THE CONSTRAINTS GENERAL RELATIVITY PLACES ON PHYSICALIST ACCOUNTS OF CAUSALITY

Erik CURIEL*

* Philosophy Department, Stanford University, Stanford, CA 94305-2155, USA.
E-mail: encuriel@csli.stanford.edu

BIBLID [0495-4548 (2000) 15: 37; p. 33-58]

ABSTRACT: All accounts of causality that presuppose the propagation or transfer or some physical stuff to be an essential part of the causal relation rely for the force of their causal claims on a principle of conservation for that stuff. General Relativity does not permit the rigorous formulation of appropriate conservation principles. Consequently, in so far as General Relativity is considered and fundamental physical theory, such accounts of causality cannot be considered fundamental. The continued use of such accounts of causality ought not be proscribed, but justification is due from those who would use them.

Keywords: causality, general relativity, conservation laws, energy.

CONTENTS

> Those who make causality one of the original *uralt* elements in the universe or one of the fundamental categories of thought, -of whom you will find that I am not one,- have one very awkward fact to explain away. It is that men's conceptions of a Cause are in different stages of scientific culture entirely different and inconsistent. The great principle of causation which we are told, it is absolutely impossible not to believe, has been one proposition at one period of history and an entirely disparate one another and is still a third one for the modern physicist. The only thing about it which has stood, to use my friend Carus' word, a κτημα εζ αει, *-semper eadem-* is the *name* of it.
>
> Charles Sanders Peirce, *Reasoning and the Logic of Things*

1. Introduction

In recent work,[1] I proved a technical result that implies a precise statement and proof of the folkloric claim that there is no true analogue in general relativity of the principle of the conservation of energy as it appears in either classical physics[2] or in special relativity; I concluded from this that energy is not a fundamental quantity according to that theory. By the claim that energy is not a fundamental quantity in general relativity, I mean that within the mathematical structure of the theory one cannot rigorously define a quantity that has any of the features one might take to be definitive of energy. I should emphasize that my result does *not* imply that one cannot talk about energy at all within the theory. The startling fact, however, is that only in certain special situations can one represent within the theory a quantity that is structurally similar to energy as it is manifested in classical physics and special relativity, and even then only by employing explicitly approximative and idealizing techniques that are not part of the theory *per se.* Consequently talk about energy reflects nothing fundamental about the theory.

In this paper, I articulate one of the most important philosophical consequences of this result: restrictions on the sorts of accounts of causality that can be considered fundamental *vis-à-vis* general relativity. I emphasize that none of my technical results or arguments pretend to bear on accounts of causality that are not beholden in some way to fundamental physical theory. Some accounts of causality purport to treat only relations among middle-sized dry goods in everyday practical affairs; others take causality to be something akin to a logical category of thought that structures our knowledge of various matters; yet others take 'causality' merely to indicate that a special type of explanation is required or is in the offing; and yet others take it as a merely subjective, psychological phenomenon, the manifestation of a brute fact about the way we are constructed to view the world. The arguments in this paper do not pretend to bear on any such accounts. Certain sorts of accounts of causality, perhaps best exemplified by Russell (1927[1954]), rest on the idea that causality is a *physical* relation holding among *physical* entities, and as such must accord with best going physical theory. It is only such accounts that concern me here. Many such accounts rely on the intuitively appealing idea that energy and other classically conserved physical quantities such as momentum are intimately connected with causality. Especially popular is the idea that the propagation and transfer of energy, and all energetic processes in general, *embody* certain sorts of causal relations. For the sake of brevity I shall refer to all such accounts of causality as *transfer accounts*. In so far as such accounts hold this and similar tenets, I argue, and in so far as one considers general relativity, or at least certain aspects of it, to constitute fundamental physical theory, transfer accounts cannot be correct.

The argument of this paper does not pretend to be a mathematical demonstration of a certain philosophical point. Whereas I take the technical results of Curiel (1999) to show indubitably that, in so far as energy can be characterized in certain precise ways, it is not possible to ascribe a localized energy density to the 'gravitational field' in general relativity, nothing I say here has such mathematical precision and certainty accruing to it. Indeed, given the nature of the point I am arguing for, I do not think such a demonstration could be had either for its truth or its falsity. For the point is about the way that a certain picture of classical physics and of special relativity naturally *suggests* that causality has a certain character -and general relativity does not share the features of those theories that made the suggestion so plausible.[3]

