Chapter 9

Philosophy of Space–Time Physics

Craig Callender and Carl Hoefer

Philosophy of space–time physics, as opposed to the more general philosophy of space and time, is the philosophical investigation of special and general relativity. Relativity theory stimulated immediate and deep philosophical analysis, both because of its novel implications for the nature of space, time and matter, and because of more general questions philosophers have about the nature of its claims. With nearly one hundred years of sustained research to draw on, this chapter cannot hope to survey all the topics that have arisen, even all the major ones. Instead, we concentrate on four topics, two with a historical and philosophical pedigree, namely, relationalism and conventionalism, and two that arise in general relativity and cosmology, namely, singularities and the so-called horizon problem. This selection should give the reader a representative taste of the field as it stands today.

Many fascinating topics, however, will not be covered. Notable examples are the topics of time travel, presentism, supertasks, and the Lorentz interpretation of relativity. For up-to-date references and discussions of these topics, the reader can turn to, respectively, Arntzenius and Maudlin (2000), Savitt (2000), Earman and Norton (1996), and Brown and Pooley (2001).

Relationism, Substantivalism and Space–time

Perhaps the most fundamental question one has about space–time is: what *is* it, really? At one level, the answer is simple; at a deeper level, the answer is complex and the continuing subject of philosophers and physicists' struggle to obtain a plausible and intelligible understanding of space–time. In large measure, this struggle can be seen as a continuation of the classical dispute, sparked by the famous Leibniz–Clarke correspondence, between *relational* and *absolutist* conceptions of space – though the terms of the debate have turned and twisted dramatically in the twentieth century.

History

The general theory of relativity's (GTR) simple answer to our question is that space–time is

(a) a four-dimensional differentiable manifold M
(b) endowed with a semi-Riemannian metric $\boldsymbol{g}$ of signature (1,3)
(c) in which all events and material things (represented by stress-energy $\boldsymbol{T}$) are located, and
(d) in which $\boldsymbol{g}$ and $\boldsymbol{T}$ satisfy Einstein's field equations (EFE).

Had gravitational physics and scientific cosmology begun with Einstein's theory rather than Newton's, this simple answer might seem perfectly natural. Attempting to obtain a deeper understanding of the theory, philosophers struggle to understand GTR's space–time in terms of ideas found in previous theories, ideas whose roots lay in experience, metaphysics and Modern philosophy and physics. The questions that arise from these grounds *seem* to make good sense, independent of their roots: Is space–time a kind of *thing* which, though different from material things and energy-forms, is in some sense just as substantial and real? Do Einstein's equations describe a sort of causal interaction between space–time and matter; or is the relation one of reductive subsumption (and if so, which way)? Can space–time exist without any matter at all? Is motion purely relational in GTR, that is, always analyzable as a change in the relative configuration of bodies, or is it absolute, that is, always defined relative to some absolute structure? Or is it partly relative and partly absolute?

None of these questions receives a clear-cut answer from GTR, which is why absolutely inclined and relationally inclined thinkers can each find grist for their ontological mills in the theory. The complexity and ambiguity of the situation leads some philosophers to argue that it is pointless to try to impose the categories of seventeenth century metaphysics on a theory that has outgrown them (Rynasiewicz, 1996, 2000). Below, we briefly survey some of the key features of GTR that intrigue and frustrate philosophical interpretation, and return at the end to the question of whether the old categories and questions still have value.

No prior geometry In all earlier theories of mechanics and/or gravitation that contained definite doctrines about the nature of space and time (or space–time), space and time were taken as "absolute" structures, fixed and unchanging. As Earman (1989) shows, even the views of the traditional relationist thinkers involved some significant prior geometric and/or temporal structures. The Euclidean structure of space, for example, was universally assumed, as well as some absoluteness of temporal structure.[1]

Not so General Relativity. The background arena in GTR is just M, which can have any of a huge variety of topologies, and whose only "absolute" features are 4-dimensionality and continuity. The rest of the spatio-temporal properties, geometric and inertial and temporal, are all encoded by $\boldsymbol{g}$, which is not fixed or prior but rather variable under the EFE.[2] This looks extraordinarily promising from a relationist viewpoint: absolute space has finally been banished!

Or has it? Although absolute space or space–time, in the sense of a pre-defined and invariant background, is absent, it is not clear that this amounts to satisfaction of a relationist's desires. Motion has not become "purely relative" in any clear sense; rather, motion is defined relative to the metric, and the metric is by no means definable on the basis of relations between material things. In fact, the EFE turns matters the other way around: given the metric $\boldsymbol{g}$, the motions of material things (encoded in $\boldsymbol{T}$) are determined. If material processes affect the structure of spacetime, perhaps this is so much the better for a substantival view of GTR's space–times.

The differing roles of M and $\boldsymbol{g}$ correspond to two different strands of argument for substantivalism, which it will be useful to distinguish. The first strand notes that it is indispensable to the mathematical apparatus of GTR that it *start* from M and build the spatio-temporal structure $\boldsymbol{g}$ on it.[3] Then, invoking the Quinean doctrine that the real is that over which we ineliminably quantify in our best scientific theories, M is claimed to represent a real, existing manifold of space–time points in the world. The second strand looks at $\boldsymbol{g}$ itself, argues that it represents a real structure in the world not reducible to or derivable from material bodies and their relations, and concludes that we have a descendent of Newton's absolute space in GTR.

Manifold and Metric Above we have indicated that M is the only "fixed background" in GTR, and that only in the sense of dimensionality and continuity, not global shape. But motions (particularly acceleration, but also velocity and position in some models) are defined by $\boldsymbol{g}$. Which one, then, represents space–time itself? Or must we say it is a combination of both? These questions open a new can of interpretive worms.

A manifold is a collection of space–*time* points, not space points. In other words, the points do not have duration; each one is an ideal point-event, a representative of a spatial location at a single instant of time. They do not exist *over* time and hence serve as a structure against which motion may be defined, as Newton's space points did. If space–time substantivalism is understood as the claim that these points are substantial entities themselves, then the so-called *hole problem* arises (Earman and Norton, 1987). The general covariance of the EFE, interpreted in an active sense, allows one to take a given model $M_1 = \langle M, \boldsymbol{g}, \boldsymbol{T} \rangle$ and construct a second *via* an automorphism on the manifold: $M_2 = \langle M, \boldsymbol{g}^*, \boldsymbol{T}^* \rangle$ which also satisfies the EFE. Intuitively, think of M_2 as obtained from M_1 by sliding *both* the metric and matter fields around on the point-manifold (Figure 9.1).

If M_2 and M_1 agree or match for all events before a certain time t, but differ for some events afterward, then we have a form of indeterminism. Relative to our chosen substantial entities, space–time points considered as the elements of M, what happens at what space–time locations is radically undetermined. This can be

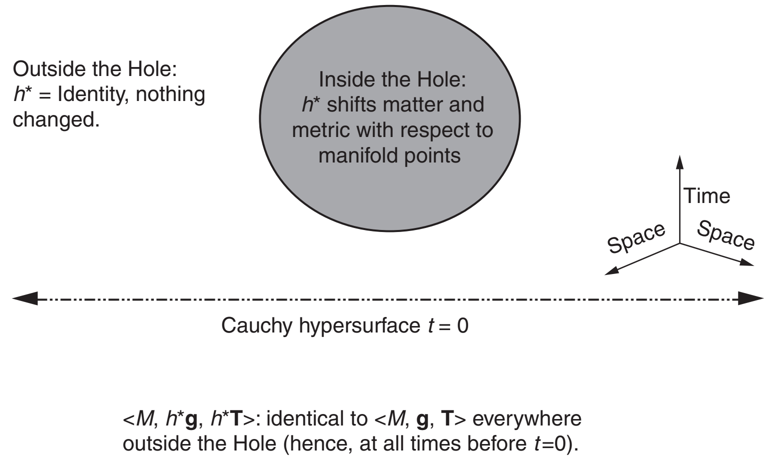

Figure 9.1 A hole diffeomorphism's effect

presented as an argument against the kind of substantivalism (*manifold* substantivalism) we started from.

