See, e.g., Butterfield (1989), Brighouse (1994), Healey (1995), Hofer (1996), Pooley (2006), and Baker (2010). For further discussion and references, see Pooley (2013, §7). The influence of this sort of view now extends beyond philosophy of physics—see, e.g., Chalmers (2014).

¹⁰ Clarke and Leibniz agreed that space was Euclidean, and went on to disagree about its ontology. In the same way, people can agree that spacetime is Newtonian in structure, but disagree as to its ontology—e.g., by disagreeing about whether spacetime and its parts are on a par ontologically with material events, or whether they should be thought of as an abstract framework for talking about geometric relations between material events. For variant ontologies of Newtonian spacetime, see Maudlin (1993).

¹¹ This is a variant on the buckling column example of Wilson (1993). Readers worried about the role that a privileged direction and unit of speed play can substitute fancier examples.

¹² Here and below, I unashamedly assume that determinism is a modal notion.

¹³ I am aware that this is painfully naive. Some of the pain will be addressed in §3.2 below.

¹⁴ There is room here to wonder whether one can coherently maintain that the spacetime of the swerve world has (just) Newtonian structure if one denies that shifts generate new possibilities. Compare with a case in which laws are posited giving one point of space modal properties that the other lack—is it really coherent to maintain that space has (just) the structure of Euclidean geometry?

¹⁵ See Lewis (1983; 1986, §4.4). For critical discussion, see, e.g., Fara (2009), Kment (2012), and Stalnaker (2012, Ch. 3).

¹⁶ Lewis (1983, 25). Roughly speaking, haecceitism can be characterized as the doctrine that “our identity is so separable from the way that we happen to be that which things would have been which had things been different is independent of which things would have had which properties” (Fara 2009, 286). There are many ways of making this idea precise—in part because the explication that looks most natural and suitable on one account of possible worlds may appear unnatural and unsuitable on another. On this point, see esp. Skow (2008a). Lewis himself thought of haecceitism in slightly different terms—as the thesis that qualitatively identical worlds can differ in how they represent *de re* some individual—and so he thought of his account not as a variety of haecceitism but as a form of anti-haecceitism that was nonetheless capable of simulating certain desirable features of haecceitism.

¹⁷ Because Lewis is a counterpart theorist, the natural way to make sense of the idea of worlds w_1 and w_2 differing in how they distribute roles over individuals is as follows: there is a possible individual x and a role R such that x has a counterpart at w_1 that occupies role R but x has no such counterpart at w_2 . But because Lewis is a qualitative counterpart theorist, this is impossible when w_1 and w_2 are qualitatively indiscernible: capacities of worlds for representation *de re* supervene on qualitative characters; see Lewis (1986, 221).

¹⁸ It is natural to think of being the occasion of a swerve as being an extrinsic property of an instant. But extrinsic qualitative similarities are often relevant to counterparthood—see Lewis (1986, esp. 231 and 245) on match of origins.

¹⁹ Lewis (1983, 397; 1986, 232). Here Lewis is speaking of a counterpart relation under which he has poor Fred there as one of his counterparts.

²⁰ Indeed, at one point Lewis seems to suggest that counterpart relations under which a thing has same-worldly counterparts other than itself will be admissible only in special contexts (1968/1983, Appendix C), elsewhere that they are admissible in some but not all contexts (1986, 232 fn. 22).

²¹ Indeed, cheap haecceitists who discuss shift arguments appear to be uninterested in the possibility of exploiting same-worldly counterpart relations in order to vindicate shiftiness—see, e.g., Butterfield (1989, 27), Brighouse (1994, 123 f.) and Pooley (2013, §7). Brighouse (1997, fn. 6) further denies that same-worldly counterparts are relevant to questions about determinism in physics.

²² Kment (2012, §5). For a related point of view, see Russell (2014, §2).

²³ Lewis (1968/1983, Postulate 5).

²⁴ For these and related examples, see Lewis (1983, §VI; 1986, §4.4).

²⁵ See Lewis (1983, 399; 1986 227 f. and 230). Here Lewis is following a suggestion of Hazen (1979, §III).

²⁶ For this claim, see, e.g., Sklar (1974, 161) and Earman (1989, 115). For another way of explicating it, see Dasgupta (2011, §1).

²⁷ There is a strong positive correlation between haecceitism and shiftiness. And whereas some see something like an emerging consensus among philosophers of physics around shiftlessness—Baker (2010, 1157) speaks of this view as being “widely held,” Greaves and Wallace (2014, 60) speak of “a widespread consensus” around it, Teh (2016, 98) of it as agreed upon by many but not all—metaphysicians appear to be something like evenly split on the question of haecceitism.

²⁸ Am I some kind of ratbag conventionalist about modality? Yes. But one does not have to be a conventionalist to make the move I make here—compare with the discussion of necessitarianism about laws in Langton and Lewis (1998, §V).