The aim of this paper is twofold. The first goal is to try to determine what sorts of accounts of causality may be consonant with our best going science. General relativity by itself cannot be considered a completely fundamental physical theory; it is hoped and expected that it will in due course be replaced by a theory that unifies gravity with the other three fundamental forces now so successfully treated by the standard model of quantum field theory. Although we have little if any hard evidence indicating specific attributes such a grander theory ought to manifest -the triumphal claims of superstring theorists notwithstanding- there is good reason to think it will mirror certain fundamental features of both general relativity and quantum field theory that are too experimentally entrenched to envision discarding, just as general relativity itself mirrors certain fundamental aspects of Newtonian gravitational theory, *e.g.* the roughly $1/r^2$ dependence of the acceleration of mutually gravitating bodies in the weak gravitational field regime. I cannot say with certainty that the absence of conservation principles of a certain sort is that kind of fundamental feature of general relativity that must find analogous expression in whatever better theory comes along to subsume it. Given that this feature of general relativity follows directly upon the presence of curvature (in a certain technical sense) in the metrical structure of spacetime, in conjunction with the fact that the primary energetic quantity of the theory is a two-index tensor and not a scalar -as fundamental a pair of features of the theory as one can well imagine- it does not seem to me a foolhardy bet. Even if in the event it turns out not to be so and principles appear in that more fundamental theory capable of supporting transfer accounts of causality, I do not think my work here will have been a waste.

This brings me to the second goal of the paper: to illuminate the need for a better understanding of general relativity itself as a physical theory. It is perhaps not widely appreciated how poorly understood general relativity is in many respects *vis-à-vis* quantum mechanics. Much the greatest obstacle is the more difficult mathematics involved in mastering the theory and employing it to make physical predictions and retrodictions we have strong reason to believe are accurate. That, in conjunction with the poor experimental access we have to phenomena in strong gravitational fields, means that we have, for example, only a superficial understanding of methods for cataloguing solutions to the Einstein field equation according to physically relevant, generic features, both because individual solutions, not to mention entire classes of them, are difficult to come by, and because the detailed interaction between global spacetime structures and observable dynamic phenomena is poorly understood. Whatever little bit we can learn about the theory itself, even if it concerns an aspect of it that will not survive the transition to a future, better theory, should still be welcomed today for its own sake, if only to help us along the path to that future theory.

These difficulties by themselves constitute no reason to ignore quantum mechanics in so far as it bears on accounts of causality, as I shall steadfastly do after these few words. I shall ignore quantum mechanics because of the character of what we do know about it: in a direction orthogonal to the considerations of the previous paragraph, quantum mechanics is far less well understood than general relativity. The so-called measurement problem appears to demand an 'interpretation' of quantum mechanics in a way not required for the modeling of phenomena in general relativity. Foundational issues in quantum mechanics at the moment are so turbid and disputed as to allow a defense of almost no clear propositions about causality, except for the obvious: proceed with extreme caution. I hope for this paper to make the case by example that matters need not stand so in general relativity.

Even though general relativity is not obscure in this sense, it surely has much yet to teach us in our attempts to investigate the physical world and to incorporate the fruits of that labor into other areas of intellectual endeavor. It would be rash to assume that we have isolated and digested all the ways achieving a proper understanding of general relativity will require us to modify in fundamental ways root concepts and the relations among them, as special relativity required us, for example, to modify our concepts of 'simultaneity' and 'mass' to attain its proper understanding. Making a start at discovering such necessary modifications is the ultimate goal of this paper.

2. Causality and Energy

The 19th century discovery of the principle of the conservation of energy not only had immediate and wide-ranging ramifications into most areas of physics as known at the time, as Helmholtz himself, one of its discoverers, was at pains to stress from the beginning;[4] it was also generally recognized to have a bearing on a proper philosophical analysis of causality. Mill, for instance, in the eighth edition of his *Logic*, took the opportunity to add a section to discuss the bearing of the principle on his own account of causality.[5] The considered opinion of several respected philosophers of this century has also been that there exists a close connection between causality and energy, even when they have not promulgated accounts of the simple form "*A* is the cause of *B* if and only if *A* and *B* stand in such-and-such a relation as regards energy." They have tended to share the view that most, if not all, causal relations involve in some essential fashion a transfer of energy between cause and effect, and that this physical fact underlies much of the explanatory force of causal laws. Quine (1973, esp. p. 5), for instance, believes that the imparting of energy is the central idea in our common "causal idiom," and that the flow of energy provides a "root notion" of causality itself. Salmon (1984, esp. p. 146), inspired by Reichenbach (1956[1991]), comes close to identifying causal processes as exactly those that transmit energy. Russell (1927[1954] and 1948[1956]) and Reichenbach (1956[1991]) are more interesting examples, for they predicate their conclusion that transfer of energy is intimately connected to causal relations, as most other such philosophers do not, upon a thorough analysis of scientific theory, including the working through in detail of examples of energy propagation and transfer in various types of physical situations.