Note however two points. First, this indeterminism is unobservable: M_1 and M_2 are qualitatively indistinguishable. Second, but relatedly, the hole problem assumes that the identities of the manifold points are given or specified, in some sense, independently of the material/observable processes occurring in spacetime (represented by $\boldsymbol{g}$ and $\boldsymbol{T}$). In fact, one way of thinking of a hole automorphism is as a (continuous) *permutation* of the points underlying physical processes, or as a *re-labeling* of the points. Not surprisingly, most responses to the hole argument have departed from these points, arguing that substantivalism can be reinterpreted in ways that do not lead to the apparent indeterminism.

Metric and matter Derivation of the metric of space–time from a (somehow!) antecedently given specification of the relational distribution of matter is a characteristic *Machian* ambition. Mach's *Science of Mechanics* (1989), at least as Einstein read it, proposed that inertia should be considered an effect of relative acceleration of a body with respect to other bodies – most notably, the "fixed stars." Transplanting this idea to the context of GTR seems to indicate having the inertial structure (which $\boldsymbol{g}$ determines) determined by the relational matter distribution. GTR does not seem to fulfill this idea in general. In some models (notably Friedman–Robertson–Walker (FRW) Big Bang models), this idea seems intuitively fulfilled. But making the idea both precise and satisfiable in GTR has proven difficult if not impossible, despite the efforts of outstanding physicists such as Einstein, Sciama, Wheeler, and Dicke. And as we noted above, superficially at least the determination relation seems to go the other way (from metric to matter).

Despite these difficulties, a Machian program for extending (or restricting) GTR has appealed to many thinkers. In addition to anti-absolutist prejudices, there are a couple of reasons for this. First, GTR does yield some non-Newtonian inertial effects of the kind Mach speculated on: the so-called "frame-dragging" effects. Second, it is difficult to take as a mere coincidence the fact that the FRW models, which seem most Machian intuitively, are also those that seem to best describe our cosmos.

And there are difficulties with taking the metric "field" as a substantial entity that either subsumes ordinary matter or is in causal interaction with it. The former idea, which can be thought of as "super-substantivalism", would involve extending our notion of the metric in an attempt to derive fine-grained properties of matter in terms of fine-grained perturbations (knots, singularities?) of the former. Attempts by Einstein and others along these lines have not led to notable successes.

The less ambitious idea of the metric and (ordinary) matter as ontological peers in mutual interaction faces challenges too. The EFE give a regularity, but in order to view the regularity causally, we would ideally like to be able to quantify the strength of the interaction, in terms of energy or momentum exchange. But the metric field's energy, if it exists at all, is poorly understood and very different from ordinary energy (Hoefer, 2000). Attempts through the end of the twentieth century to detect the most intuitively causal-looking interactions – absorption of gravitational wave energy – were uniformly negative.

Current work

The relationism/substantivalism issue was dominated through the late 1980s and 1990s by responses to the Earman and Norton hole argument. The argument has pushed philosophers who have more sympathy with substantivalist views than relationist views to make more precise their ontological claims (Maudlin, 1990; Butterfield, 1989; Stachel, 1993). Depending on whether they view substantivalism as primarily attractive due to the Quinean indispensability argument or rather the metric-based considerations, philosophers re-work their views in different ways.

The hole argument inspired those with relationist leanings to revive the idea, advocated by Reichenbach earlier in the twentieth century but effectively killed by Earman (1970) and Friedman (1983), that GTR can be interpreted as fully compatible with relationism. Teller (1991), Belot (1999) and Huggett (1999) are examples of this approach. What makes this position possible is

(a) focusing on the point-manifold-indispensability argument for substantivalism primarily, and
(b) taking a liberal attitude toward the idea of *relations* between material things.

If the manifold is taken as only representing the continuity, dimensionality and topology of space–time (as some substantivalists would agree anyway), then what's really indispensable is the metric. Can it be interpreted relationally? The philosophers who argue that it can are *not* claiming a Machian reduction of metrical structure to material relations. Instead, they claim that the metric itself can be interpreted as giving the structure of actual *and possible* spatiotemporal relations between material things. g is not a thing or substance. Where matter is present, it is crucial to the definition of local standards of acceleration and non-acceleration; the EFE record just this relationship. In many ways, the desires of traditional relationists (especially Leibniz, Huygens and Mach) are – arguably – met by GTR when interpreted this way.

Current work has served to clarify the various types of substantivalist view that may be brought forth, and the strengths and weaknesses of them. To a lesser extent, relationist alternatives have also been clarified. Others feel unhappy about both alternatives, and their reasons stem from a conviction that the ontological categories of absolutism, substantivalism and relationism have no clear meanings in GTR and have thus outlived what usefulness they ever had.

Rynasiewicz has published two provocative papers on this topic. His 1996 paper argues that the categories of absolute and relational simply do not apply in GTR, so it is a mug's game trying to see which one "wins" in that theory. Tracing relational and absolutist ideas from Descartes through to Einstein and Lorentz, the core of his case is that the metric field of GTR is a bit like a Cartesian subtle matter and a bit like Newton's absolute space, but in the end not enough like either (though it is a lot like *ether*). In his 2000 paper, he does a similar historical/conceptual analysis of the notions of absolute and relative motion, concluding that the notions are impossible to define in GTR. While it is possible to mount counterarguments in defense of the traditional notions (Hoefer, 1998), it is impossible to deny that GTR is an awkward theory to comprehend using traditional concepts of space and time.

Robert Disalle (1994, pp. 278–9) argues along similar lines. He offers a positive way to understand space–time after we have freed ourselves from the outmoded categories. In his 1995, he argues that a chief mistake of the tradition is thinking of space–time structure as an entity that we postulate to *causally explain* phenomena of motion. It can't do the job of explaining motions because it is simply an expression of the *facts* about those motions – when certain coordinating definitions are chosen to relate spatio-temporal concepts with physical measurements and processes. The point is nicely made by analogy. When pre-nineteenth century thinkers asserted the Euclidean nature of space, they were claiming that observations of length, angle and distance will always conform to the rules of Euclid's geometry. But saying space is Euclidean is not giving a *causal explanation* of rulers and compasses behaving as they do, and it is not the postulation of a new, substantial "thing" in which rulers etc. are embedded. Nor is it, however, a claim that all spatial facts are reducible to observables or measurement outcomes. One can, says Disalle, be a realist about space (or space–time)'s structure, without making the mistake of inappropriate reification.

Future work

The notions of relational and substantival spacetime may have reached a sort of impasse when it comes to the interpretation of GTR's overall structure, as presented in entry-level textbooks. This hardly means that we have an adequate understanding of space–time's ontology, a comfortable resting place for philosophical curiosity. A search of the abstracts of recent work in the foundations of GTR and quantum gravity will show numerous occurrences of words like *relational* and *absolute*, *Leibniz* and *Mach*. This is because philosophers and physicists alike still want to deepen their understanding of the world's ontology. There is still important work that can be done on classical GTR. For example, what is the status of energy conservation laws? Does matter–energy really get exchanged between ordinary matter and empty space–time? How might relationalists understand parity nonconservation? Are there Machian replacements for or restrictions of GTR that are observationally equivalent over the standard range of current tests? (See Barbour and Bertotti (1982) and Barbour (1999).)

Conventionalism about Space–time

Some of the most basic principles of science (and perhaps mathematics) seem to be true as a matter of definitional choice. They are not *quite* purely analytic or trivial; they can not be demonstrated true simply on the basis of prior stipulative definitions and logical rules. Further, incompatible-looking alternative principles are conceivable, even though we may not be able to see how a useful framework could be built on them. Such principles are often held to be true by *convention*.

