

General Relativity Needs No Interpretation*

Erik Curiel^{†‡}

I argue that, contrary to the recent claims of physicists and philosophers of physics, general relativity requires no interpretation in any substantive sense of the term. I canvass the common reasons given in favor of the alleged need for an interpretation, including the difficulty in coming to grips with the physical significance of diffeomorphism invariance and of singular structure, and the problems faced in the search for a theory of quantum gravity. I find that none of them shows any defect in our comprehension of general relativity as a physical theory. I conclude by comparing general relativity with quantum mechanics, a theory that manifestly does stand in need of an interpretation in an important sense. Although many aspects of the conceptual structure of general relativity remain poorly understood, it suffers no incoherence in its formulation as a physical theory that only an ‘interpretation’ could resolve.

When science starts to be interpretive it is more unscientific even than mysticism. (D. H. Lawrence, “Self-Protection”)

1. Introduction. Several philosophers, mathematicians and physicists have in recent years concluded that general relativity requires an interpretation.¹ They arrive at seemingly the same destination having set out from diverse and varied starting points, motivated by problems ranging from the attempt to find a quantum theory of gravitation to the attempt to come to grips with the existence of generic singular structure in solutions to the Einstein

*Received November 2007; revised February 2009.

[†]To contact the author, please write to: Center for Philosophy of Science, University of Pittsburgh, 817 Cathedral of Learning, Pittsburgh, PA 15260; e-mail: erik@strangebeautiful.com.

[‡]I thank John Norton for penetrating remarks and hard questions, Jeremy Butterfield for insightful comments on an earlier version of the article, Chris Pincock for a useful conversation on interpreting physical theories, and David Malament for many helpful discussions over several years on essentially every issue I discuss in the article and especially for pressing me to clarify my notion of ‘interpretation’.

1. See, e.g., Belot (1996) and Rovelli (2000).

Philosophy of Science, 76 (January 2009) pp. 44–72. 0031-8248/2009/7601-0005\$10.00
Copyright 2009 by the Philosophy of Science Association. All rights reserved.

field-equation. I argue in this article that general relativity does not require such a thing, at least not in any substantive, interesting sense.

The argument is in three parts. In the first, in order to make the discussion concrete, I propose in Section 2 a weak, necessary condition for the conclusion that a scientific theory stands in need of an interpretation in a physically significant sense that is at the same time not philosophically trivial. Of course, to conclude from this that general relativity stands in need of no interpretation has only so much interest and force as the necessary condition I propose. The one I do propose seems to me to have the virtues of weakness, clarity and manifest physical significance.

To the best of my knowledge, those who conclude that general relativity requires interpretation found their arguments on issues about or results in one or more of a fixed number of areas of active research. In the second part of the argument, I bring to bear the condition posited in Section 2 to address those issues in turn: the invariance of the theory under the group of spacetime diffeomorphisms (Section 3); the generic prevalence of singular structure in solutions to the Einstein field-equation (Section 4); and the search for a theory of quantum gravity (Section 5). In each case I conclude that, whatever else one may think of the importance of such problems and issues in their own right, none of them forces on us the need for a new theoretical structure in the terms of which we must interpret general relativity if we are to comprehend it. This should not come as a surprise, I think, for none of those problems and issues requires for their formulation concepts or terms not already available and well understood in general relativity itself. Along the way, I point out analogous ‘problems’ in other, nonquantum theories that *prima facie* do not point to a gap in our comprehension of those theories that only an interpretation of some sort could bridge. This suggests that the drive to find an interpretation for general relativity based on analogous problems may have nothing to do with the content of general relativity as a physical theory, but perhaps more to do with unresolved, purely psychological dissonances its conceptual structures leave us with.

In the third part of the argument, Section 6, I compare the situation for general relativity with that for quantum mechanics, a theory almost everyone agrees requires interpretive elucidation of some sort. I focus on the problem of measurement. I argue that it embodies one of quantum mechanics’s failures to satisfy the weak condition I propose in Section 2 for a theory not to require an interpretation; in consequence, quantum mechanics does stand in need of an interpretation in that sense.

2. Interpretation of a Physical Theory. The question we consider ought not be whether general relativity requires an ‘interpretation’ simpliciter—until we agree on what one means by the term, a disagreement on the

matter would be only an argument about what words we feel most comfortable using to describe the questions the theory presents us with. *De gustibus non disputandum est*. The interesting question is rather what one may mean by claiming that general relativity stands in need of one, and whether in any of the interesting senses it does indeed stand lacking.

It will be useful to distinguish three senses of ‘interpretation’ possibly relevant to my arguments. Because interpretation is, broadly speaking, a semantical concept, it will be convenient to speak in semantical terms. Fix a scientific framework, which is to say, something like a formalization of a scientific theory in conjunction with the body of empirical, experimentally derived and substantiated knowledge the formalized theory is supposed to represent. An interpretation of the framework, no matter what else it may be, in so far as it will say something about the way the formal apparatus of a theory hooks up to the world, will pertain to the semantic relations between the elements of the framework’s theoretical formalism on the one hand and the content of its empirical knowledge on the other. There are, broadly speaking, three ways it may do so: either as a fixation, an expansion or an explication of those semantic relations. Thus, one can give (elements of) the framework an interpretation in one of the three following senses:²

Concrete. The fixation of a semantics for the formalism, in the sense that the formalism under the semantics expresses the empirical knowledge the framework contains—for example, the fixation of a Tarskian family of models, or, less formally, the contents of a good, comprehensive text-book.

Categorial. The explication of concepts in the theory that the semantics of a concrete interpretation alone does not fix—for example, a demonstration that the theory is deterministic in any of a variety of senses.

Metalinguistic. The explication of the semantics of a concrete interpretation, when the representational nature of the concrete interpretation is itself not understood—for example, the Copenhagen Interpretation of quantum mechanics.

Now, all scientific frameworks have concrete interpretations, or else we would not even think of them as embodying *physical* theories in the first place, so those who argue that general relativity requires an interpretation cannot mean it needs a concrete one—we know general relativity has

2. I think it would be of interest to work out these different senses of interpretation with some rigor and precision. This article is not the appropriate forum for that job, however, and in any event the rough and ready characterization I give them here suffices for the task at hand.

sound concrete interpretations.³ Because *all* physical theories have known interpretative problems of the categorial sort, neither can the proponents of the need for an interpretation for general relativity mean that the alleged problem they are pointing to is a categorial one; otherwise the claim is either trivial or historically uninformed, in so far as workers in the field have been engaged in interpretative projects of a categorial sort since long before the cry was heard that an interpretation for general relativity was needed (e.g., Earman's [1986, 170–198] investigation of determinism in the context of relativistic spacetimes).

The only interesting sense of interpretation left is of the metalinguistic sort, and I think that it or something close to it must be what the proponents of the need for an interpretation have in mind. Let us try to characterize the idea a little more clearly, enough at least for the purpose at hand. To reformulate the idea slightly, a scientific framework is, roughly speaking, a (partly) formal theoretical structure that provides semantical resources rich enough for the construction of appropriate schematic representations of the types of physical system the theory at issue treats. In order to have the resources for the construction of schematic representations of physical systems, it must already have a concrete interpretation of some form, even if only a primitive and ad hoc one. The need for a metalinguistic interpretation arises when there is a deep puzzle about the way that significance accrues to the formal elements of the framework from its concrete interpretation—when we know what the elements' significance is, in a crude, operational way, but we have no understanding of the nature of the way the elements of the framework actually do represent the physical systems the theory treats. The way that Hermitian operators in standard quantum mechanics represent observables is perhaps the canonical example of such a problem: we know they do in some way or other represent observables, and we know how to use them to construct good models of systems that we can use to predict the (probabilistic) outcomes of experiments, but we have no clear understanding at all of the nature of the representational relations between, on the one hand, the operator as part of the formalism and, on the other, the actual values we measure for physical quantities in experiments. This is a deep and compelling sense in which a theory can stand in need of an interpretation, and it is a sense, moreover, that does not *prima facie* hold for every physical theory automatically. It would be a deep result about gen-

3. See, e.g., those of Synge (1960), Misner, Thorne, and Wheeler (1973), and Wald (1984). One can fairly argue over the virtues and demerits of each with respect to depth, rigor and thoroughness, and with respect to a whole set of particular philosophical problems and issues, but it would be churlish at best and ignorant at worst to deny that they do provide concrete interpretations of the theory.

eral relativity were it shown that we stand in need of an interpretation for it in this sense.⁴

Let us say, then, that we have a framework some aspect of the semantic structure of which is, for one reason or another, poorly understood. One natural route of attack in the attempt to grasp it better is to try to find another framework that we do understand, in the terms of which we can construct a representation of the poorly understood part of the first. In order to be useful, the representation must in some important sense faithfully recapitulate in the second framework the poorly comprehended structure of the first—it must preserve the structure. This idea, the ‘preservation of structure’, plays the crucial role. Klein’s model of Lobachevskian geometry in the Euclidean plane provides a wonderful example of its significance,⁵ as does Helmholtz’s (1870) depiction, in the terms of our own naive language of sense-perceptions, of the possible perceptual experiences of a person residing in a space of constant curvature, and Peirce’s ([1898] 1993, 252–253) depiction of the possible perceptual experiences of a person traveling through a multiply connected space. In each case, the intrinsic structure of an unfamiliar, *prima facie* mysterious notion is faithfully rendered in the terms of familiar, understood frameworks. Of course, none of these are examples of the interpretation of a physical theory in the sense at issue here, but they do still typify one central part of the problem relevant to arguments about general relativity: the interpretation of a highly abstract mathematical structure with no ‘obvious’ direct experiential analogue, when that abstract structure is being considered as a possible element of a physical theory.