²⁹ Examples could be multiplied. Some others I have come across recently: Giulini (2009, §§2.3 and 3.5), Spradlin *et al.* (2002, §4), Tong (2005, §1.1.2), Witten (1999, Lecture 2).

³⁰ It is important to distinguish two questions. (1) Are two solutions *representationally equivalent*, in the sense that they are equally well-suited to represent any given physical situation? (2) Given a family of representationally equivalent solutions, how large is the family of physical situations that they are suited to represent?

Regarding question (1), it is important to note that two solutions can be related by a symmetry without being representationally equivalent—for discussion and examples, see Belot (2013). But in the cases to be discussed below, we will be interested in symmetries that relate isomorphic solutions. And I take it for granted that isomorphic solutions are always representationally equivalent (readers who don’t should consult Weatherall 2016).

³¹ The gauge/physical distinction is relative to an interpretation: one can always construct perverse interpretations under which symmetries typically considered to be paradigm examples of gauge symmetries come out as physical symmetries; and for relationalists even shifts will count as gauge symmetries.

³² Weatherall (2016) is difficult to classify. To the extent that in the final section of his paper, Weatherall assimilates his position to those of Butterfield (1989) and Brighouse (1994), he commits himself to shiftlessness concerning classical physics. But to the extent that in the opening section of the paper he seems to restrict attention to question (1) of fn. 30 above and to exclude question (2) from consideration, it would appear that he should be neutral between shiftiness and shiftlessness.

³³ For discussion and references, see Belot (2013).

³⁴ On this point, see Ismael and van Fraassen (2003), Healey (2009), and Dasgupta (2016).

³⁵ See, esp., Anderson (1964), Earman (1974), and Friedman (1973; 1983).

³⁶ For this approach, see, e.g., Brown and Sypel (1995), Budden (1997), Greaves and Wallace (2014), Healey (2009), and Saunders (2003). For critical discussion, see Skow (2008b).

³⁷ See Greaves and Wallace (2014) and Teh (2016).

³⁸ For discussion and references, see Struyve (2011, §6.5).

³⁹ See, e.g., Manton and Sutcliffe (2004, §§8.2 and 8.11), Shnir (2005, §6.5), Tong (2005, §2.1.3), Weinberg (2012, §§3.3 and 5.5), and Witten (1999, §9.5).

⁴⁰ See, e.g., Manton (1982) and Weinberg (1979).

⁴¹ Note the overlap between the authors cited in fnn. 39 and 40 above. Students are sometimes explicitly alerted to the existence of the shift in counting (see Tong 2005, §1.1.2).

Readers who have in mind the thoroughly unphysical nature of global phase in quantum mechanics may be left wondering what is going on here. Why do physicists count the global phase of a classical monopole as a physical characteristic? There appear to be two reasons. (i) The standard $4k$ -dimensional moduli space of monopoles carries an elegant hyper-Kähler geometric structure that is spoiled if one neglects global phase and moves to the $4k - 1$ -dimensional alternative (this point is stressed by Tong 2005, §1.1.2). (ii) Perturbing a monopole by a time-dependent transformation of its global phase endows it with electric charge just as perturbing a monopole by a time-dependent transformation of its position endows it with momentum (see Manton and Sutcliffe 2004, §8.11; or Weinberg 2012, §5.5).

⁴² Witten (1999, §§2.2 and 9.5). Conjecture: the underlying idea is that a symmetry counts as physical if and only if it is associated with a non-trivial conserved quantity—in the Yang–Mills–Higgs case, the conserved quantity associated with global phase transformations is trivial in the vacuum sector but nontrivial in the monopole sector.

⁴³ To keep things simple, henceforth we will restrict to globally hyperbolic solutions of the Einstein equations (no time travel or other funny business allowed), with vanishing cosmological constant (unless

otherwise noted), and with matter fields (if any) that supervene on geometry (such as a pressureless perfect fluid—see §§3.13 f. of Sachs and Wu 1977).

⁴⁴ Compare with the toy triangle example from above: there are several ways to endow a given set of three points with the geometry of a 3-4-5 right-angled triangle.

⁴⁵ That is: a generalized shift is a diffeomorphism from V to itself that is an isometry from (V, g_1) to (V, g_2) .

⁴⁶ The kernel of the argument is due to Einstein and played an important role in complicating his route to discovery of general relativity—for a brief discussion and references, see Janssen (2014, §3). Earman and Norton (1987) reconditioned the hole argument and launched it on a controversial career as an argument against substantivalism about the spacetime of general relativity—for a survey of the ensuing literature, see Earman (1989, Ch. 9) and Pooley (2013, §7). Here I focus on the uncontroversial core of the argument.

⁴⁷ Compare: two ways of building a 3-4-5 triangle out of three given points might agree that a certain point lies opposite the longest side, while disagreeing about the roles of the other two points.