One example in mathematics is the famous parallel postulate of Euclidean geometry. Physical examples are less common and typically fraught with controversy. Perhaps Newton's famous 2nd law, $\mathbf{F} = m\mathbf{a}$, is an example. This may be thought a poor choice, for surely, as the center of his mechanics, the 2nd law is far from true by definition. But in the Newtonian paradigm, the 2nd law served as ultimate arbiter of the questions

(a) whether a force existed on a given object; and
(b) if so, what its magnitude was.

Any failure of the **a** of an object to conform to expectation was grounds for assuming that an unknown or unexpected force was at work, not grounds for questioning the 2nd law.

Of course, there is no guarantee that one can *always* maintain any putative conventional truth, come what may. Rather, one can usually imagine (or experimentally find, as in the example at hand) circumstances in which unbearable tensions arise in our conceptual frameworks from the insistence on retention of the conventional principle, and one is effectively forced to give it up. (Duhem (1954) gives a classic discussion of these matters.) If this is right, then the original claim of conventionality looks like something of an exaggeration. Are there in fact *any* choices in the creation of adequate physical theory that are genuinely free, conventional choices (as, e.g. choice of units is), without being completely trivial (as, again, choice of units is)? Many philosophers have thought that space–time structures give us true examples of such conventionality.

History

Before the eighteenth century all philosophers of nature assumed the Euclidean structure of space; it was thought that Euclid's axioms were true *a priori*. The work of Lobachevsky, Riemann and Gauss destroyed this belief; they demonstrated, first, that consistent non-Euclidean constant-curvature geometries were possible, and later that even variably curved space was possible. It was also apparent that our experience of the world could not rule out these new geometries, at least in the large. But what, exactly, does it mean to say that space is Euclidean or Riemannian? A naïve-realist interpretation can, of course, be given: there exists a *thing*, space, it has an intrinsic structure, and that structure conforms to Euclid's axioms. But some philosophers – especially empiricists such as Reichenbach – worried about how space is related to observable properties. These philosophers realized that our physical theories always contain assumptions or postulates that *coordinate* physical phenomena with spatial and temporal structures. Light rays in empty space travel in straight lines, for example; rigid bodies moved through empty space over a closed path have the same true length afterward as before; and so on. So-called *axioms of coordination* are needed to give meaning and testability to claims about the geometry of space.

The need for axioms of coordination seems to make space for conventionalism. For suppose that, under our old axioms of coordination, evidence starts to accumulate that points toward a non-Euclidean space (triangles made by light rays having angles summing to less than 180°, for example). We could change our view of the geometry of space; but equally well, say conventionalists, we could change the axioms of coordination. By eliminating the postulate that light rays in empty space travel in straight lines (perhaps positing some "universal force" that affects such rays), we could continue to hold that the structure of space itself is Euclidean. According to the strongest sorts of conventionalism, this preservation of a conventionally chosen geometry can *always* be done, come what may. Poincaré (1952) defended the conventionality of Euclidean geometry; but he also made an empirical conjecture, now regarded as false: that it would always be *simpler* to construct mechanics on assumption of Euclidean geometry.

Discussions of conventionalism took a dramatic turn because of the work of Einstein. With its variably curved space–time, GTR posed new challenges and opportunities for both sides on the conventionality of geometry. Cassirer, Schlick, Reichenbach, and Grünbaum are some notable figures of twentieth century philosophy who argued for the conventionality of space–time's geometry in the context of GTR. Recent scholars have tended to be skeptical that any non-trivial conventionalist thesis is tenable in GTR; Friedman (1983) and Nerlich (1994) are prominent examples here.

But it was in 1905, rather than 1915, that Einstein gave the greatest wind to conventionalist's sails. In the astounding first few pages of the paper that introduced the special theory of relativity (STR), Einstein overthrew the Newtonian view of space–time structure, and in passing, noted that *part* of the structure with which he intended to replace it had to be chosen by convention. That part was *simultaneity.* Einstein investigated the operational significance of a claim that two events at different locations happen simultaneously, and discovered that it must be defined in terms of some clock synchronization procedure. The obvious choice for such a procedure was to use light-signals: send a signal at event *A* from observer 1, have it be received and reflected back by observer 2 (at rest relative to 1) at event *B*, and then received by 1 again at event *C*. The event *B* is then simultaneous with an event *E*, temporally mid-way between *A* and *C* (Figure 9.2).

Or is it? To suppose so is to assume that the velocity of light on the trip from *A* to *B* is the same as its velocity from *B* to *C* (or, more generally, that light has

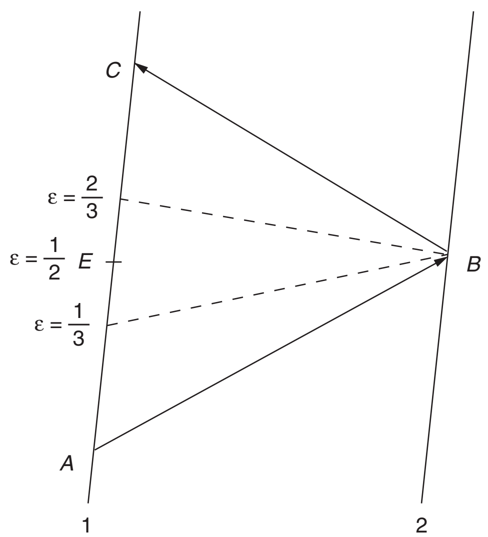

Figure 9.2 Simultaneity conventions in STR

the same velocity in a given frame in all directions). This seems like a very good thing to assume. But can it be verified? Einstein thought not. All ways of directly measuring the one-way velocity of light seemed to require *first* having synchronized clocks at separated locations. But if this is right, we are going in circles: we need to know light's one-way velocity to properly synchronize distant clocks, but to know that velocity, we need antecedently synchronized clocks.

To break the circle, Einstein thought we needed simply to adopt a conventional choice: we *decide* that event E is simultaneous with B (i.e. that light's velocity is uniform and direction-independent). Other choices are clearly possible, at least for the purposes of developing the dynamics and kinematics of STR. Following Reichenbach, these are synchronizations with $\varepsilon \neq \frac{1}{2}$ (ε being the proportion of the round-trip time taken on the outbound leg only). Adopting one of these choices is a recipe for calculational misery of a very pointless kind. But the Einstein of 1905, and many philosophers of an operationalist/verificationist bent since then, thought that such a choice cannot be criticized as *wrong*. Ultimately, they say, distant simultaneity is not only frame-relative, but partly conventional.

Taking up the challenge of establishing a one-way velocity for light, Ellis and Bowman (1967) argued that *slow clock transport* offers a means of synchronizing distant clocks that is independent of the velocity of light. In STR, when a clock is accelerated from rest in a given frame up to some constant velocity, then decelerated to rest again at a distant location, there are time-dilation effects that prevent us from regarding the clock as having remained in synch with clocks at its starting point. And calculation of the size of the effect depends on having established a distant-simultaneity convention (i.e. a choice of ε). So it looks as though carrying a clock from observer 1 to observer 2 will not let us break the circle. But Ellis and Bowman noted that the time dilation effect tends to zero as clock velocity goes to zero, and this is independent of ε-synchronization. Therefore, an "infinitely slowly" transported clock allows us to establish distant synchrony, and measure light's one-way velocity. Conventionalists were not persuaded, and the outcome of the fierce debate provoked by Ellis and Bowman's paper was not clear.

In 1977, David Malament took up the conventionalist challenge from a different perspective. One way of interpreting the claim of conventionalists such as Grünbaum is this: the observable *causal structure* of events in an STR-world does not suffice to determine a unique frame-dependent simultaneity choice. By "causal structure" we mean the network of causal connections between events; loosely speaking, any two events are causally connectable if they *could* be connected by a material process or light-signal. In STR, the "conformal structure" or light-cone structure at all points is the idealization of this causal structure. It determines, from a given event, what events could be causally connected to it (toward the past or toward the future). Grünbaum and others believed that the causal structure of space–time by no means singles out any preferred way of cutting up space–time into "simultaneity slices".