These examples suggest a way to characterize a metalinguistic interpretation of a physical theory as, in part, a mapping that preserves structure. Say we have two scientific frameworks, one, \mathcal{T} , to be interpreted in the terms of the other, \mathcal{S} . An injective mapping $\phi : \mathcal{S} \rightarrow \mathcal{T}$ *preserves semantic structure* (or is a *semantics-preserving mapping*) when, for every appro-

4. There is another class of questions about a scientific framework related to the interpretative kinds I list here, those pertaining to the relations of a theory to other theories, whether, for example, one theory reduces to or emerges from another, or whether one can cogently or at least consistently translate between models of a system the theory provides and models provided by other theories of the same system, and so on. While these also are questions of great interest, they are not my concern in this article. The issue here is our understanding of general relativity as a scientific framework in its own right. Some, such as Belot (1998), argue that one cannot comprehensively understand a theory without understanding all (or, at least, many) such relations it may stand in to our other best theories. Be that as it may, part of the primary contention of this article is that there are several important senses in which one can understand a theory that do not depend on the relation of the theory to other theories.

5. See, e.g., Hilbert and Cohn-Vossen ([1932] 1983, 242–259)

priate model μ in \mathcal{T} of a physical system, the model $\phi^{-1}[\mu]$ defined in \mathcal{S} by the preimages of all elements of μ composed in the same formal, syntactic structure as in \mathcal{T} is an appropriate model of the same physical system. One of the most remarkable aspects of these sorts of interpretations is that they often make themselves otiose in the end. Once we have an understanding of the semantic content of a system based on that of other, already understood systems, we often no longer require the resources of those other systems in order to employ the resources of the first in explicable, comprehensible and unambiguous ways—we come to understand the semantic relations that compose the concrete interpretation of the original system.

The idea of a mapping that preserves semantic structure provides the terms to state a necessary condition for a physical theory's standing in need of a metalinguistic interpretation. A physical theory requires a metalinguistic interpretation only if we cannot understand, without recourse to the resources of an external framework, how its concrete interpretation renders semantic content to its schematic representations of physical systems, that is, only if we require a semantics-preserving mapping from a well understood framework in order to understand the way its models represent their target physical systems. I take this to be the interesting, substantive claim made by the proponents of the need for an interpretation for general relativity, namely, that the required interpretation be metalinguistic. This minimal condition on the adequacy for an explication of 'interpretation' relevant to these arguments has the virtues of weakness, clarity and manifest physical content. I do not think it sets up a straw man.

It should be clear that the sense of theoretical interpretation I take to be relevant to the arguments does not rely on or assume a shallow instrumentalism. I am not saying that a theory requires no interpretative elucidation of any sort if it has a well understood concrete interpretation already in hand. On the contrary, there are always in general an endless number of interesting and deep interpretative issues of a categorial sort to consider. That, however, is not the sense of interpretation relevant to us here, precisely because the fact that general relativity raises interpretative issues of a categorial sort in no way differentiates it from any other physical theory. In fact, the criterion I posit is strongly anti-instrumentalist, in the sense that it assumes we do not adequately understand a scientific theory when we have *only* operational knowledge of the way its theoretical representations hook up with the physical world—when, that is, we have an adequate concrete interpretation but do not understand in any other than an operational way the representational relations that the concrete interpretation employs to assign physical significance to the elements of the theory.

I can now state the thesis of the article in concise form: because general relativity has a sound concrete interpretation with a well understood semantics, the only substantive, interesting interpretative questions accruing to the theory are of the categorial sort, which accrue equally well to all physical theories; the call for an interpretation for the theory, therefore, is either misguided or redundant.

Before beginning the meat of the arguments, I want to put aside what may seem to be an immediate problem for the criterion. The idea of ‘extratheoretical resources’ is itself only so clear as is the idea of ‘extratheoretical’, that is, as the idea of the boundary of a theory itself, what one counts as internal and external to the theory. This is notoriously difficult to characterize with clarity and precision even on a case by case basis for individual theories much more in a generic way supposedly applicable to a family of theories. No physical theory, I wager, can be characterized with such clarity and exhaustive finality that no question could ever arise as to whether a particular theoretical structure or empirical proposition properly belongs to the theory or not. Luckily for me, I do not need to settle the issue from the start to make my arguments. I cannot give a formal, global, compelling criterion for ‘belongs to the theory’, not even for the single case of general relativity by itself, but I can say in all the cases I look at what at issue is in the theory and what is not without ambiguity, and that suffices for the argument.⁶

3. Diffeomorphic Freedom. The group of diffeomorphisms of the space-time manifold, in some sense or other, can be thought of as a group of symmetries of the physical phenomena treated by general relativity. Invariance of solutions to the Einstein field-equation under the action of diffeomorphisms, however, is not a true symmetry in the sense captured by Noether’s theorem, nor in the sense of a gauge as employed in, for example, Yang-Mills theory, nor in any other well understood sense per-

6. Chris Smeenk pointed out to me in conversation that general relativity does seem to require blatantly extratheoretical resources to construct some types of models, those of the so-called Parametrized Post-Newtonian (PPN) type. (See, e.g., Misner, Thorne, and Wheeler 1973, Chapter 39.) While I think this is an extraordinarily interesting case for those studying the construction and interpretation of inter- and trans-theoretical models, it does not constitute an example in which general relativity requires extratheoretical resources in order to comprehend the semantics of any of its models. The extratheoretical resources used to construct PPN models enter the models only in the service of making possible particular mathematical approximations and simplifications; they do not affect the semantics of the model as a representation of a physical system, which is based entirely on and explicable entirely in the terms of the resources of general relativity.

taining to other physical theories. In consequence there has arisen the issue of the meaning of this diffeomorphic freedom.⁷

I think the most unproblematic and uncontroversial claim one can make about diffeomorphic freedom is that it embodies an irremediable mathematical ambiguity in the apparatus provided by general relativity for the modeling of experiments: the choice of the presentation of the spacetime manifold and metric that one uses to model an experiment is fixed only up to diffeomorphism. A comparison is edifying. Classical mechanics, as embodied respectively in Lagrangian and Hamiltonian mechanics, shares similar ambiguities, slightly different in each formulation of the theory. In Lagrangian mechanics, one is free to choose the Lagrangian function itself up to the addition of an exact antivertical 1-form on the tangent bundle of configuration space (or, in more traditional terms, up to the addition of a total-time derivative of a function of configuration coordinates) without changing the family of solutions the Lagrangian determines.⁸ In Hamiltonian mechanics, one is free to choose any symplectomorphism between the space of states and the cotangent bundle of configuration space, that is, to choose the symplectomorphic presentation of phase space (or, in more traditional terms, the family of canonical coordinates one uses to parametrize phase space), without changing the family of solutions the Hamiltonian function determines.⁹ One feels no lack of understanding of Lagrangian mechanics, no lacuna in its conceptual resources, merely because one is free to choose the form of the Lagrangian more or less freely, just as one is not driven to investigate the ontic status of points in phase space, or of the physical quantities whose values one uses to label those points, in particular which quantities get

7. The problem of diffeomorphic freedom often shows itself in the contemporary literature draped in the garb of the Hole Argument, as in, for example, Earman and Norton (1987), Stachel ([1980] 1989), and Belot (1996). (See Einstein 1914 and Einstein and Grossmann 1914 for two versions of the original argument, and Norton 1989, 1993 for historical and critical discussion.) The lesson of the Hole Argument these days is often thought to bear on the problem of the existence of spacetime points and the debate between substantivalists and relationalists more broadly speaking. The argument's lesson, so claimed, is that one cannot identify spacetime points without reliance on metrical structure, that there is no 'bare manifold of points', as it were, under the metric field (e.g., Belot 1996, and Gaul and Rovelli 2000). In Curiel 2009b I give a detailed rebuttal to contemporary renderings of the Hole Argument in the context of the debate between substantivalists and relationalists, arguing that it is irrelevant to that debate. Because the ontic status of spacetime is most often raised as evidence for the need for an interpretation for general relativity in the context of discussions of quantum gravity, I will wait until Section 5 to discuss it here.