⁴⁸ A possible exception: Paul Ehrenfest. In a letter to Einstein of 1 April 1917, de Sitter attributes to Ehrenfest a certain view that *might* be that generally covariant theories are indeterministic in a physically harmless fashion (CPAE 8.321): “the world as a whole can perform random motions without us (within the world) being able to observe it” (in this connection see also de Sitter 1917, 1222 f.). Unfortunately, Ehrenfest’s letters to Einstein of this period have not survived. In his discussion of the hole argument in a letter to Ehrenfest of 29 December 1915 (CPAE 8.174) Einstein makes a tantalizing allusions to randomness; see also his letter to Ehrenfest of 5 January 1916 (CPAE 8.180). For some relevant context, see Kox (1988, §3).

⁴⁹ An exception: Brighouse (2008).

⁵⁰ See, e.g., Butterfield (1989), Brighouse (1994), and Hofer (1996). For further discussion and references, see Pooley (2013, §7).

⁵¹ See Maudlin (1989) and Bartels (1996). These approaches differ concerning what sort of geometric properties points possess essentially: under Maudlin’s approach, it is the geometric relations each point bears to all others; under Bartels’s approach, each point has certain intrinsic properties (corresponding to the sort of geometric invariants that suffice to build up the metric tensor). There is good reason to prefer the latter approach. By ordinary standards, general relativity *is* indeterministic: there are (very special) possible ways a world can start out that admit geometrically different continuations (see, e.g., Chruściel and Isenberg 1993). But this sort of indeterminism is effaced on Maudlin’s approach: if it is impossible for the same point to exist in two worlds with different overall geometries, then in particular it is impossible for two worlds with different overall geometries to share an initial segment.

⁵² See, e.g., Earman (1989, Ch. 9) and Dasgupta, (2011).

⁵³ The Gospel According to Bob: Wald (1984, 438).

⁵⁴ For details, discussion, and references, see Belot (2008). Note however, that this basic approach requires some fine-tuning. As it stands, it would lead one to expect that physicists would view the Yang–Mills–Higgs theory as having a number of distinct vacuum states related to one another by global phase transformations. But they do not: see, e.g., Englert (2005, §§1 and 3.1), ’t Hooft (2007, §5.3), and Witten (1999, §2.2). This is in contrast with the way that global phase transformations are treated in the monopole sector of the same theory. So there is a puzzle here. See the discussion of fn. 42 above.

⁵⁵ NB: the swerve theory is not of this form. And if we feed the swerve theory into the simplified machine described above, the we get the fairly bizarre ruling that temporal shifts count as gauge while spatial shifts count as physical. Best to just set this theory aside in the present context.

⁵⁶ Roughly speaking: in a field theory with a degenerate Lagrangian, the family of gauge symmetries is so large that it is parameterized by arbitrary functions on spacetime. See Olver (1993, 334 ff.).

⁵⁷ Exception: if the original drawing features a vertical line, then of course merely shifting the glass in the vertical direction will not result in a new drawing.

⁵⁸ In theories like Maxwell’s theory and its generalizations, the the story is interestingly different.

⁵⁹ On this point, see Ringström (2009, §17.2).

⁶⁰ See, e.g., Komar (1958) and Bergmann (1961).

⁶¹ See, e.g., Crnković and Witten (1987) for an analysis in the covariant phase space framework; Beig (1994) for a treatment within (a modern version of) Dirac's theory of constrained Hamiltonian systems; and Belot (2008, §3.1) for an elementary treatment.

⁶² A widely used scheme involves slicing solutions into instants of constant mean curvature. Let Σ be an instant of time (a Cauchy surface) in a solution (V, g) . We define a real-valued function K on Σ as follows: choose a family of observers that are at relative rest as they pass through Σ ; for $x \in \Sigma$, the *mean curvature* at x is the number $K(x)$ that measures the rate at which the observer in our family that passes through x sees space expanding as she passes through x (the result is independent of the family of observers chosen). We call Σ a *constant mean curvature (CMC) slice* if K is a constant function on Σ . Many but not all vacuum space times can be foliated in a unique way by CMC slices. For references to the large literature on CMC slicing, see, e.g., Andersson (2004, §2.3).

⁶³ Warning: this is only one of several generic notions of asymptotic flatness in general relativity, each with several species falling under it. An important alternative notion is asymptotic flatness at null infinity, standardly characterized in terms of Penrose's scheme of conformal compactification (see, e.g., Wald 1984, Ch. 11). Below it is argued that if asymptotic flatness is imposed at spatial infinity, then there are good reasons to treat time translations at spatial infinity as physical rather than gauge symmetries. There are similarly good reasons to treat BMS supertranslations at null infinity as physical rather than gauge if asymptotic flatness is imposed at null infinity: consider, e.g., the roles that these symmetries play in Hawking *et al.* (2016) and Hollands *et al.* (2016).