Malament showed that, in an important sense, they were wrong. The causal/conformal structure of Minkowski space–time *does* pick out a unique frame-relative foliation of events into simultaneity slices. Or rather, more precisely, the conformal structure suffices to determine a unique relation of *orthogonality.* If we think of an ε-choice as the choice of how to make simultaneity slices relative to an observer in a given frame, then Malament showed that the conformal structure is sufficient to define a unique, orthogonal foliation that corresponds to Einstein's $\varepsilon = \frac{1}{2}$ choice. For many philosophers, this result marked the end of the debate over conventionality of simultaneity. (But see Janis, 1983 and Redhead, 1993 for conventionalist responses.)

Current work

A recent paper by Sarkar and Stachel (1999) tries to re-open the issue of conformal structure and simultaneity relations. Stachel and Sarkar note that one of Malament's assumptions was that the causal connectability relation is taken as *time-symmetric*, i.e. that it does not distinguish past-future from future-past directions of connection. They argue that it is possible to distinguish the backward from forward light cones using only the causal-connectability relation Malament starts from. If this is granted, and we do not impose the condition that any causally-definable relation must be time-symmetric, then the uniqueness result Malament proved fails. Many different cone-shaped foliations become definable. Stachel and Sarkar advocate the backward-lightcone surface as an alternative simultaneity surface choice that could be made. It remains true, however, that only the genuine orthogonality relation ($\varepsilon = \frac{1}{2}$) is transitive and location-independent. These are two of the core features of classical simultaneity. To put forward Stachel and Sarkar's alternative relation as a genuine candidate for a distant-synchrony relation is therefore, at best, awkward and out of line with core intuitions about simultaneity.

Still, many philosophers of physics feel dissatisfied with even this much of a concession to conventionalism. They suspect that, even if it may have been in some sense possible to do physics with $\varepsilon \neq \frac{1}{2}$ in 1905, more recent quantum field theory has surely ruled that out. Zangari (1994) argued that the mathematics of spinor fields in Minkowski space–time – used in describing spin-$\frac{1}{2}$ particles, for example – is only consistent in frames with standard synchrony. Gunn and Vetharaniam (1995) claimed that Zangari was mistaken, and that using a different formalism, the Dirac equation could be derived in a framework including $\varepsilon \neq \frac{1}{2}$ frames. Karakostas (1997) has argued that both of the preceding authors' arguments are flawed, though Zangari's main claim is correct. And most recently, Bain (2000) argues that none of these authors has it exactly right. There is *always* a way to do physics using arbitrary coordinates (including those corresponding to non-standard simultaneity choices); but whether that amounts to the conventionality of simultaneity in an interesting sense remains a tricky question.

In trying to see one's way through the dense thicket of technical claims and counter-claims in these papers, it helps to fall back on the notion of *general covariance*. Kretschmann hypothesized in 1918 that *any* physical theory could be expressed in a generally covariant form, i.e. in a form that is valid in arbitrary coordinates. Nonstandard-synchrony frames do correspond to coordinate systems allowed under general covariance. Karakostas does not deny Kretschmann's claim. Instead, he notes that generally covariant treatments of spinor fields can be done, but they have to introduce a geometric structure (a "frame" or "*vierbein* field") that effectively picks out the orthogonal (= standard simultaneity) direction for a given observer in a given frame. This is a typical sort of move when theories with absolute space–time structures are given in a generally covariant form. Geometric objects or fields replace privileged coordinates or frames, but the "absoluteness" is only shifted, not removed. In the case of spinor fields, it seems that *something* that effectively encodes the Einstein-standard synchrony relation is mathematically necessary. Can the conventionalist respond by claiming that this necessary structure is, withal, not a *simultaneity* structure? Bain claims that she can; for spinor fields have nothing to do with rods and clocks, and the measurement of light's one-way velocity – i.e. with the original point conventionalists made.

Conventionalist claims – concerning both geometry and simultaneity – seem to be constantly in danger of collapsing into triviality: the trivial claim that, if we are mathematically clever and not afraid of pointless hard work, we can choose any perverse sort of coordinate system we like, and then claim that the coordinates reflect the geometric/simultaneity relations we have "chosen." Perhaps we can do this; but to suppose that this amounts to a genuine choice of spatio-temporal facts is to be somewhat disingenuous about the content of such facts. To be sure, axioms of coordination are needed to link pure geometric concepts to observable phenomena. But the axioms we choose are themselves constrained in many ways by the need to cohere with further practices and metaphysical assumptions. In practice, these constraints seem to fully determine, or even over-determine our "choices" regarding geometry. What keeps the debate concerning conventionality of *simultaneity* alive is the way in which our "conventional choices" play only a completely trivial role *qua* axioms of coordination. Just as one can do physics with any choice of ε, one can also do physics without *any* choice of clock synchronization.

Future work

Relativity theory (STR and GTR) provides the natural home for at least limited forms of conventionalism, though it remains a subject of dispute just how significant the conventionality is. The work of Karakostas, Bain and others points in the direction future work on these topics will take: toward new physics. One would also expect that advances in the general methodology of science will continue to bear on these issues.

Black Holes and Singularities

Our best theories tell us that stars eventually run out of nuclear fuel. When they do so, they leave equilibrium and undergo gravitational collapse, ending as white dwarves if the collapsing core's mass $M < 1.4$ solar masses, as neutron stars if $5 > M > 1.4$, or as black holes if $M > 5$. Black holes are regions of space–time into which matter can enter but from which matter can not escape. Their end states are singularities, which for now we might associate with a "hole" in space–time or a point where the space–time metric "blows up" and is ill-defined. There is some astronomical evidence for the existence of black holes, and they are relevant to a number of questions that interest philosophers, such as whether time travel is possible and whether the past and future are finite (Weingard, 1979). However, we here focus on singularities, as they are more general since they can exist without black holes, and they also pose several different philosophical questions that are the subject of active research.

History

Singularities are hardly novel to GTR. The classical Coulomb field when combined with STR goes to infinity at points. Collapsing spherical dust clouds and other highly symmetric solutions provide examples of singularities in Newtonian gravitational theory. But singularities in GTR are especially puzzling, as we will see.

The existence of singularities to EFE was known from the theory's inception. Hilbert, for instance, wrote about the notorious singularities in the Schwartzchild solution as early as 1917. The line element of this solution has singularities at $r = 0$ and $r = 2M$. Einstein in 1918 worried about them only because he took them as a threat to Machianism. Singularities in the solutions to the field equations didn't cause general alarm for many more decades because they were not very well understood (Earman, 1999; Earman and Eisenstaedt, 1999). They were viewed as unacceptable pathologies, but it was assumed that they were defects of only certain models. From 1918 until the mid-1950s, it was not realized that the singularities in these space–times were "essential" in some sense. There were two other options.

First, a singularity might be merely a "coordinate singularity" and not a feature of the space–time. To illustrate the distinction, consider coordinizing a sphere. It is a theorem that no single coordinate system can cover the sphere without singularity. This represents a problem for the coordinate system, not the sphere. The sphere is a perfectly well-defined geometric object; moreover, there are ways of covering the sphere without singularity using two different coordinate patches. The Schwartzchild solution caused particular mischief in this regard during the first half of the twentieth century; it famously emerged that only one ($r = 0$) of its two apparent singularities is genuine – the "Schwartzchild radius" ($r = 2M$) is a mere artifact of the coordinates.