8. See, e.g., Curiel 2009a.

9. Ibid.

nominated ‘configuration’ and which ‘momentum’, merely because one is free to choose whatever symplectomorphism one likes in its presentation.¹⁰

The choice of Lagrangian or the choice of symplectomorphism rests on nothing more than pragmatic considerations of the type adumbrated by Carnap (1956) in his discussion of the choice of a linguistic framework for the investigation of philosophical and physical problems,¹¹ considerations determined by what Robert Geroch calls, somewhat archly yet entirely aptly, ‘psychology’.¹² One chooses on the basis of nothing more than what puts one at ease in any of a variety of ways, from pragmatic considerations such as what will be simple or useful for a particular investigation, to those based on historical custom and aesthetic predilection. It is clear in these cases that the existence of inevitable, more or less arbitrary, nonphysical elements in the presentation of the models of a theory by itself does not require of one the provision of an interpretation of either Lagrangian or Hamiltonian mechanics. More to the point, it is clear in these cases that the physical significance of the theory’s models is not masked or polluted by the unavoidable arbitrariness in the details of their presentations. In the same way, the diffeomorphic freedom in the presentation of relativistic spacetimes does not ipso facto demand an interpretation, in so far as it in no way prevents us from focusing on and investigating what is of true physical relevance in systems that general relativity models, what one may think of as the intrinsic physics of the systems. As to what ‘intrinsic physics’ may mean: it is what Alain Connes was trying to get at, I think, during a conversation on the nature of general relativity as a physical theory, when he pointedly asked me what information I would communicate to beings in a different universe in an attempt to describe to them the spacetime metric of our own.¹³ It is at a minimum, stated in schema, what all possible observers in all possible

10. The analogy of diffeomorphic freedom in general relativity to symplectomorphic freedom in Hamiltonian mechanics is especially striking, as, in essence, any argument one can run in the one context one can run as well (or as poorly) in the other merely by substituting ‘phase space’ for ‘spacetime manifold’ and ‘symplectomorphism’ for ‘diffeomorphism’. Does that show anything of intrinsic physical significance?

11. This is not to say that I consider the choice of a Lagrangian or a symplectomorphism to be the choice of a Carnapian linguistic framework, only that the sorts of considerations that go into each choice are similar.

12. Geroch uses the term in conversation and lectures almost to the point of mannerism, but it is a useful mannerism, always illuminating.

13. I had no good answer at the time, and I’m still not confident I have one. I strongly suspect that Rovelli (2002) provides the tools to begin to construct a good answer. Connes argued that it should be something like the spectrum of the Hermitian operator representing the quantum version of the classical metric. I do not know whether he would endorse that answer today.

states as represented in a model of an experiment would agree on, no matter the presentation of the model used.

In response to this line of thought, an opponent might claim that the analogy is no good. General relativity, he or she could say, is a fundamental theory of a class of physical systems (relativistic spacetimes), whereas Lagrangian and Hamiltonian mechanics are general frameworks within which one formulates physical theories such as general relativity. The ambiguity in the presentation of models that diffeomorphism invariance yields must indicate a fact of *physical* significance, in so far as it is a fundamental, ineliminable structure in a *physical* theory. To dismiss diffeomorphism invariance as having no intrinsic physical significance, the opponent would conclude, is not to take seriously enough on its own terms the formal structures of our best physical theory of spacetime structure: the theory is trying to tell us something, and we are not listening. Rovelli (2000, 118), for example, implicitly suggests a response like this, in his favorable reference to the argument of Stein (n.d.) that Poincaré's failure to take the Lorentz transformations seriously enough as a formal representation of (part of) the physical dynamics of physical systems makes it plausible that he could not have discovered special relativity himself in the first place, and that in the event that failure drove his subsequent refusal to accept Einstein's proposal of special relativity.

I agree with Rovelli on the plausibility of Stein's conjecture that Poincaré's refusal to take the math seriously enough in the context of physical theories hindered his work as a physicist, but I do not think that the same lesson applies here. The important question in both cases is not *whether* we ought to take the math seriously enough, but *what* math we must take seriously. In the years before Einstein's proposal of special relativity, no one knew what the Lorentz transformations meant in the sense that no one knew how to understand them as (part of) a representation of the behavior of physical systems. It was not clear how to devise and perform experiments to probe their manifestation in physical phenomena, or even whether experiments could reveal their effects at all—that is to say, no one knew how to extract what was of intrinsic physical significance from models in which the transformations played a role, whether, indeed, the role the transformations played did reflect anything of intrinsic physical significance. In that state of affairs, we *did* require a metalinguistic interpretation of the transformations to enable us to advance our understanding of that part of the physical world. Einstein provided that interpretation, in the form of the theory of special relativity: the theory itself provided the metalinguistic interpretation of the Lorentz transformations by fixing a cogent semantics for them. Now that the theory is in place and the role of the Lorentz transformations in it well understood, no further clarification of their significance *in* the theory is required.

In a similar vein, the comprehension of special relativity's dismissal of the idea of absolute simultaneity did not require an interpretation of the theory, in any sense of the term; it required only that investigators come to terms with the fact that the fundamental principles of the theory do not allow for the rigorous, physically significant explication of at least some of the fundamental terms of Newtonian physics. Special relativity did have to demonstrate that it could represent experiments in such a way that the semantics of the experimental physics of the day could be used to provide concrete interpretations (in the sense of Section 2) of the models it produced, including those experiments that Newtonian mechanics both could and could not handle, and to do so, moreover, without the need for extratheoretical resources. It did that, and so it became clear that 'absolute simultaneity' in a global sense is not a notion with any natural or even merely reasonable explication in the theory. That demonstration also showed something of deeper significance, that absolute simultaneity is not a notion we should rely on in a search for deeper, better theories, but that fact does not bear on our understanding either of Newtonian mechanics or of special relativity as theories in their own rights.

The same holds for diffeomorphism invariance in general relativity. Before the theory was established, it was not clear what it could have meant for the representation of physical systems to be invariant under the full group of spacetime diffeomorphisms, as Einstein's struggles with the Hole Argument poignantly show. The establishment and comprehension of general relativity itself provides a metalinguistic interpretation for the significance of diffeomorphism invariance, a significance we do understand well in the context of any of the consistent, cogent, concrete interpretations (in the sense of Section 2) we have of the theory: to transform a model of a physical system by the action of a spacetime diffeomorphism does nothing more than change the presentation of the model, but does not alter the intrinsic physics that the model depicts. It is an inevitable ambiguity in our mode of presentation in the theory. Perhaps in some other context a structure purporting to explicate the same aspect of our understanding of physical phenomena will have a different significance—in which case it is overwhelmingly likely that we would have learned something of profound importance about the physical world—but that has no bearing on the quality of our understanding of its role in general relativity.

Perhaps not remarkably, quantum mechanics shares this sort of inevitable ambiguity in the presentation of models. The Hilbert space of a quantum system, including the algebra of self-adjoint operators, does not determine its underlying 'configuration space' any more than does the symplectic structure on the cotangent bundle in Hamiltonian mechanics. Indeed, as is well known, every separable Hilbert space is isomorphic to

every other separable Hilbert space of the same dimension, and the algebra of self-adjoint operators on any presentation of one is itself fixed only up to unitary equivalence. In this case, however, because we do not understand the conceptual resources of quantum mechanics as a theory, this ambiguity may point to the need for a physical clarification, that is, a metalinguistic interpretation, as I discuss in Section 6. We do not understand enough to say either way.