⁶⁴ For a treatment in this style, see Christodoulou (2008, Ch. 3). For other approaches, see, e.g., Andersson (1987) and Ashtekar *et al.* (1991).

⁶⁵ This formulation glosses over a number of important subtleties—as noted above, there are in fact a number of specific notions of asymptotic flatness at spatial infinity, differing, e.g., as to the precise sense in (and rate at) which the physical metric is required to approach the Minkowski metric in the neighbourhood of spatial infinity. For a helpful discussion, see Ashtekar *et al.* (2008).

⁶⁶ For a covariant phase space treatment, see Ashtekar *et al.* (1991). For treatments in the tradition of Dirac's theory of constrained Hamiltonian systems, see Beig and Ó Murchadha (1987) or Compère and Dehouck (2011). For an elementary treatment, see Belot (2008, §3.1). Warning: there is some variation in the precise boundary conditions employed by these authors.

⁶⁷ In part because of differences between the framework of covariant phase spaces and Dirac's framework of constrained Hamiltonian systems, in part because different authors deploy different boundary conditions. Subtleties of the sort alluded to in fn. 54 above also play a role, especially in the treatment of supertranslations at spatial infinity within Dirac-style approaches.

⁶⁸ Slicing by maximal surfaces—a variant on the CMC slicing scheme described in fn. 62 above. See Christodoulou (2008, §3.2.1).

⁶⁹ This is true in all of the approaches mentioned in fn. 66 above.

⁷⁰ CPAE 6.43, §2. In his reply to Kretchmann of 1918, Einstein explicitly identifies the principle of relativity with what we would call general covariance (CPAE 6.4). In a letter to Mie of 1918, Einstein remarks that locally generally covariant theory in which isolated systems correspond to asymptotically flat solutions “is a monstrosity” (CPAE 8, 470).

⁷¹ For a more intrinsic way of making this picture precise, see Andersson (1987, §2.2).

⁷² Here is a way to get the general idea. Recall from above that there is a geometrically natural way to partition any asymptotically flat (V, g) into instants of time asymptotic to a family of flat, parallel hypersurfaces of the background Minkowski geometry. This recipe has a unique output (except in the special case when (V, g) has the flat geometry of Minkowski spacetime). So a generalized shift relating asymptotically flat solutions must map the family of flat hyper-surfaces associated with the first metric to the family of flat hypersurfaces associated with the second metric—so it must basically determine a Lorentz symmetry of the background geometry. Further, it can make sense to say that a generalized shift induces a time translation at infinity as well: suppose that two isometric metrics are associated with the same family of flat hypersurfaces; then if one of these metrics includes an instant of time with a certain geometry asymptotic to a given flat hypersurface of the background metric, then the other will include an instant with the same geometry asymptotic to a flat hypersurface of the background metric that differs from the given one by a—typically, nontrivial—time-translation. A somewhat different story is required for spatial translations.

⁷³ Wald (1984, 467 fn. 2). See also, e.g., Wald and Zoupas (2000).

⁷⁴ This is, presumably, the sort of thing that Penrose had in mind in the passage quoted above.

⁷⁵ This is just a restatement of the usual prescription: classical first-class constraints are enforced at the quantum level by requiring that the operators corresponding to them annihilate physical states.

⁷⁶ On this point, see Friedman and Sorkin (1982, 617 f.).

⁷⁷ See, e.g., Hawking and Ellis (1972, §7.6); or Ringström (2009, Ch. 15).

⁷⁸ This monumental result is due to Christodoulou and Klainerman. For a brief treatment, see Christodoulou (2008, Ch. 4).

⁷⁹ The classic discussion of stability considerations in general relativity is Hawking (1970). For further discussion and references, see Fletcher (2016). For the importance of stability considerations in underwriting physical explanations in general, see Batterman (2002).

⁸⁰ One could instead consider the union over all possible spatial topologies of all such spaces of initial data. But the result would be rather trivial: for the topologies relevant here, the only way that a sequence of instantaneous states defined on a sequence of manifolds Σ_i will converge to an instantaneous state defined on a manifold Σ , is if all but finitely many of the Σ_i have the same topology as Σ ; see Petersen (2006, §10.3.2).

⁸¹ For this theme, see, e.g., Andersson (1987), Ashtekar *et al.* (1991), Bartnik (2005), Choquet-Bruhat *et al.* (1979), and Regge and Teitelboim (1974).

⁸² On this point, see Ringström (2009, §17.2).

⁸³ What follows amounts to two thoughts. Einstein was right: to impose asymptotic boundary conditions on a theory like general relativity is to violate the principle of general covariance. And: it is better to be honest rather than sneaky about that.

⁸⁴ Here, general covariance means: acting on a solution by a spacetime diffeomorphism yields a solution.

⁸⁵ For a serviceable if imperfect characterization of the notion of background structure, see Belot (2011).