Second, like the singularities in classical gravitational theory, relativistic singularities might be due to an artificial symmetry of the solution. The singular nature of a solution of Newton's equations representing a perfectly spherical collapse of dust is real enough. It is no artifact of the coordinates chosen. But the feeling is, what chance is there that this is our world? Our world does not have its matter arranged like dust formed in a *perfect* sphere. Change the distribution somewhat and the singularity disappears. Why worry? Similarly, when it became clear that (for example) the Schwarzchild and Friedman solutions contained genuine singularities, the hope was that these arose from the artificial symmetries invoked; after all, the Schwartzchild solution represents the geometry exterior to a spherically symmetric massive body and the Friedman solutions represents a homogeneous and isotropic matter distribution. The singular solutions were hoped to be in some sense "measure zero" in the space of all the solutions of EFE.

These hopes were dashed by singularity theorems in the 1950s by Raychaudhuri and Komar, and especially by theorems in the 1960s and early 1970s by Penrose, Geroch and Hawking. These theorems appear to demonstrate that singularities are generic in space–times like ours. They assume what seem to be plausible conditions on the stress-energy of matter to force geodesics to cross; they then employ global conditions on the geometry to show that these geodesics terminate in a singularity.

These advances in the 1960s and 1970s were made possible in part by the new, minimal definition of a singularity. Without going into the details, a space–time is said to be singular according to these theorems just in case it contains a maximally extended timelike geodesic that terminates after the lapse of finite proper time. Briefly put, a space–time is singular iff it is timelike geodesically incomplete. (This definition can be extended to cover null and spacelike curves, and can be extended in other ways too – to so-called 'b-incompleteness' – but we will not go into this here.) The idea behind this definition is that it must be a serious fault of the space–time, one worthy of the name singularity, if the life of a freely falling immortal observer nevertheless terminates in a finite time.

However fruitful this definition, it has proved to be controversial, as have the significance of the singularity theorems. The current work in philosophy on these topics, largely driven by Earman (1995), focuses on these two questions.

Current work

This section focuses on the analysis of singularities. We concentrate on this topic not because we feel that it is any more important than other questions – indeed, we feel the opposite, that (for instance) the question of the significance of singularities for GTR is far more important – rather, we so concentrate because it is a necessary point of entry into the literature. One cannot successfully evaluate the significance of singularities without first knowing what they are.

Naively, one has the idea that a singularity is a hole in space–time surrounded by increasing tidal forces that destroy any approaching object. This picture cannot be correct for general relativistic space–times. The reason is simple: the singularities here are singularities in the metric space itself, so there is literally no location for a hole. General relativity requires a manifold with a smooth Lorentz metric, so by definition there are no locations where the metric is singular. Fields *on* space–time can be singular at points; but space–time itself has nowhere to be singular.

Following Geroch (1968), commentators have identified several quite distinct meanings of singularity. To name a few, and sparing details, consider the following conditions proposed for making a space–time singular:

(a) curvature blowup: a scalar curvature invariant, e.g. Ricci, tensor goes unbounded along a curve in space–time
(b) geodesic incompleteness: see above
(c) missing points: points are "missing" from a larger manifold, arising from the excision of singular points.

All three, we suppose, are involved in our intuitive idea of a space–time singularity. And for a Riemannian space, (b) and (c) are co-extensive. The Hopf–Rinow theorem states that, for connected surfaces, the conditions of being a complete metric space and being geodesically complete are equivalent. A metric space is complete if every Cauchy sequence of points in it converges to a point in that space. Intuitively, incompleteness is associated with missing points. For instance, the plane minus the origin, the surface $\mathbb{R}^2 - \{(0,0)\}$, is not complete because the Cauchy sequence $\{(1/n, 0)\}$ converges to a point excised from the plane. It is also not geodesically complete since there are no geodesics joining points $(-1,0)$ and $(1,0)$, so here we see a connection between geodesic incompleteness and missing points.

However, a relativistic space–time is not a Riemannian space, but a pseudo-Riemannian one, and the Hopf–Rinow theorem does not survive the change. None of the three definitions are co-extensive: the literature shows that while (c) implies (b), (b) does not imply (c); (a) implies (b), but (b) does not imply (a); and (a) seems to imply (c), but (c) does not imply (a). The official definition, (b), thus seems to act as a kind of symptom of the other two pathologies. Even here there are counterexamples. A curve might be incomplete even if the curvature is behaving normally, as happens in Curzon space–time; and as Misner shows, a curve might be incomplete even in a compact, and hence, complete and "hole-less", space–time.

It is of interest to see how hard it is to even make sense of definition (c). As mentioned above, a relativistic space–time has no room for singular points in the metric. Definition (c) would then have us look for the traces of an excised point, i.e. look for what is not there. How do you find points which are not on the space–time but which have been removed? Looking at the topology will not help since, in general, a variety of non-singular metrics can be put on any given topology (for instance, the Schwartzchild topology of $R^2 \times S^2$ is compatible with plenty of non-singular topologies). Although this way of understanding singularities is still active, it may be that the whole idea of a singularity as some localizable object is misleading.

Once we have an understanding of singularities in GTR, the next question to ask is about their significance. Do they "sow the seeds of GTR's demise" as is often alleged? Or are they harmless, perhaps even salutary, features of the theory? Earman (1996) provides an argument for tolerating singularities; but many physicists claim that they represent a genuine deficiency of the theory.

Future work

The topic of singularities is really a new one for philosophers of science. We can scarcely mention all the areas open to future endeavor. The majority of our focus below depends, perhaps naturally, on relatively recent ideas in physics.

Good arguments needed Earman (1996) pieces together and criticizes various arguments for the widespread belief that singularities sow the seeds of GTR's demise. A survey of the literature shows that there is a dearth of good argument supporting this belief. Can a good argument be articulated on behalf of this opinion that does not rely on misleading analogies with other pitfalls in the history of science?

Are there really singularities? The singularity theorems do not fall out as deductive consequences of the geometries of relativistic space–times. To say anything, the stress-energy tensor must be specified, and, in fact, all the theorems use one or another energy condition. The so-called weak energy condition, for instance, states that the energy density as measured by any observer is non-negative. But, is it reasonable to suppose these hold? The philosopher Mattingly (2000) sounds a note of skepticism, pointing out that various classical scalar fields and quantum fields violate all the conventional energy conditions. Even if Mattingly's skepticism is not vindicated, a better understanding of the relation between the energy conditions and real physical fields is certainly worth having.

Quantum singularities Philosophers may also wish to cast critical eyes over some of the methods suggested for escaping singularities with quantum mechanics. It is sometimes said that one should define a quantum singularity as the vanishing of the expectation values for operators associated with the classical quantities that vanish at the classical singularity. Then it is pointed out that the radius of the universe, for example, can vanish in what is presumed to be an infinite density and curvature singularity, even though the expectation value does not vanish (Lemos, 1987). This is sometimes taken as showing that quantum mechanics smoothes over the classical singularity. But, is this really so? Callender and Weingard (1995), for example, argue that this quantum criterion for singular status fits poorly with some interpretations of quantum mechanics.

Black hole thermodynamics Hawking's (1975) "discovery" that a black hole will radiate like a blackbody strengthened Beckenstein's work supporting an analogy between classical thermodynamics and black holes. The field known as black hole thermodynamics was spawned, and there are now thought to be black hole counterparts to most of the concepts and laws of classical thermodynamics. For instance, the black hole's surface gravity divided by 2π acts like the temperature and its area divided by 4 acts like the entropy. Physicists enticed by this analogy often claim that it is no analogy at all, that black hole thermodynamics *is* thermodynamics and that (for instance) the surface gravity really *is* the temperature. The signifiance of these startling claims and the analogy are certainly worthy of investigation by philosophers of science.

Information loss A related topic is the black hole "information loss paradox" that arises from Hawking's (1975) result. Take a system in a quantum pure state and throw it into a black hole. Wait for the black hole to evaporate back to the mass it had when you injected the quantum system. Now you have a system of a black hole with mass M plus a thermal mixed state, whereas you started with a black hole with mass M plus a pure state. Apparently, you have a process that converts pure states into mixed states, which is a non-unitary transformation prohibited by quantum mechanics (such a transformation allows the sum of the probabilities of all possible measurement outcomes to not equal 1). See Belot et al. (1999) and Bokulich (2000) for some philosophical commentary on this topic.