It is a striking fact, surely one worth puzzling over, that all of our physical theories suffer inevitable ambiguity of one sort or another in the models they render of physical systems. It may point to profound questions about our capacity to comprehend the physical world, or point to constraints on the form our comprehension can take or on the content it can achieve. This fact, however, in no case (except that of quantum mechanics) requires the use of extratheoretical machinery for us to grasp what is of true physical significance in the models the theories give us. In particular, nothing in general relativity by itself and nothing in the empirical knowledge we have gained by the application of its models to the representation of physical systems demands or suggests that one must attempt to understand diffeomorphic invariance in any sense other than the one sketched here. One can impose extratheoretical conceptual resources on the theory so as to render the diffeomorphic freedom *prima facie* mysterious but, without having explained why the obvious and clear conceptual resources the theory makes available for the modeling of physical systems do not suffice, to argue on the basis of reasons external to the theory that general relativity requires a metalinguistic interpretation would be circular.

4. Singular Structure. The idea of singular structure in relativistic spacetimes is a broad one, encompassing at least the presence of incomplete, inextendible timelike curves, causal loops, failures of global hyperbolicity, and so on. It has sometimes been held that the seemingly pathological character of these phenomena taken together with their generic appearance in otherwise seemingly reasonable solutions to the Einstein field-equation jointly show that general relativity must be inadequate as a physical theory, in the sense that external conceptual resources will have to be brought to bear on it in order to have good physical sense made of those solutions. There are too many such phenomena, and they raise variously too many important and deep problems, to allow for detailed discussion here, so I will discuss only one, sketching why the seemingly pathological consequences of the existence of incomplete, inextendible curves do not demand a metalinguistic interpretation for general relativity. The discussion of the rest of the phenomena follows similar lines; the interested reader should consult Curiel 1999.

Recall that a spacetime is said to be singular if it itself is maximal, in the sense that it cannot be isometrically embedded in a larger spacetime, and if it contains an incomplete, inextendible timelike geodesic. The first condition, of maximality, ensures that the geodesic's incompleteness is nontrivial in a technical sense with direct physical significance: the incompleteness does not arise from a poor initial choice of spacetime model. Maximally incomplete timelike geodesics offer up many physically and psychologically unsettling possibilities. Since a timelike geodesic is the possible worldline of an observer or a particle, it is *prima facie* possible that an observer or particle could traverse such a singular curve, which seems to imply that a particle could pop in or out of existence *ex nihilo* or *ad nihilum* with no known physical process or mechanism capable of effecting such a thing dynamically, or even that an observer could experience only a finite total interval of time in which to live without ever aging past a certain limit and without ever dying. The spacetimes having these maximally inextendible curves, moreover, are not *recherché* outliers, but include essentially all of the most physically important solutions to the Einstein field-equation, including: the Friedman-Robertson-Lemaître-Walker (FLRW) family of spacetimes, representing universes uniformly filled with homogeneous, isotropic dust or fluid, which to a high degree of approximation model the large-scale structure of our universe well; Schwarzschild spacetime, the unique static, spherically symmetric spacetime representing an isolated static, spherically symmetric body such as a star or a black hole; Reissner-Nordström spacetime, the generalization of Schwarzschild spacetime allowing for the central spherical body to have an electric charge; and Kerr-Newman spacetime, the generalization of Reissner-Nordström spacetime allowing the central body to have angular momentum.¹⁴

The template of the argument that the generic prevalence of singular structure in solutions to the Einstein field-equation demands that general relativity be given an interpretation, reconstructed in the terms of this article, goes like this.

1. General relativity allows for models of spacetime with singular structure of a particular form *S*.
2. In those models, a physical system behaves in a manifestly pathological manner *P* that seems to derive from or depend on the presence of *S*.
3. *P*, if taken at face value, conflicts with dearly held principle *Z*.
4. Therefore, we cannot take *P* at face value.

14. See Hawking and Ellis 1973 and Wald 1984 for a rigorous exposition of these spacetimes.

5. Therefore, we do not understand the way that general relativity's models represent the physical world under its standard concrete interpretations (in the sense of Section 2).
6. Therefore, general relativity requires a metalinguistic interpretation.

We will, as I said, discuss only the prediction of incomplete, inextendible timelike curves, which fixes S in the template. P is taken either to be the vanishing or appearance of particles or the finitude of an observer's existence, when the particle or observer is taken to traverse the incomplete curve. For Z , I discuss here only two concrete proposals: (1) consistency with other physical theories, and (2) comprehensibility of ontology.

I consider each briefly in turn, to show that, in so far as it is a problem at all, it is one at most of the categorical sort, not one indicating that we lack understanding of the concrete interpretations we have for general relativity. (Other problems the presence of singular structure supposedly entails for general relativity are dealt with by considerations similar to those I present against these two; see Curiel 1999 for details.)

On physical grounds, curve incompleteness has been objected to because it seems to imply that (nonvirtual) particles could be annihilated or created right in the middle of spacetime, with no known physical force or mechanism capable of doing such a thing, thus indicating a possible inconsistency with other physical theories.¹⁵ This complaint, however, does not in fact say that we do not understand the way significance accrues to the *prima facie* concrete content of a model general relativity provides of a physical system; it says rather only that there is no obvious way to square the content of the model with what we know on the basis of other physical theories. Based on our best physical knowledge of the behavior of particles of all sorts, this behavior is not in fact allowed; therefore, the proper conclusion to draw seems to be that general relativity allows models of spacetimes that could not be instantiated on pain of falsifying other physical theories. This happens all the time between theories, however, without marking a lack of understanding, because different theories have different regimes of applicability. For all we know, moreover, the physical world really does act like this—general relativity may be trying to tell us something. There is, in any event, no problem with our understanding of the concrete content of the spacetime model at issue here—we can say using only the conceptual resources general relativity provides for us, without ambiguity and with a perfect grasp of all the semantic relations, what the outcome of any experiment would be in such a spacetime. Thus, the alleged

15. Cf., e.g., Hawking 1967, 189, for such an argument with specific regard to particles, and Clarke 1975 and Ellis and Schmidt 1977 for similar arguments with respect to the demand that spacetime itself be maximal.

inconsistency with other theories does not show that we lack understanding of the way general relativity represents physical systems.

The presence of incomplete, inextendible curves on a spacetime manifold poses an interesting ontological problem. The curve itself, of course, any finite extent of it, is as physical a thing as one could want, and its existence can be characterized in more or less the same way as any localized object in spacetime. The *incompleteness* of the singular curve, however, as a property of the curve, does not exist in any localized sense in spacetime.¹⁶ In general, one cannot even associate the incompleteness of an incomplete, inextendible curve with a bounded region of spacetime. The existence of such an incomplete curve is, in a technical sense, a global feature of the spacetime manifold and its metric. Still, the incompleteness of the curve surely exists in some sense or other, as evidenced by the wide variety of manifestly physical phenomena that can be nomically associated with it, such as the presence of a black hole—it just does not exist locally in space and time, so to speak, the way a (classical) particle does. This nonlocalizability of singular structure may suggest that there are elements of general relativity's concrete models whose representational content we do not understand, in so far as it seems that singularities are real things and yet we cannot fix an ordinary ontology for them using only the resources the theory makes available to us. In fact, however, the nonlocalizability of incomplete, inextendible curves is no different from that of any other global topological structure, such as the spacetime manifold's paracompactness or the fact that it is Hausdorff separable or the value of its Euler characteristic. The ontological status of these sorts of global structures is a fascinating problem, to be sure, but it is one of a categorial character, shared by all other theories of spatiotemporal structure. That we cannot localize the paracompactness of the spacetime manifold does not stop us from clearly grasping the concrete content of all spacetime models the theory presents us with, and nor does the fact that we cannot localize the incompleteness of an incomplete, inextendible curve.

5. Quantum Gravity. Many eminent workers in the field of quantum gravity suggest that at least part of the reason we have so far had little if any success in the search for a viable theory points to a lack of understanding of the conceptual resources of general relativity. We require a proper interpretation of general relativity, they argue, in order to find a path to the deeper theory. In these arguments, they tend to focus on four issues: the significance of time and temporal evolution in the theory; what to count as observables in the theory; the meaning of diffeomorphic freedom;

16. See Curiel 1999 for a precise description of this fact.

and the ontic status of spacetime points.¹⁷ We have already discussed the third issue, so I treat here only the first, second and fourth. I will not enter into the technical details of the problems, as they are not required to make the case.