⁸⁶ In fact, there are several competing characterizations of the asymptotically de Sitter sector and its asymptotic symmetries. But they each embody the principle that generalized shifts asymptotic to the identity at infinity should be treated as gauge symmetries, and others as physical symmetries. See, e.g., Annin *et al.* (2011), Ashtekar *et al.* (2015), and Marolf and Kelly (2012).

⁸⁷ See, e.g., Wald (1983) and Ringström (2014).

⁸⁸ See, e.g., Belot (2005; 2006; 2014) on the idea that some interpretative questions about physical theories concern symmetries of the space of worlds of the theory; Halvorson (2012; 2016) on the need to think of the content of a theory as involving maps defined on the set of worlds of the theory, rather than being identified with that set; North (2011) on the structure of classical mechanics; and Stalnaker (2012, esp. §§3.6 and 3.9) on reasons to think that a possible world cannot be fully specified without giving some information about its relation to other possible worlds.

⁸⁹ I suspect that everything said below about this nonstandard approach could be taken over by primitivists or necessitarians about laws.

⁹⁰ See, e.g., Marmo and Saletan (1977), Hojman and Montemayor (1980), and Henneaux and Shepley (1982).

Abbreviations

CD. Augustine, *De Civitate Dei*. Cited by chapter and paragraph number. Available in H. Bettenson (trans.), *City of God*. (London: Penguin Classics, 2003.)

CPAE. *The Collected Papers of Albert Einstein*. Cited by volume and document number. (Princeton: Princeton University Press. 1987–.)

LC. The Leibniz-Clarke Correspondence. Cited by author's initial, letter number, and paragraph number. Available in H.G. Alexander (ed.), *The Leibniz-Clarke Correspondence* (Manchester: University of Manchester Press, 1956.)

References

- Anderson, J. (1964) "Relativity Principles and the Role of Coordinates in Physics." In H.Y. Chiu and W. Hoffman (eds), *Gravitation and Relativity*. New York: W.A. Benjamin, Inc., 175–194.
- Andersson, L. (1987) "Momenta and Reduction for General Relativity." *Journal of Geometry and Physics* 4: 289–314.
- . (2004) "The Global Existence Problem in General Relativity." In P. Chruściel and H. Friedrich (eds), *The Einstein Equations and the Large Scale Behavior of Gravitational Fields*. Basel: Birkhäuser, 71–120.
- Anninos, D., G. Ng, and A. Strominger (2011) "Asymptotic Symmetries and Charges in de Sitter Space." *Classical and Quantum Gravity* 28: 175019.
- Ashtekar, A., L. Bombelli, and O. Reula (1991) "The Covariant Phase Space of Asymptotically Flat Gravitational Fields." In M. Francaviglia (ed), *Mechanics, Analysis and Geometry: 200 Years after Lagrange*. Amsterdam: North–Holland, 417–450.
- Ashtekar, A., B. Bonga, and A. Kesavan (2015) "Asymptotics with a Positive Cosmological Constant: I. Basic Framework." *Classical and Quantum Gravity* 32: 025004.
- Ashtekar, A., J. Engle, and D. Sloan (2008) "Asymptotics and Hamiltonians in a First-Order Formalism." *Classical Quantum Gravity* 25: 095020.
- Baker, D. (2010) "Symmetry and the Metaphysics of Physics." *Philosophy Compass* 5: 1157–1166.
- Bangs, L. (1977) "How Long Will We Care?" *Village Voice* XXII no. 35: 1, 16–17.
- Bartels, A. (1996) "Modern Essentialism and the Problem of the Individuation of Spacetime Points." *Erkenntnis* 45: 25–43.
- Bartnik, R. (2005) "Phase Space for the Einstein Equations." *Communications in Analysis and Geometry* 13: 845–885.
- Batterman, R. (2002) *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford: Oxford University Press.
- Beig, R. (1994) "The Classical Theory of Canonical General Relativity." In J. Ehlers and H. Friedrich (eds), *Canonical Gravity: from Classical to Quantum*. Berlin: Springer–Verlag, 59–80.
- Beig, R. and N. Ó Murchadha (1987) "The Poincaré Group as the Symmetry Group of Canonical General Relativity." *Annals of Physics* 174: 463–498.
- Belot, G. (2005) "Dust, Time, and Symmetry." *British Journal for the Philosophy of Science* 56: 255–291.
- . (2006) "The Representation of Time and Change in Mechanics." In J. Butterfield and J. Earman (eds), *Philosophy of Physics Part A*. Amsterdam: North–Holland, 133–227.
- . (2008) "An Elementary Notion of Gauge Equivalence." *General Relativity and Gravitation* 43: 2865–2884.
- . (2011) "Background–Independence." *General Relativity and Gravitation* 42: 199–215.
- . (2013) "Symmetry and Equivalence." In R. Batterman (ed), *The Oxford Handbook of Philosophy of Physics*. Oxford: Oxford University Press, 318–339.
- . (2014) "Does Time Flow Equally? Is it the Same at Rome and at London?" Unpublished.
- Bergmann, P. (1961) "Observables in General Relativity." *Reviews of Modern Physics* 33: 510–514.
- Brighouse, C. (1994) "Spacetime and Holes." In R. Burian, M. Forbes, and D. Hull (eds), *PSA 1994*, vol. 1. East Lansing MI: Philosophy of Science Association, 117–125.
- . (1997) "Determinism and Modality." *British Journal for the Philosophy of Science* 48: 465–481.
- . (2008) "Understanding Indeterminism." In D. Dieks (ed), *The Ontology of Spacetime II*. Amsterdam: Elsevier, 153–173.
- Brown, H. and R. Sypel (1995) "On the Meaning of the Relativity Principle and other Symmetries." *International Studies in the Philosophy of Science* 9: 235–253.
- Budden, T. (1997) "Galileo's Ship and Spacetime Symmetry." *British Journal for the Philosophy of Science* 48: 483–516.
- Burgess, J. (2015) *Rigor and Structure*. Oxford: Oxford University Press.
- Butterfield, J. (1989) "The Hole Truth." *British Journal for the Philosophy of Science* 40: 1–28.
- Chalmers, D. (2014) "Three Puzzles About Spatial Experience." Forthcoming in A. Pautz and D. Stoljar (eds), *Themes from Ned Block*. Cambridge MA: MIT Press.