Cosmic censorship Perhaps the biggest open question relevant to singularities and many other topics in gravitational physics is the status of Penrose's cosmic censorship hypothesis; for a recent assessment, see Penrose (1999). This hypothesis is often glossed as the claim that naked singularities cannot exist; that is, that singularities are shielded from view by an event horizon, as happens in spherically symmetric gravitational collapse. Naked singularities are unpleasant because they signal a breakdown in determinism and predictability. If a naked singularity occurs to our future, then no amount of information on the space-like hypersurface we inhabit now will suffice to allow a determination of what happens at all future points. Singularities are, intuitively, holes out from which anything might pop. A singularity that we can see means we might see anything in the future, since the causal past will not sufficiently constrain the singularity.

Stated as the claim that naked singularities cannot exist, however, the hypothesis is clearly false, since there are plenty of relativistic space–times that violate it. Though formulated in a variety of non-equivalent ways (Earman, 1995, ch. 3), it is common to speak of weak and strong versions of the claim. Weak cosmic censorship holds that gravitational collapse from regular initial conditions never creates a space–time singularity visible to distant observers, i.e. any singularity that forms must be hidden within a black hole. Strong cosmic censorship holds that any such singularity is never visible to any observer at all, even someone close to it. By "regular initial data" we mean that the space–times are stable with respect to small changes in the initial data. Elaborating this definition further obviously requires some care.

The only consensus on the topic of cosmic censorship is that the hypothesis is both important and not yet proven true or false. Regarding the latter, there are plenty of counter-examples to both formulations of the hypothesis, though especially to strong cosmic censorship; see, for example, Singh (1998). Current and future work will dwell on whether these examples really count. In the background there is, you might say, the "moral" cosmic censorship hypothesis, which claims that the only naked singularities that occur are Good ones, not Bad ones. The exact formulations of Good and Bad depend, as one would expect, on the character of the particular investigator: prudish investigators hope GTR doesn't offer so much as a hint of nakedness, whereas the more permissive will lower their standards.

It is important to know whether some version of the hypothesis is true. If a cosmic censor operates, then many topics dear to philosophers will be affected. A cosmic censor will naturally affect what kinds of singularities we can expect, and therefore influence the question of their significance for GTR (Earman, 1996); it would mean the time-travel permitting solutions of EFE such as Gödel's will not be allowed; that the possibility of spatial topology change (Callender and Weingard, 2000) will not be possible, and so on. And a lack of a cosmic censor will also bear on much of the physics of potential interest to philosophers, e.g. black hole thermodynamics hangs crucially on the existence of a cosmic censor.

There are also philosophical topics about cosmic censorship that need further exploration. To name two, what is meant by "not being allowed" in the statement of cosmic censorship and how do white holes (the time reverses of black holes) square with the hypothesis and the time symmetry of EFE?

Horizons and Uniformity

History

The observed isotropy (or near isotropy) and presumed homogeneity of our universe suggest that we inhabit a world whose large scale properties are given by the well-known Friedman standard model. In this model, the world "began" in a hot dense fireball known as the Big Bang, and matter has since expanded and cooled ever since. The rate of expansion and cooling depend on the equation of state for the cosmological fluid, and the ultimate fate of the universe (closed or open) depends on the curvature. Part of the corroboration of this model comes from the observed uniformity of the cosmic microwave background radiation (CMBR). Neglecting some recently detected small inhomogeneities (which are themselves not defects but harmonic oscillations expected in some Big Bang models), these observations show that the temperature of this radiation is uniform to at least one part in 10,000 in every direction we look.

When coupled with the Friedman model, the uniformity of the CMBR produces a puzzle. To see this, we need to resolve an apparent contradiction between one's naïve view of the Big Bang singularity with the fact that in the Friedman model not all bodies can communicate with each other, even merely a fraction of a second after the Big Bang. Consider two nearby co-moving particles at the present time. The scaling factor, a, is the distance between the particles, say one light-second. Because the universe is expanding $da/dt > 0$. Now one would expect that, since $a \rightarrow 0$ as $t \rightarrow 0$ for all the particles, any particle could have been in causal contact with any other at the Big Bang. Since they are all "squashed together," a light pulse from one could always make it to any other particle in the universe. This is not so.

First off, there is *no* point on the manifold where $t = 0$ and $a = 0$; this point is not well-defined and it is not clear, anyway, that all the particles in the universe occupying the same point really makes sense. So this "point" does not count. But now it is a question of how fast the worldlines are accelerating away from each other and whether light signals from each can reach all the others. Will light emanating from a body "just after" the Big Bang singularity be able to reach an arbitrary body X by time t, where t is some significantly long time, possibly (in a closed universe) the end of time? In general, the answer is No, for there are some (realistic) values of the expansion parameter that do not allow the light signal to catch up to X by t. The space–time curvature is the key here. Imagine that you and a friend are traveling in opposite directions on a flat plane. Assuming nothing travels faster-than-light, can you evade a light pulse sent out in your direction from your friend? No: though you may give it a good run for its money, eventually it will catch you if the universe is open. Now imagine that you're moving on a balloon and the balloon is being quickly inflated. Then, depending on the speed of inflation and your velocity, you may well be able to escape the light signal, possibly for all time.

The curvature due to the expansion and deceleration causes the worldlines of galaxies to curve. In two spatial dimensions, our past lightcone becomes pear-shaped rather than triangular (Figure 9.3).

Note that due to this curvature we cannot "see" the entire Big Bang. A useful picture of the causal situation emerges if we "straighten out" the curvature, much as we do when we use a Mercator projection when we draw a flat picture of the earth (Figure 9.4).

Here the top of the large triangle is the point we are at right now, and the two shaded triangles are the past null cones of two points, separated by an angle A,