In brief, the problem of time bears on attempts to formulate a theory of quantum gravity in a canonical framework, as, for example, in the attempt to impose a Dirac-style quantization on general relativity considered as a constrained Hamiltonian system. The role time plays in general relativity differs from that in canonical theories in ways so profound as to make it difficult even to begin to see how to reconcile them. One cannot in any principled way impose a preferred temporal frame in a generic, relativistic spacetime. No congruence of timelike curves has any privileged status over any other. A privileged temporal frame, however, lies at the heart of a canonical theory, picked out by the parameter of the dynamical evolution conjugate to a system's Hamiltonian operator. That we do not see how to reconcile the role time plays in general relativity with that in canonical theories, however, does not by itself show that we lack understanding of its role in general relativity, or, for that matter, of its role in canonical theories. The two may be truly incompatible in a fundamental way, or we may not yet have found the way to reconcile them within the constraints imposed by our understandings of each in the confines of their respective theories. We cannot say in our current understanding of the physical world.

On a closely related note, a *prima facie* strong case can be made that general relativity requires a metalinguistic interpretation by asking what the theory provides or even allows for in the way of canonical observables when one tries to formulate it as a constrained Hamiltonian system. It is extraordinarily difficult, at best, to construct canonical observables for a spacetime when modeled in that framework (Rovelli 1991b). Indeed, a theorem of Torre (1993) states that in the special case of closed, vacuum spacetimes there are *no* observables that are local in a physically important sense. If we cannot even define the observables for the theory, the thought goes, then surely we do not understand how significance accrues to the elements of the models the theory renders for observable systems, and thus general relativity requires a metalinguistic interpretation. The weakness of this argument, of course, is that the sort of observable at issue is

17. See, e.g., Unruh and Wald 1989; Rovelli 1991a; Kuchař 1992; and Isham 1993 for reviews of different forms of the problem of time and for proposed solutions; see, e.g., Rovelli 1991b for a discussion of what to count as an observable; see, e.g., Rovelli 2000 for an exposition of the questions about diffeomorphic freedom and the ontic status of spacetime points; see Belot 1996 and Belot and Earman 2001 for discussion of the philosophical issues from points of view opposed to mine.

of a very special sort, that defined in the context of a theory formulated as a constrained Hamiltonian system. No comparable problem arises, for example, when one formulates general relativity in the Lagrangian framework, much less when one simply formulates general relativity as a theory in its own right, independent of any such larger framework in the terms of which one may try to rewrite it. This is shown clearly, for example, by Rovelli's 2002 construction of a complete set of simple, realistic observables that are invariant under the action of diffeomorphisms and are available in completely generic spacetimes with only the most minimal constraints on the forms of matter allowable.

Because of the incompatibilities between the roles played by time and temporal evolution in general relativity on the one hand and canonical theories on the other, we already had strong reason to think that general relativity and the canonical framework do not easily fit together. The fact that it is difficult to find canonical observables for general relativity when one tries to force it into the canonical framework strengthens the plausibility of the conclusion. Precisely because the fact does not pertain to our understanding of general relativity as a theory in its own right, however, it cannot be used to argue that general relativity stands in need of a metalinguistic interpretation. To attempt to force general relativity into the straitjacket of an external framework that does not naturally suit it, to find that the structure of general relativity does not conform to the character of the conceptual resources of the external framework, and then to conclude that we must have recourse to extratheoretical resources in order to understand the significance of general relativity as a physical theory based on that is to argue in a circle, for the puzzle arises only upon the demand that one incorporate extratheoretical elements into the theory in the first place. General relativity by itself is a perfectly good physical theory without the imposition of those extratheoretical resources. It is surely a puzzle why general relativity does not easily conform with one of the fundamental frameworks we have for the formulation of physical theories, and that is a problem we ought to investigate. It is not, however, a puzzle that shows that we require a metalinguistic interpretation of general relativity in order to understand it.

Now, let us move on to the fourth item on the list, the ontic status of spacetime points: diffeomorphism invariance per se may not demand interpretation, as I argued in Section 3, but several workers in the field hold that it grounds the construction of an argument that does point to a problem requiring interpretation, the infamous Hole Argument, which supposedly bears on the problem of the existence of spacetime points and the debate between substantivalists and relationalists more broadly speak-

ing.¹⁸ The argument's lesson, so claimed, is that one cannot identify spacetime points without reliance on metrical structure, that there is no 'bare manifold of points', as it were, under the metric field.¹⁹ I am not going to discuss the Hole argument here, as the heart of the issue is the ontic status of spacetime points, irrespective of the particular arguments one uses to try to show that the structure of general relativity militates in favor of one side or another in the contemporary debate between substantialists and relationalists. When one puts the issue this way, however, it should be clear that this is at worst (or best) a categorial problem. This shows itself in the fact that one does not need to take a stand on the issue of the existence of spacetime points at all in order to understand the way that the concrete interpretations of general relativity (in the sense of Section 2) work. If one wants to do this sort of thing, moreover, one can either make spacetime points part of the ontology of a concrete interpretation of general relativity or not; the semantics can easily accommodate either choice. Whether or not to include spacetime points as part of the ontology of the concrete content of general relativity as a physical theory is a categorial problem, just as it is in every other theory of spatiotemporal structure.

Still, workers in the field hold that the role of time, diffeomorphic freedom, the ontic status of spacetime points and other issues like them *are* problems because they make it difficult to attempt to 'quantize' general relativity in any of the standard ways. The thought then runs that a resolution of the perceived conceptual problems in general relativity holds out the promise of guidance towards a theory of quantum gravity. There are two questions to keep separate here: first, whether it is plausible that the resolution of conceptual difficulties in general relativity would provide clues to a theory of quantum gravity; and second, whether the difficulties that the standard procedures of quantization face when applied to general relativity can in any plausible sense be attributed to a failure on our part in understanding the theory in a sense that only a metalinguistic interpretation could resolve.

Now, if there is one feature of a viable theory of quantum gravity that one can foresee with some assurance, it is that the theory will deal with issues of space, time and spacetime in ways far different than does general relativity, on account of the fact that general relativity has no superposition principle, no uncertainty principle and no measurement problem. If we lack an understanding of the significance of time or of any of the

18. See note 7 for several citations relevant to the Hole Argument.

19. See, e.g., Belot 1996 and Gaul and Rovelli 2000. In Curiel 2009b I give a detailed rebuttal to contemporary renderings of the Hole Argument in the context of the debate between substantialists and relationalists, arguing that it is irrelevant to that debate.

rest as physical phenomena in abstraction from its treatment by any particular theory, and we also believe that general relativity is not a fundamental theory of the world, it is difficult to see why an interpretation of general relativity that clarified those problems should give us any richer understanding of the structure of the physical world in such a way as to illuminate a path to a viable theory of quantum gravity.²⁰ Those who hold that we must find an interpretation for general relativity in order to make progress toward a theory of quantum gravity a fortiori do not view general relativity as a ‘fundamental’ theory, whatever exactly that may come to. This seems to represent a tension in their view, for it is difficult to see, even apart from this particular case, how an interpretation of a ‘higher’ or ‘emergent’ physical theory could shed light on the conceptual resources of a more fundamental theory to which it somehow reduces. On the face it, the clarification of conceptual structure seems always to point in the opposite direction.

Consider the move from pre-Newtonian physics to the system Newton presents in *Principia*. In light of our present state of knowledge, we are tempted to say something like the following: though it is often claimed that the Ptolemaic and the Copernican systems are observationally indistinguishable, in fact they are distinguishable in so far as each predicts either an angular momentum for each of the planets different from that predicted by the other, or an inertial mass for each planet different from that predicted by the other. If we assume that the inertial mass of each planet is the same in each theory, then, because the effective axis and angular velocity of orbital rotation of each planet is different in each, one theory will predict for it a different angular momentum than does the other; likewise, if one adjusts the inertial mass of each planet in the different theories so as to make the angular momentum predicted by the one accord with that by the other, then one will have eo ipso made the masses each theory attributes to the planets different from those the other does. (Even in the latter case, the angular momentum of each planet would have a different axis in the two theories, but we let that pass.)

Strictly speaking, however, the Ptolemaic and the Copernican systems could not in their own time, not even in principle, have made those differing predictions, for the simple reason that neither had the theoretical resources to represent either inertial mass or angular momentum in anything like the senses that the work of Galileo, Huygens and Newton rendered to the concepts. Neither system had a notion of kinematics or of dynamics that would have allowed for the understanding of those quantities, for, among other lacks, neither had a kinematical represen-

20. Belot (1998) makes a similar point with regard to the role that the vector potential plays in classical Maxwell theory and that in quantum electrodynamics.

tation of a measure of a body's inertia adequate for the representation of the motion of a body under an impressed force, and neither had a representation of anything analogous to equations of motion in the sense of Newton's Laws adequate for the computation of dynamic quantities such as angular momentum. It thus could not have been either wrong or right to make the claim, strictly within the context of either of the two historical systems, that the one predicted different planetary angular momenta or inertial masses than the other, for one cannot meaningfully formulate the propositions in either of them. Indeed, with regard to the more fundamental putative difference between the two systems—whether the Earth or the Sun is at rest—neither had a concept of motion sophisticated enough to lend the question any clear sense, much more to make it amenable to investigation by either theoretical or observational means.