- Choquet-Bruhat, Y., A. Fischer, and J. Marsden (1979) "Maximal Hypersurfaces and Positivity of Mass." In J. Ehlers (ed), *Isolated Gravitating Systems in General Relativity*. Amsterdam: North-Holland, 396–456.
- Christodoulou, D. (2008) *Mathematical Problems of General Relativity*, vol. I. Zurich: European Mathematical Society.
- Chruściel, P. and J. Isenberg (1993) "Nonisometric Vacuum Extensions of Vacuum Maximally Globally Hyperbolic Spacetimes." *Physical Review D* 48: 1616–1628.
- Compère, G. and F. Dehouck (2011) "Relaxing the Parity Conditions of Asymptotically Flat Gravity." *Classical and Quantum Gravity* 28: 245016.
- Crnković, Č. and E. Witten (1987) "Covariant Description of Canonical Formalism in Geometrical Theories." In S. Hawking and W. Israel (eds), *Three Hundred Years of Gravitation*. Cambridge: Cambridge University Press, 676–684.
- Dasgupta, S. (2011) "The Bare Necessities." *Philosophical Perspectives* 25: 115–160.
- . (2016) "Symmetry as an Epistemic Notion (Twice Over)" *British Journal for the Philosophy of Science* 67: 837–878.
- de Sitter, W. (1917) "On the Relativity of Inertia. Remarks on Einstein's Latest Hypothesis." *Koninklijke Akademie van Wetenschappen te Amsterdam. Proceedings of the Section of Sciences* 19: 1217–1225.
- Earman, J. (1974) "Covariance, Invariance, and the Equivalence of Frames." *Foundations of Physics* 4: 267–289.
- . (1989) *World Enough and Space-Time*. Cambridge MA: MIT Press.
- Earman, J. and J. Norton (1987) "What Price Substantivalism? The Hole Story." *British Journal for the Philosophy of Science* 38: 515–525.
- Eddington, A. (1930) "On the Instability of Einstein's Spherical World." *Monthly Notices of the Royal Astronomical Society* 90: 668–678.
- Englert, F. (2005) "Broken Symmetry and Yang–Mills Theory." In G. 't Hooft (ed), *50 Years of Yang–Mills Theory*. Singapore: World Scientific, 65–95.
- Fara, D. (2009) "Dear Haecceitism." *Erkenntnis* 70: 288–297.
- Field, H. (1985) "Can We Dispense with Spacetime?" In P. Asquith and P. Kitcher (eds), *PSA 1984*, vol. 2. East Lansing MI: Philosophy of Science Association, 33–90.
- Fletcher, S. (2016) "Similarity, Topology, and Physical Significance in Relativity Theory." *British Journal for the Philosophy of Science* 67: 365–389.
- Friedman, J. and R. Sorkin (1982) "Half-Integral Spin from Quantum Gravity." *General Relativity and Gravitation* 14: 615–620.
- Friedman, M. (1973) "Relativity Principles, Absolute Objects, and Symmetry Groups." In P. Suppes (ed) *Space, Time, and Geometry*. Dordrecht: Reidel, 296–320.
- . (1983) *Foundations of Spacetime Theories*. Princeton: Princeton University Press.
- Giulini, D. (2009) "Concepts of Symmetry in the Work of Wolfgang Pauli." In H. Atmanspacher and H. Primas (eds), *Recasting Reality: Wolfgang Pauli's Philosophical Ideas and Contemporary Science*. Berlin: Springer-Verlag, 33–82.
- Greaves, H. and D. Wallace (2014) "Empirical Consequences of Symmetry." *British Journal for the Philosophy of Science* 65: 59–89.
- Halvorson, H. (2012) "What Scientific Theories Could Not Be." *Philosophy of Science* 79: 183–206.
- . (2016) "Scientific Theories." In P. Humphreys (ed), *The Oxford Handbook of Philosophy of Science*. Oxford: Oxford University Press, 585–608.
- Hawking, S. (1970) "Stable and Generic Properties in General Relativity." *General Relativity and Gravitation* 1: 393–400.
- Hawking, S. and G. Ellis (1973) *The Large-Scale Structure of Space-Time*. Cambridge: Cambridge University Press.
- Hawking, S., M. Perry, and A. Strominger (2016) "Soft Hair on Black Holes." *Physical Review Letters* 116: 231301.
- Hazen, A. (1979) "Counterpart-Theoretic Semantics for Modal Logic." *The Journal of Philosophy* LXXVI: 319–338.
- Healey, R. (1995) "Substance, Modality, and Spacetime." *Erkenntnis* 42: 287–316.