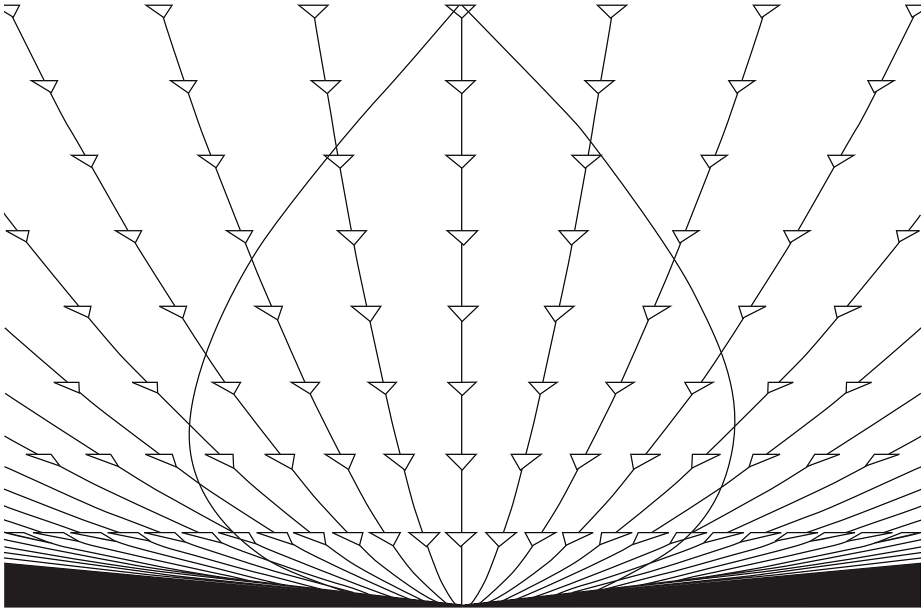

Figure 9.3 Our past light cone

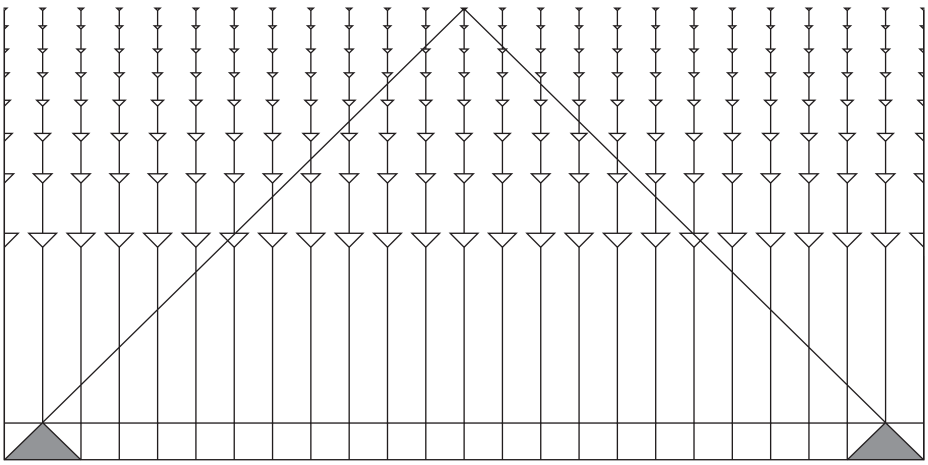

Figure 9.4 Straightening out the curvature

that we can see in our past. If A is sufficiently great the shaded regions do not intersect. But since the past null cone of a point represents all the points with which it might have had causal contact, this means that no point in the shaded regions could have had causal contact with each other (ignoring the possibility of faster-than-light travel). A *particle horizon* is defined as the maximum coordinate distance that one can see from a given point in space–time. From the diagram, one can see that the points at the top of the shaded regions have horizons that preclude them from seeing each other's pasts.

The puzzle about horizons arises from the fact that a Friedman model like that pictured can be said to fairly represent our universe, where the shaded regions are points in our past where matter decoupled from radiation. Since they share no common causal past, this means that they have no mechanism in common that will make the microwave radiation's temperature the same. How, then, did they arrive at the same temperature? It seems that, short of denying that the early universe can be approximately represented by a Friedman model, the only answer is that the universe was "born" in a highly isotropic and homogenous state. This necessary special initial state is the cause of the horizon problem.

Current work

In physics, the main response to this puzzle is to change the physics of expansion. Though there are other responses, the one known as inflation is almost universally maintained. In inflationary scenarios, the standard Friedman expansion is jettisoned in an early epoch in favor of a period of exponential expansion; the universe then undergoes a phase transition that slows it down back to the more moderate Friedman expansion. The details of this period vary with different proposals (there are more than fifty). Inflation does not remove particle horizons; instead, it increases the size of each point's past null cone so that pairs will overlap. The shaded past lightcones in the diagram would intersect while remaining proper subsets of each other. The hope is that the common causal past between two points will be large enough so that it accounts for their uniform temperature.

Work by Penrose (1989), Earman (1995) and Earman and Mosterin (1999) have severely criticized inflation for failing to deliver on its original promises. The theory, they say, does not rid cosmology of the need for special initial conditions to explain the apparent uniformity of the cosmic background radiation, nor does it enjoy much in the way of empirical success.

Future work

The horizon problem shares some general features with other well-known "problems" in physics. The problem of the direction of time (well, one of them) asks for an explanation of the thermodynamic arrow of time and ends up requiring the postulation of a very special initial condition of low entropy (Price, 1996). Philosophically, it is non-trivial whether requiring "special" boundary conditions is a genuine defect of a theory. For the situation with entropy and the direction of time, many do not see the special posit as a genuine failing of the theory to provide a scientific explanation (Callender, 1997); in the cosmological case with horizons, however, it is orthodoxy now that it is a genuine failing of the standard model that it cannot explain the uniformity of the cosmic background radiation. But is it? The failing is certainly not one of empirical disconfirmation, since given the "special" initial conditions the model is empirically adequate – we even get a deterministic explanation through time of why we see the features we do. Earman and Mosterin do much to criticize inflation as a solution to this problem, but the larger issue, common to this topic and others – whether there really is a problem here at all – is left open.

A related issue is whether the notion of "specialness" can be sharpened. In the cosmological case considered here, it is especially problematic to specify in exactly what sense the boundary conditions are "special," as Penrose (1989) emphasizes. By contrast, in the thermodynamic case this is somewhat clearer since one is talking about a statistical theory (statistical mechanics) equipped with a standard probability measure with respect to which the needed initial conditions do indeed occupy small measure. To be sure, there are problems in this case too in justifying this "natural" probability measure, but they appear to pale in comparison to the problems of defining a probability for cosmic inflation.

Finally, the question of whether there was inflation will probably ultimately be decided by observation and experiment rather than philosophical argument. Recent and future improvements in observational cosmology (e.g. CMB measurements, measurements of type 1a supernovae at high redshifts) have opened up the possibility of empirical support or disconfirmation of some inflation scenarios. The epistemology of this optimistic, burgeoning branch of physics is yet another field ripe for philosophical analysis.

Conclusion

The specter hanging over all future work in this field is quantum gravity. It is widely believed that general relativity is inconsistent with quantum field theory; "quantum gravity" is the research program that seeks a third theory that unifies, or at least makes consistent, these two theories. Though no such theory yet exists, there are some well-developed approaches such as string theory and canonical quantum gravity as well as some less developed theories such as topological quantum field theory and twister theory; for a philosophical slant, see Callender and Huggett (2001) and references therein. We believe it is fair to say that all of these theories are quite radical in their implications for space and time. If any of them, or remotely similar descendents, succeed, they may well have dramatic consequences for virtually all of the issues discussed above.

Acknowledgments

Figures 9.3 and 9.4 are reproduced, with permission, from Ned Wright's Cosmology Tutorial (http://www.astro.ucla.edu/%7Ewright/cosmolog.htm).

Notes

1 Whether this really spoils the relational ambitions of a hypothetical physics set in them is a difficult question. If space's structure is nothing more than what is implied by all the distance/angle relations between physical things – as one form of relationism holds – then the geometry of space must be an empirical matter, not something we can fix *a priori*.
2 Actually, there is a further element of absoluteness in GTR, namely the demand that the metric field have signature (1,3) and hence be "locally" like Minkowski space–time. See Brown (1997) for an illuminating discussion of this posit.
3 In fact, Hartry Field's (1980) uses the point-manifold, interpreted realistically, to eliminate platonic entities from the mathematics of physics. So, from his perspective, the manifold is not just indispensable in GTR, but in all of physical science.

References

Arntzenius, F. and Maudlin, T. (2000): "Time Travel and Modern Physics," *Stanford Online Encyclopedia of Philosophy*, http://plato.stanford.edu/entries/time-travel-phys.

Bain, J. (2000): "The Coordinate-Independent 2-Component Spinor Formalism and the Conventionality of Simultaneity," *Studies in History and Philosophy of Modern Physics*, 31B, 201–26.

Barbour, J. (1999): *The End of Time: The Next Revolution in our Understanding of the Universe*. London: Weidenfeld & Nicolson.

Barbour, J. and Bertotti, B. (1982): "Mach's Principle and the Structure of Dynamical Theories," *Proceedings of the Royal Society of London* A382, 295–306.

Belot, G. (1999): "Rehabilitating Relationalism," *International Studies in the Philosophy of Science*, 13, 35–52.

Belot, G. Earman, J. and Ruetsche, L. (1999): "The Hawking Information Loss Paradox: The Anatomy of a Controversy," *British Journal for the Philosophy of Science*, 50, 189–229.

Bokulich, P. (2000): "Black Hole Remnants and Classical vs. Quantum Gravity," available at the *PSA 2000* program page http://hypatia.ss.uci.edu/lps/psa2k/program.html.

Brown, H. (1997): "On the Role of Special Relativity in General Relativity," *International Studies in the Philosophy of Science*, 11, 67–80.