In light of the groundbreaking work of Galileo on inertial mass and motion, of Huygens and Newton on inertial mass and rotational motion, and of Newton on dynamics and his theory of universal gravity, all of which was in the service of constructing a new theoretical framework rather than extending or deepening the old ones, it became clear that, *within the context of the new system of mechanics*, the analogue of the Copernican system is the correct one, as Newton ([1686] 1999, Book III, Theorem XII) himself pointed out, albeit in a characteristically circumspect and exact manner. It became clear that this was correct, moreover, just in so far as the new system had the conceptual resources to make meaningful the controverted pseudopropositions, as it were, about rest and motion, and to make clear the propositions we can now pose about differences between inertial mass and angular momentum and put to experimental test. It was not, however, a clarification of the notions in the context of the old theory that showed the way to the new theory. Rather, the recognition that the concepts as represented in the old theory were inadequate for the new tasks at hand, which included the development of the idea of the equations of motion of a physical system, provided impetus for and guidance in the search for the new theory. This may be the case with, say, time as represented in general relativity vis-à-vis what a theory of quantum gravity will require for its proper treatment of time. But, again, this would not show that our understanding of general relativity itself is in any way deficient. It may show that its conceptual resources are inadequate for the physical questions we now want to ask, but that is a different matter.²¹

21. In the event, I strongly suspect general relativity is *not* inadequate for any tolerably well posed physical questions we now have in its established regime of applicability. No 'theory' of quantum gravity is mature and cogent enough to pose questions of manifest physical content that general relativity cannot handle, or, at least, quantum

Now, for the second question—whether the difficulties that the standard procedures of quantization face when applied to general relativity can in any plausible sense be attributed to a failure on our part in understanding the theory—I would reply that in principle, and a very important principle at that, we ought not to be ‘quantizing’ anything in the first place. Maxwell and Boltzmann did not ‘statisticalize’ classical fluid dynamics in order to arrive at a viable account of molecular kinetics. One would not even have known where to have started if one were to have attempted to follow that idea. In the event, Maxwell and Boltzmann had to attempt to formulate the relevant kinetic and statistical concepts and the relations among them, and to clarify them all, with no particular regard at first for whether or not those concepts constituted in some sense a ‘proper reduction’ or ‘clarification’ of the conceptual apparatus of classical fluid dynamics. They based their investigations to a large extent on the known, well entrenched body of experimental knowledge they were trying to produce a better theory of. This is not to say that they did not use higher-order theories such as fluid dynamics in their work, to help them, for instance, in imposing useful organization on parts of the body of experimental knowledge. In doing so, however, they did not attempt directly to refine or clarify the conceptual resources of those theories in the context of the theories themselves in their search for a better one.

Boltzmann’s statistical treatment of the Second Law, for example, did not depend on the resources of any particular higher-order theory. He could deal with the phenomena it pertains to outside the context of any definite theory, just as Carnot ([1824] 1960) himself did in his analysis of heat cycles, which abstracted the idea of a heat engine from contemporary theory and abstracted in turn from this the even more schematic notion of a thermodynamical system, even if he chose in places to use the language of the caloric theory, even if he felt in places he *had* to use that language. Only after the new statistical mechanics had been put into something like a definite form were Maxwell and Boltzmann in a position to verify that, for example, the collision equation, in the context of appropriate approximations and idealizations, does lead to the Navier-Stokes equations.²² The new statistical mechanics then rendered an ex-

field theory posed on the classical background of the spacetimes of general relativity. The only questions that a putative theory of quantum gravity could pose that could not be handled by one of those two frameworks is about the quantum nature of the metric and curvature themselves; we know so little about phenomena at the relevant scales, however, that it is difficult for me to see how questions about such things can have manifest physical content. They are at the moment, of necessity, questions that spur us on to try to acquire the knowledge we need to turn them into questions with real physical significance.

22. See, for instance, Sommerfeld 1964, 293–318.

plication and conceptual clarification of the notions of the older theories, such as those of pressure and temperature. A deeper understanding of pressure and temperature within the context of the older theories themselves did not contribute to the construction of the new statistical mechanics.

That we find ourselves in the position of having to attempt to work from the top down in the case of quantum gravity, trying to base our search on the conceptual resources of theories thought to be less fundamental rather than tackling the formulation of such a theory head on—that we find ourselves having to ‘quantize’ rather than ‘classicize’—shows that we lack experimental knowledge, not necessarily conceptual understanding, at least with respect to the resources of general relativity. We need hard experimental data that we do not know how to accommodate in the context of general relativity to move forward in the understanding of the physical regime that general relativity treats. In this vein, I find it suggestive that I have never read or heard of an experimentalist who complains about a lack of conceptual clarity hindering the experimental search for clues to a theory of quantum gravity; one hears only of theoreticians complaining about the difficulties in making theoretical progress.

That general relativity in and by itself does not suggest a way to formulate itself as some sort of a classical limit of a quantum theory does not show that general relativity itself, as a complete physical theory, stands lacking with respect to semantic content, with respect to our understanding of the way its concrete interpretations (in the sense of Section 2) yield models of physical systems it is appropriate for the treatment of. I think it rather suggests that our lack of otherwise inexplicable experimental data and our lack of a secure and cogent conceptual understanding of quantum mechanics holds us back, and, because of that very fact, the clarification of quantum mechanics’ resources, it seems to me, perhaps does hold out the promise for clues to a theory of quantum gravity, if anything apart from radically new experimental data does.

6. Comparison to Quantum Mechanics.

In my childhood, the legend was current that only twelve men in the world understood Einstein’s theory. Nowadays, relativity is quite tame; but *nobody* yet understands the quantum theory. (Stein 1972, 367–368)

The theory that philosophers and physicists have until now felt most required an interpretation is quantum mechanics. With regard to the impetus for demanding an interpretation of a theory, the difference between

general relativity on the one hand and quantum mechanics on the other could not be more complete: it is the difference between a benign ambiguity in the machinery of the former theory that models experiments, and an insidious ambiguity as to what counts, according to the latter theory, *as* an experiment.²³ Quantum mechanics requires at the end of the day a *deus ex machina* to arrive at a mathematical model of an experiment that offers definite (albeit, probabilistic) predictions for the readings of measurement devices—to model a physical system, one must decide at what point in a system’s physical evolution the wave-function ‘collapses’ in order to extract a definite (probabilistic) prediction for the reading of a measuring apparatus, which is to say in more classical terms that one must decide what moment in the system’s interaction with its environment constitutes an ‘experiment’.²⁴ In the terms of this article, we lack an adequate understanding of the concrete models of quantum mechanics, and so we stand in need of a metalinguistic interpretation for it.

General relativity, by contrast, while offering up many superficially distinct ways of modeling any given experiment, does not require one to classify different moments in a system’s evolution as constituting an ‘experiment’ in order to arrive at definite predictions for the readings of experimental apparatus. All moments are treated on an equal footing by the theory. This is not to say that I think physics to be only the definite prediction of readings for experimental instruments. As Wigner once remarked, quantum mechanics not only predicts scattering amplitudes; it also allows us to calculate the specific heat of substances. Both classical and quantum theories share this capacity to illuminate features of intrinsic physical significance about the world without regard to their manifestation

23. Thermodynamics, in particular the Second Law, has a sort of ambiguity about it similar to that of quantum mechanics. Locally, it is not clear that the Second Law is even valid unless one is extremely judicious in picking one’s system of study, which includes the selection of the interval of that system’s physical evolution during which one will probe it and a *fortiori* selection of the types of interactions the system will have with its environment, as the appropriate ones for the making of thermodynamical measurements—which is to say that, so far as the Second Law is concerned, it is not clear what constitutes an ‘experiment’ and what does not. I think this does show that the Second Law requires interpretation in the sense relevant to this article.