- . (2009) “Perfect Symmetries.” *British Journal for the Philosophy of Science* 60: 697–720.
- Henneaux, M. and L. Shepley (1982) “Lagrangians for Spherically Symmetric Potentials.” *Journal of Mathematical Physics* 23: 2101–2107.
- Hofer, C. (1996) “The Metaphysics of Space-Time Substantivalism.” *The Journal of Philosophy* XCIII: 5–27.
- Hojman, S. and R. Montemayor (1980) “S-Equivalent Lagrangians for Free Particles and Canonical Quantization.” *Hadronic Journal* 3: 1644–1657.
- Hollands, S., A. Ishibashi, and R. Wald (2016) “BMS Supertranslations and Memory in Four and Higher Dimensions.” Unpublished. arXiv:1612.03290
- Ismael, J. and B. van Fraassen (2003) “Symmetry as a Guide to Superfluous Theoretical Structure.” In K. Brading and E. Castellani (eds), *Symmetries in Physics: Philosophical Reflections*. Cambridge: Cambridge University Press, 371–392.
- Janssen, M. (2014) “No Success Like Failure . . .’: Einstein’s Quest for General Relativity, 1907–1920.” In M. Janssen and C. Lehner (eds), *The Cambridge Companion to Einstein*. Cambridge: Cambridge University Press, 167–227.
- Kelly, W. and D. Marolf (2012) “Phase Spaces for Asymptotically de Sitter Cosmologies.” *Classical and Quantum Gravity* 29: 205013.
- Kment, B. (2012) “Haeceitism, Chance, and Counterfactuals.” *The Philosophical Review* 121: 573–609.
- Komar, A. (1958) “Construction of a Complete Set of Independent Observables in the General Theory of Relativity.” *Physical Review* 111: 1182–1187.
- Kox, A.J. (1988) “Hendrik Antoon Lorentz, the Ether, and the General Theory of Relativity.” *Archive for History of Exact Sciences* 38: 67–78.
- Langton, R. and D. Lewis (1998) “Defining ‘Intrinsic.’” *Philosophy and Phenomenological Research* 58: 333–345.
- Lewis, D. (1968) “Counterpart Theory and Quantified Modal Logic.” *The Journal of Philosophy* 65: 113–126. Reprinted in (with additional Postscripts) in D. Lewis (1983) *Philosophical Papers*, vol. 1. Oxford: Oxford University Press, 26–46.
- . (1983) “Individuation by Acquaintance and by Stipulation.” *The Philosophical Review* 92: 3–32.
- . (1986) *On the Plurality of Worlds*. Oxford: Blackwell.
- Manton, N. (1982) “A Remark on the Scattering of BPS Monopoles.” *Physics Letters B* 110: 54–56.
- Manton, N. and P. Sutcliffe (2004) *Topological Solitons*. Cambridge: Cambridge University Press.
- Marmo, G. and E. Saletan (1977) “Ambiguities in the Lagrangian and Hamiltonian Formalism: Transformation Properties.” *Il Nuovo Cimento B* 40: 67–89.
- Maudlin, T. (1989) “The Essence of Spacetime.” In A. Fine and J. Leplin (eds), *PSA 1988*, vol. 2. East Lansing MI: Philosophy of Science Association, 82–91.
- Maudlin, T. (1993) “Buckets of Water and Waves of Space.” *Philosophy of Science* 60: 183–203.
- North, J. (2011) “Classical Mechanics, Structure, and the World.” Unpublished.
- Olver, P. (1993) *Applications of Lie Groups to Differential Equations* (2nd ed). Berlin: Springer-Verlag.
- Penrose, R. (1982) “Some Unsolved Problems in Classical General Relativity.” In S.-T. Yau (ed), *Seminar on Differential Geometry*. Princeton: Princeton University Press, 631–668.
- Petersen, P. (2006) *Riemannian Geometry* (2nd ed). Berlin: Springer-Verlag.
- Pooley, O. (2006) “Points, Particles, and Structural Realism.” In D. Rickles, S. French, and J. Saatsi (eds), *The Structural Foundations of Quantum Gravity*. Oxford: Oxford University Press, 83–120.
- . (2013) “Substantivalist and Relationalist Approaches to Spacetime.” In R. Batterman (ed), *The Oxford Handbook of Philosophy of Physics*. Oxford: Oxford University Press, 522–586.
- Regge, T. and C. Teitelboim (1974) “Role of Surface Integrals in the Hamiltonian Formulation of General Relativity.” *Annals of Physics* 88: 286–318.
- Ringström, H. (2009) *The Cauchy Problem in General Relativity*. Zurich: European Mathematical Society.
- . (2014) “On the Future Stability of Cosmological Solutions to Einstein’s Equations with Accelerated Expansion.” In *Proceedings of the International Congress of Mathematicians. Seoul 2014*, volume II. Seoul: Kyung Moon SA, 983–999.
- Russell, J. (2014) “On Where Things Could Be.” *Philosophy of Science* 81: 60–80.
- Sachs, R. and H. Wu (1977) *General Relativity for Mathematicians*. Berlin: Springer-Verlag.