Brown, H. and Pooley, O. (2001): "The Origin of the Space–Time Metric: Bell's 'Lorentzian Pedagogy' and its Significance in General Relativity," in C. Callender and N. Huggett (eds.), *Physics Meets Philosophy at the Planck Scale*, Cambridge University Press, 256–74. Also available at gr-qc/9908048.

Butterfield, J. (1989): "The Hole Truth," *British Journal for the Philosophy of Science*, 40, 1–28.

Callender, C. (1997): "What is 'The Problem of the Direction of Time'?" *Philosophy of Science*, Supplement, 63(2), 223–34.

Callender, C. and Huggett, N. (eds.) (2001): *Physics Meets Philosophy at the Planck Scale*. New York: Cambridge University Press.

Callender, C. and Weingard, R. (1995): "Bohmian Cosmology and the Quantum Smearing of the Initial Singularity," *Physics Letters*, A208, 59–61.

Callender, C. and Weingard, R. (2000): "Topology Change and the Unity of Space," *Studies in History and Philosophy of Science*, 31B, 227–46.

Disalle, R. (1994): "On Dynamics, Indiscernibility and SpaceTime Ontology," *British Journal for the Philosophy of Science*, 45, 265–87.

Disalle, R. (1995): "SpaceTime Theory as Physical Geometry," *Erkenntnis*, 42, 317–37.

Duhem, P. (1954): *The Aim and Structure of Physical Theory*, trans. P. Wiener, Princeton, NJ: Princeton U.P.

Earman, J. (1970): "Who's Afraid of Absolute Space?" *Australasian Journal of Philosophy*, 48, 287–319.

Earman, J. (1989): *World enough and space–time: absolute versus relational theories of space and time*. Cambridge, Mass.: MIT Press.

Earman, J. (1995): *Bangs, Crunches, Whimpers and Shrieks: Singularities and Acausalities in Relativistic Spacetimes*. New York: Oxford University Press.

Earman, J. (1996): "Tolerance for Space–time Singularities," *Foundations of Physics*, 26(5), 623–40.

Earman, J. (1999): "The Penrose–Hawking Singularity Theorems: History and Implications," in H. Goemmer, J. Renn, J. Ritter and T. Sauer (eds.), *The Expanding Worlds of General Relativity*, Einstein Studies, vol. 7. Boston: Birkhauser.

Earman, J. and Eisenstaedt, J. (1999): "Einstein on Singularities," *Studies in History and Philosophy of Modern Physics*, 30, 185–235.

Earman, J. and Mosterin, J. (1999): "Inflation," *Philosophy of Science*, 66, 1–49.

Earman, J. and Norton, J. (1987): "What Price Space–Time Substantivalism? The Hole Story," *British Journal for the Philosophy of Science*, 38, 515–25.

Earman, J. and Norton, J. (1996): "Infinite Pains: The Trouble with Supertasks," in A. Morton and S. Stich (eds.), *Benacerraf and his Critics*, Cambridge, MA: Blackwells, 231–61.

Ellis, B. and Bowman, P. (1967): "Conventionality in Distant Simultaneity," *Philosophy of Science*, 34, 116–36.

Field, H. (1980): *Science without Numbers: A Defence of Nominalism*. Oxford: Blackwell.

Friedman, M. (1983): *Foundations of Space–Time Theories: Relativistic Physics and Philosophy of Science*. Princeton, N.J.: Princeton University Press.

Geroch. R. (1968): "What is a Singularity in General Relativity?" *Annals of Physics*, 48, 526–40.

Gunn, D. and Vetharaniam, I. (1995): "Relativistic Quantum Mechanics and the Conventionality of Simultaneity," *Philosophy of Science*, 62, 599–608.

Hawking, S. W. (1975): "Particle Creation by Black Holes," *Communications in Mathematical Physics*, 43, 199–220.

Hoefer, C. (1998): "Absolute versus Relation Space–Time: For Better or Worse, the Debate Goes On," *British Journal for the Philosophy of Science*, 49, 451–67.

Hoefer, C. (2000): "Energy Conservation in GTR," *Studies in the History and Philosophy of Modern Physics*, 31, 187–99.

Huggett, N. (1999): "Why Manifold Substantivalism is Probably Not a Consequence of Classical Mechanics," *International Studies in the Philosophy of Science*, 13, 17–34.

Janis, A. (1983): "Simultaneity and Conventionality", in R. Cohen (ed.), *Physics, Philosophy and Psychoanalysis*, Reidel: Dordrecht, 101–10.

Karakostas, V. (1997): "The Conventionality of Simultaneity in the Light of the Spinor Representation of the Lorentz Group," *Studies in the History and Philosophy of Modern Physics*, 28, 249–67.

Lemos, N. (1987): "Conservation of Probability and Quantum Cosmological Singularities," *Physical Review*, D36, 2364.

Mach, E. (1989): *The Science of Mechanics: A Critical and Historical Account of its Development*, trans. T. McCormack, LaSalle, Ill: Open Court.

Malament, D. (1977): "Causal Theories of Time and the Conventionality of Simultaneity," *Noûs*, 11, 293–300.

Mattingly, J. (2000): "Singularities and Scalar Fields: Matter Theory and General Relativity," available at the *PSA 2000* program homepage, http://hypatia.ss.uci.edu/lps/psa2k/program.html.

Maudlin, T. (1990): "Substance and Space–Time: What Aristotle Would Have Said to Einstein," *Studies in History and Philosophy of Science*, 21, 531–61.

Nerlich, G. (1994): *The Shape of Space*. Cambridge, New York: Cambridge University Press.

Penrose, R. (1989): "Difficulties with Inflationary Cosmology," *Annals of the New York Academy of Sciences*, 571, 249–64.

Penrose, R. (1999): "The Question of Cosmic Censorship," in *Journal of Astrophysics and Astronomy*, 20(3–4), 233–48.

Poincaré, H. (1952): *Science and Hypothesis.* New York: Dover Books.

Price, H. (1996): *Time's Arrow and Archimedes' Point: New Directions for the Physics of Time.* Oxford, Oxford University Press.

Redhead, M. (1993): "The Conventionality of Simultaneity," in J. Earman, A. Janis, G. Massey and N. Rescher (eds.), *Philosophical Problems of the Internal and External Worlds: Essays on the Philosophy of Adolf Grünbaum*, Pittsburgh: University of Pittsburgh Press, 103–28.

Rynasiewicz, R. (1996): "Absolute Versus Relation Space–Time: An Outmoded Debate?" *Journal of Philosophy*, 93, 279–306.

Rynasiewicz, R. (2000): "On the Distinction between Absolute and Relative Motion," *Philosophy of Science*, 67, 70–93.

Sarkar, S. and Stachel, J. (1999): "Did Malament Prove the Non-Conventionality of Simultaneity in the Special Theory of Relativity?" *Philosophy of Science*, 66, 208–20.

Savitt, S. (2000): "There is No Time Like the Present (in Minkowski Space–Time)," *Philosophy of Science*, 67, Proceedings, S563–S575.

Singh, T. P. (1998): "Gravitational Collapse, Black Holes and Naked Singularities," available at gr-qc/9805066.

Stachel, J. (1993): "The Meaning of General Covariance," in J. Earman, A. Janis, G. Massey and N. Rescher (eds.), *Philosophical Problems of the Internal and External Worlds: Essays on the Philosophy of Adolf Grünbaum*, Pittsburgh: University of Pittsburgh Press, 129–60.

Teller, P. (1991): "Substance, Relations, and Arguments About the Nature of Space–Time," *Philosophical Review*, 100, 363–97.
Weingard, R. (1979): "Some Philosophical Aspects of Black Holes," *Synthese*, 42, 191–219.
Zangari, M. (1994): "A New Twist in the Conventionality of Simultaneity Debate," *Philosophy of Science*, 61, 267–75.