24. The force of this remark does not depend on one’s having something like Born’s understanding of quantum mechanics. In order to reconcile the evolution of superposed states encoded in Schrödinger’s equation with the extraordinarily precise measurements we make, say, of the spin of neutrons or of the positions of supercooled atoms, something has to give somewhere no matter how one tries to understand standard quantum mechanics on its own terms. (Bell [1987, *passim*] makes the same point.) Geroch (1984), for example, makes an interesting case against this view, but I find more compelling the arguments of Stein (1972) in general, and those of Stein (1984) against positions such as Geroch’s in particular.

in any particular experimental arrangement—all the more remarkable in light of our lack of understanding of quantum mechanics.²⁵

Quantum mechanics demands a metalinguistic interpretation because it is not clear how to model physical phenomena, how to model the outcomes of experiments simpliciter: the predictions of ‘pure’ quantum theory are in some sense in contradiction with the outcomes of experiments, but not in such a way as to invalidate the theory but rather to substantiate it—an extraordinary state of affairs.²⁶ In particular, nothing in the framework of quantum mechanics itself, no well understood concept or term, holds out any promise of answering the question, What is a measurement? This problem has two independent components, the combination of which lends the problem its depth and urgency. First, the question points to a lack of understanding about the production of models of physical systems in the theory. The production of such models, and their concomitant comparison to experiment, is one of the most important ways that theoretical structures can accrue to themselves empirical content; but our comprehension of the way that such models represent physical systems is the sole guarantee we can have of their propriety and soundness. Second, the question by its very formulation seems to require the introduction of extratheoretical conceptual resources in order to come to grips with it: one can appreciate the full import of the question only by way of a notion, ‘measurement’, that one cannot express in strictly quantum mechanical terms, perhaps not even in strictly physical terms.²⁷ None of the problems raised by the proponents of the need for an interpretation for general relativity shares these features: there simply is no analogous problem in general relativity. We know how to model in the terms of the theory experiments that manifest and probe every phenomenon suggested or predicted by the theory, with no inconsistency of any kind, for we understand with no lack of cogency the fundamental, physical terms and principles of the theory in which one articulates its models and draws conclusions on their basis—we understand the semantics of all good concrete interpretations we have of the theory, and there are no funda-

25. Perhaps this shows that the conceptual resources involved in a theory’s prediction of the experimental values of dynamic quantities such as scattering angles, or at least the ways that those resources come into play in that task, differ in a deep way from those in the computation of kinematic quantities such as specific heats?

26. “Mathematicians, who need only simple axioms about otherwise undefined objects, have been able to write extensive works on quantum measurement theory—which experimental physicists do not find it necessary to read” Bell ([1981] 1987, 117).

27. Bell ([1975] 1987, [1981] 1987, [1984] 1987) makes this point eloquently.

mental inconsistencies between those different concrete interpretations.²⁸ All the conceptual problems we face in general relativity, considered as an independent theory in its own right, are of the categorial sort. In quantum mechanics, we do not even know what the fundamental terms and principles are. ‘Measurement’? ‘Interaction’? ‘Observation’?

Rovelli (2000, his emphasis) says, “In [the] effort [to find a theory of quantum gravity], physics is once more facing *conceptual* problems: *What is matter? What is causality? What is the role of the observer in physics? What is time? What is the meaning of ‘being somewhere’? What is the meaning of ‘now’? What is the meaning of ‘moving’? Is motion to be defined with respect to objects or with respect to space?*” I agree wholeheartedly with this assessment.²⁹ None of these problems, however, *is* a problem in the context of general relativity, at least no more so than in any other physical theory—they are categorial problems for the theory. *All* of them *are* problems in quantum mechanics, and, indeed, are problems in a way not shared by any other theory we know: they are problems that lie at the bottom of the need of a metalinguistic interpretation for quantum mechanics. The quantum picture of the world is just too different from the classical, however, to make it anything more than merely possible—certainly not *prima facie* probable, and it seems to me, in the event, quite unlikely—that the resolution of subtle questions of categorial interpretation in a classical theory would lead to resolution or even clarification of any of the conceptual problems of quantum mechanics. In the absence of empirical data that shows us definitively where general relativity or quantum mechanics goes wrong or is otherwise inadequate—for we have none—it makes sense to focus on the theory that we manifestly do not understand.

7. Conclusion. The word ‘interpretation’ has a well entrenched usage in the literature, pertaining to problems of the sort I have gestured at in our understanding of quantum mechanics, such as it is. I think we should keep it that way. It is best not to assimilate such a useful, focused idea to others of a completely different stripe by using the same word promiscuously to cover all cases where we lack understanding of the physical world and the way our theories describe it.

Many of the issues the proponents of a need for an interpretation raise,

28. Rovelli (2002), for instance, provides a concrete set of modeling tools one could use to formulate this proposition with some precision and rigor and to prove it; this would be an interesting exercise.

29. As a personal prejudice, I would remove ‘causality’ from the list of fundamental problems, as it does not seem to me a notion that one needs or even wants in physics, but that is beside the point here.

as we have seen, turn on difficulties in making sense of the relations between general relativity and other physical theories or between general relativity and abstract frameworks within which one formulates theories. These are not problems about the interpretation of general relativity, not even ones of a categorial sort. They rather, it seems to me, betray a lack of comprehension on our part about the character of the physical world itself. We have produced several ways of trying to represent and so to understand the world, in this case, say, general relativity and the canonical framework; they appear incompatible on the face of it, and yet we understand each on its own fairly well. This suggests to me that the resources of each on its own is adequate for the representation of a subset of that part of the world that it treats, and that, moreover, the subset each treats differs from that treated by the other. Neither suffices to treat everything we know about that part of the world, and we do not know how to put them together. It seems to follow that we lack knowledge and understanding of that part of the world, and more particularly of how to construct a theoretical representation that encompasses all we do know of that part of the world; it does not follow that we lack understanding of the theories we already have for the representation of that part of the world. Part of the job of trying to understand that part of the world better, moreover, will surely require us to learn more about the content of our theories that do represent it well. In so far as we understand the representational resources of the theories we already have, that job will require, at most, interpretative work of the categorial sort, nothing more.

I do not conclude from this discussion that general relativity on its own, irrespective of its relations to other theories, poses no questions of important and profound philosophical interest: quite the contrary. I conclude only that the problems there are do not suffer from conceptual incoherence in the very attempt to formulate them, and so do not indicate a lack of comprehension about the nature of the theory as a *physical* theory—all the interpretative problems are categorial, not metalinguistic. Admittedly we may not—and I am sure we in fact do not—possess adequate understanding of many of the features of the physical world that general relativity represents. It may turn out, for example, that an element of general relativity that we find unproblematic in our current state of understanding of the physical world will turn out to require radical alteration as our empirical knowledge increases, as with the notion of ‘simultaneity’ in the move from Newtonian mechanics to special relativity as demanded by the ever-increasing body of experimental data that Newtonian mechanics, as a foundation for Maxwell theory, could not accommodate. Diffeomorphism invariance, for instance, may turn out to be just such an element. Were we to work out what, if anything, there may be of intrinsic physical significance in the phenomena whose representation

in general relativity includes in part the diffeomorphic invariance of its structure, then perhaps we would see that our understanding of that part of those phenomena had been flawed. This, however, would *not* show that we lacked an adequate understanding of the role diffeomorphism invariance plays in general relativity that only a suitable metalinguistic interpretation could have rectified. It would rather show that general relativity is inadequate as a physical theory in its treatment of those features of the physical world it purports to represent by, in part, the use of diffeomorphism invariance.

Attempts to improve our understanding of such parts of the physical world should not focus on the significance of the representations of them that general relativity provides; the attempts should, on the contrary, focus on the ways that general relativity *fails* to provide adequate representations of them, if it does so. In the event, we know of no phenomena that general relativity treats for which it provides inadequate models. Unlike the time prior to the formulation of special relativity, no body of otherwise inexplicable experimental data, in this area of physics at least, cries out for the comprehension that only a dramatic improvement in our conceptual resources could provide. This, to my mind, is the great problem facing physics today, or at least that part of it comprehended by general relativity: not that we understand too little or do not understand clearly enough; rather that we understand too well but do not *know* enough. The next great advance in these areas of physics, I suspect, is not awaiting a Swiss patent-office clerk to revolutionize the conceptual resources we have for understanding the physical world. We are not ready for that yet. We rather await a Leverrier, a Fizeau, a Hertz, a Michelson and Morley, to show us what our current theories cannot adequately describe, to provide us the body of otherwise inexplicable empirical knowledge that will furnish the raw material for and *drive* the hoped-for conceptual revolution.