- Saunders, S. (2003) "Physics and Leibniz's Principles." In K. Brading and E. Castellani (eds), *Symmetries in Physics: Philosophical Reflections*. Cambridge: Cambridge University Press, 289–308.
- Seiberg, N. (2007) "Emergent Spacetime." In D. Gross, M. Henneaux, and A. Sevrin (eds), *The Quantum Structure of Space and Time*. Singapore: World Scientific, 163–178.
- Shnir, Y. (2005) *Magnetic Monopoles*. Berlin: Springer-Verlag.
- Sklar, L. (1974) *Space, Time, and Spacetime*. Berkeley: University of California Press.
- Skow, B. (2008a) "Haecceitism, Anti-Haecceitism, and Possible Worlds." *The Philosophical Quarterly* 58: 98–107.
- . (2008b) "Local and Global Relativity Principles." *Philosophers' Imprint* 8: 1–14.
- Spradlin, M., A. Strominger, and A. Volovich (2002) "De Sitter Space." In C. Bachas, A. Bilal, M. Douglas, N. Nekrasov, and F. David (eds), *Unity from Duality: Gravity, Gauge Theory, and Strings*. Berlin: Springer-Verlag, 423–453.
- Stalnaker, R. (2012) *Mere Possibilities: Metaphysical Foundations of Modal Semantics*. Princeton: Princeton University Press.
- Struyve, W. (2011) "Gauge Invariant Accounts of the Higgs Mechanism." *Studies in History and Philosophy of Modern Physics* 42: 226–236.
- Teh, N. (2016) "Galileo's Gauge." *Philosophy of Science* 83: 93–118.
- 't Hooft, G. (2007) "The Conceptual Basis of Quantum Field Theory." In J. Butterfield and J. Earman (eds), *Philosophy of Physics, Part A*. Amsterdam: North-Holland, 661–729.
- Tong, D. (2005) "TASI Lectures on Solitons." Unpublished. arXiv:hep-th/0509216v5
- Wald, R. (1983) "Asymptotic Behavior of Homogeneous Cosmological Models in the Presence of a Positive Cosmological Constant." *Physical Review D* 28: 2118–2120.
- . (1984) *General Relativity*. Chicago: University of Chicago Press.
- Wald, R. and A. Zoupas (2000) "General Definition of 'Conserved Quantities' in General Relativity and Other Theories of Gravity." *Physical Review D* 61: 084027.
- Weatherall, J. (2016) "Regarding the 'Hole Argument'." Forthcoming in *British Journal for the Philosophy of Science*.
- Weinberg, E. (1979) "Parameter Counting for Multimonoopole Solutions." *Physical Review D* 20: 936–944.
- . (2012) *Classical Solutions in Quantum Field Theory*. Cambridge: Cambridge University Press.
- Wilson, M. (1993) "There's a Hole and a Bucket, Dear Leibniz." *Midwest Studies in Philosophy* 18: 202–241.
- Witten, E. (1999) "Dynamics of Quantum Field Theory." In P. Deligne, P. Etingof, D. Freed, L. Jeffrey, D. Kazhdan, J. Morgan, D. Morrison, and E. Witten (eds), *Quantum Fields and Strings: A Course for Mathematicians*, vol. 2. Providence RI: American Mathematical Society, 1119–1409.
- Zee, A. (2010) *Quantum Field Theory in a Nutshell* (2nd ed). Princeton: Princeton University Press.