REFERENCES

- Bell, J. ([1975] 1987), “The Theory of Local Beables”, presented at the sixth GIFT Seminar, Jaca, June 2–7, and reproduced in *Epistemological Letters*. Reprinted in *Speakable and Unsayable in Quantum Mechanics*. Cambridge: Cambridge University Press, 52–62.
- ([1981] 1987), “Quantum Mechanics for Cosmologists”, in C. Isham, R. Penrose, and D. Sciama (eds.), *Quantum Gravity 2*. Oxford: Clarendon Press, 611–637. Reprinted in *Speakable and Unsayable in Quantum Mechanics*. Cambridge: Cambridge University Press, 117–138.
- ([1984] 1987), “Beables for Quantum Field Theory”, *CERN-TH 4035/84*. Reprinted in *Speakable and Unsayable in Quantum Mechanics*. Cambridge: Cambridge University Press, 173–180.
- (1987), *Speakable and Unsayable in Quantum Mechanics*. Cambridge: Cambridge University Press.
- Belot, G. (1996), “Why General Relativity *Does* Need an Interpretation”, Lindley Darden (ed.), *PSA 1996: Proceedings of the 1996 Biennial Meeting of the Philosophy of Science*

- Association*. Part I: *Contributed Papers*. East Lansing, MI: Philosophy of Science Association, 80–88.
- (1998), “Understanding Electromagnetism”, *British Journal for the Philosophy of Science* 49: 531–555.
- Belot, G., and J. Earman (2001), “Pre-Socratic Quantum Gravity”, in C. Callender and N. Huggett (eds.), *Philosophy Meets Physics at the Planck Scale*. Cambridge: Cambridge University Press, 213–255.
- Carnap, R. (1956), “Empiricism, Semantics, and Ontology”, in *Meaning and Necessity: A Study in Semantics and Modal Logic*, 2nd ed. Chicago: University of Chicago Press, 205–221. Published earlier in *Revue Internationale de Philosophie* 4: 20–40.
- Carnot, S. ([1824] 1960), *Reflections on the Motive Power of Fire, and on Machines Fitted to Develop That Power*. Reprint. Translated by E. Mendoza. Originally published as *Réflexions sur la puissance motrice du feu et sur les machines propres à développer cette puissance* (Paris: Chez Bachelier). New York: Dover.
- Clarke, C. (1975), “The Classification of Singularities”, *General Relativity and Gravitation* 6: 35–40.
- Curiel, E. (1999), “The Analysis of Singular Spacetimes”, *Philosophy of Science* 66 (Proceedings): S119–S145.
- (2009a), “The Conceptual Structure of Lagrangian and Hamiltonian Mechanics”, unpublished manuscript in preparation for submission.
- (2009b), “On the Existence of Spacetime Structure”, *Studies in the History and Philosophy of Modern Physics*, forthcoming.
- Earman, J. (1986), *A Primer on Determinism*. Dordrecht: Reidel.
- Earman, J., and J. Norton (1987), “What Price Spacetime Substantivalism? The Hole Story”, *Philosophy of Science* 38: 515–525.
- Einstein, A. (1914), “Die Formale Grundlage der Allgemeinen Relativitätstheorie”, *Königlich Preussische Akademie der Wissenschaften—Phys.-math. Klasse Sitzungsberichte* 1914: 1030–1085.
- Einstein, A., and M. Grossmann (1914), “Kovarianzeigenschaften der Feldgleichungen der auf die Verallgemeinerte Relativitätstheorie Gegründeten Gravitationstheorie”, *Zeitschrift für Mathematik und Physik* 63: 215–225.
- Ellis, G., and B. Schmidt (1977), “Singular Space-Times”, *General Relativity and Gravitation* 8: 915–953.
- Gaul, M., and C. Rovelli (2000), “Loop Quantum Gravity and the Meaning of Diffeomorphism Invariance”, in J. Kowalski-Glikman (ed.), *Towards Quantum Gravity*, Lecture Notes in Physics, vol. 541. Berlin: Springer-Verlag, 277–324, arXiv:gr-qc/9910007v2.
- Geroch, R. (1984), “The Everett Interpretation”, *Noûs* 18: 617–633.
- Hawking, S. (1967), “The Occurrence of Singularities in Cosmology III”, *Philosophical Transactions of the Royal Society A* 300: 187–210.
- Hawking, S., and G. Ellis (1973), *The Large Scale Structure of Space-Time*. Cambridge: Cambridge University Press.
- Helmholtz, H. von (1870), “Über den Ursprung und die Bedeutung der Geometrischen Axiome”, in *Populäre Wissenschaftliche Vorträge von H. Helmholtz*, vol. 3. Braunschweig: Vieweg, 23–51.
- Hilbert, D., and S. Cohn-Vossen ([1932] 1983), *Geometry and the Imagination*. Reprint. Translated by P. Nemenyi. Originally published as *Anschauliche Geometrie* (Berlin: Springer). New York: Chelsea.
- Isham, C. (1993), “Canonical Quantum Gravity and the Problem of Time”, in L. Ibort and M. Rodríguez (eds.), *Integrable Systems, Quantum Groups and Quantum Field Theories*. London: Kluwer, 157–288, arXiv:gr-qc/9210011v1.
- Kuchař, K. (1992), “Time and Interpretations of Quantum Gravity”, in G. Kunstatter, D. Vincent, and J. Williams (eds.), *General Relativity and Relativistic Astrophysics*. Singapore: Scientific, 211–314.
- Misner, C., K. Thorne, and J. Wheeler (1973), *Gravitation*. San Francisco: Freeman.
- Newton, I. ([1686] 1999), *Philosophiæ Naturalis Principia Mathematica*. Berkeley: University of California Press.

- Norton, J. (1989), "How Einstein Found His Field Equations, 1912–1915", in D. Howard and J. Stachel (eds.), *Einstein and the History of General Relativity*, Einstein Studies, vol. 1. Berlin: Birkhäuser, 101–159.
- (1993), "General Covariance and the Foundations of General Relativity: Eight Decades of Dispute", *Reports on Progress in Physics* 56: 791–858.
- Peirce, C. S. ([1898] 1993), *Reasoning and the Logic of Things: The Cambridge Conference Lectures*. Edited by K. Ketner. Cambridge, MA: Harvard University Press.
- Rovelli, C. (1991a), "Is There Incompatibility between the Ways Time Is Treated in General Relativity and in Standard Quantum Mechanics?", in A. Ashtekar and J. Stachel (eds.), *Conceptual Problems of Quantum Gravity*. Boston: Birkhäuser, 126–140.
- (1991b), "What Is Observable in Classical and Quantum Gravity?", *Classical and Quantum Gravity* 8: 297–316.
- (2000), "Quantum Spacetime: What Do We Know?", in C. Callender and N. Huggett (eds.), *Physics Meets Philosophy at the Quantum Scale*. Cambridge, Cambridge University Press, 101–124, arXiv:gr-qc/9903045v1.
- (2002), "Gps Observables in General Relativity", *Physical Review D* 65: 044017, arXiv:gr-qc/0110003v2.
- Sommerfeld, A. (1964), *Thermodynamics and Statistical Mechanics*. Lectures on Theoretical Physics, vol. 5. Translated by J. Kestin. Edited by F. Bopp and J. Meixner. New York: Academic Press.
- Stachel, J. ([1980] 1989), "Einstein's Search for General Covariance, 1912–1915", paper delivered to the Ninth International Conference on General Relativity and Gravitation at Friedrich-Schiller-Universität, Jena, 14–16 July 1980. Printed in D. Howard and J. Stachel (eds.), *Einstein and the History of General Relativity*, Einstein Studies, vol. 1. Berlin: Birkhäuser, 63–100.
- Stein, H. (1972), "On the Conceptual Structure of Quantum Mechanics", in R. Colodny (ed.), *Paradigms and Paradoxes: The Philosophical Challenge of the Quantum Domain*, University of Pittsburgh Series in the Philosophy of Science, vol. 5. Pittsburgh: University of Pittsburgh Press, 367–438.
- (1984), "The Everett Interpretation of Quantum Mechanics: Many Worlds or None?", *Noûs* 18: 635–652.
- (n.d.), "Physics and Philosophy Meet: The Strange Case of Poincaré", manuscript.
- Synge, J. (1960), *Relativity: The General Theory*. Amsterdam: North-Holland.
- Torre, C. (1993), "Gravitational Observables and Local Symmetries", *Physical Review D* 48: R2372–R2376.
- Unruh, W., and R. Wald (1989), "Time and the Interpretation of Canonical Quantum Gravity", *Physical Review D* 40: 2598–2614.
- Wald, R. (1984), *General Relativity*. Chicago: University of Chicago Press.