The idea that energy and causality are somehow intimately related to each other has much to commend itself. From a naive viewpoint, the natural measure of the quantity of available energy 'stored' in a dynamical system -how much *work* it can do on other dynamical systems- by itself suggests such a relation: a cause is, roughly speaking, something that produces a change in something else, and changes often require work. Of more philosophical interest, the fact that energy appears to propagate and be exchanged continuously holds out the hope of offering an explanation of one of the most deeply rooted and widely held beliefs about the character of causality, that it itself is manifested continuously along well-defined, spatiotemporally localized paths or chains: that between any localized entity *A* and any other localized entity *B* such that *A* causes *B* there is always a third localized entity *Q* 'spatiotemporally between' *A* and *B* such that *A* causes *Q* and *Q* causes *B*,[6] and such that the causal efficacy 'travelled' from *A* to *B* via *Q* and *not* via any other localized entity *Z*. All propositions of this sort I shall group together under the rubric 'the principle of causal continuity'. I am firmly convinced that most of the attraction of transfer accounts of causality can be reduced down to this notion of causal continuity and its 'explanation' by reference to the continuity of energetic processes.[7]

I think Russell and Reichenbach were right to try to ground such views of causality upon a thorough analysis of fundamental physical theory, for such views presuppose several substantive theses about the structure of the physical world: that, for example, a dynamical quantity called 'energy' exists and has some of the properties we naively associate with it; in particular, that packets of it can be transferred between dynamical systems in such a way that one can not only keep track of the identity of certain packets of it for brief periods of time, but also so that one can determine what was the source and what the sink of a given bit; that particular segments of the time-evolution of dynamical systems can be distinguished as those during which the system is 'isolated from its environment'; and so on. None of these seem on their face particularly contentious. A conception of causality, however, that accepted such theses and that had the stated goal of being fundamental -in the sense that its motivation or justification was ultimately to be found in fundamental physical theory- *just because* it presupposed such substantive theses about the physical character of the world, would require for its justification a thorough investigation of the best going current physical theory or theories, at the least in order to demonstrate that such assumptions as these do not conflict with any of the precepts of our best current science. One would have to ascertain for instance not only that the best physical theories did not preclude this account of causality from the start, as they would, say, were energy not a well-defined quantity in them, but to ascertain also that the particular types of interactions demanded by such an account were representable in the theories. Best of all would be an argument that showed how such an account of causality could be read off directly from the mathematical or conceptual structure of the best going physical theories, in the way meant when classical particle mechanics is said to require no interpretation.

3. Causal Relations and Energetic Processes

To make a start on sorting out the different ways energy can plausibly be thought to bear on causal relations, consider that exemplar of causation, the naive picture in classical physics of the motion of impenetrable, perfectly elastic bodies and their impinging on each other. The usual story says that when such a body in motion, *A*, strikes such a body at rest, *B*, under 'normal conditions', the first body imparts motion to the second by transferring to it some of its kinetic energy: *A*'s striking *B* causes *B* to start into motion. The first step in the analysis of the possible relations of energy to causality will be to tease out of this brief telling of the old chestnut the propositions that warrant the causal claim.

Assume that *A* and *B* are completely isolated from other physical systems -perhaps they inhabit the near-perfect void of intergalactic space- which is to say, they are not interacting with any other dynamical systems in such a way as sensibly to affect their dynamical evolution. Fix an inertial coordinate system in the representation of which *B* is at rest. By hypothesis, before the collision *A* propagates so that its center of mass traces out a straight line that, continued indefinitely, would pass through *B*'s center of mass. As *A*'s center of mass successively occupies each point of this line before the collision, *A*'s kinetic energy approximately equals its total energy, which remains unchanged -it is *conserved*. This is one of the indications that *A* indeed is not interacting with any other dynamical system, though it is only a necessary and not a sufficient condition: if one swings a ball in a circle at a constant rate on the end of a string, the ball interacts with the string, but the ball exchanges no energy with it, only momentum.

The fine details of a collision between two elastic bodies are complicated beyond human comprehension; I am happy to report that for our purposes we do not need to comprehend the collision in fine detail. The net result of the impact is that *A*'s direction of motion alters and its speed of propagation diminishes, and that *B* starts into motion with precisely the direction and speed to compensate for *A*'s lost energy and momentum: the sum of *A*'s and *B*'s kinetic energies after impact equals *A*'s kinetic energy before impact, and the vectorial sum of *A*'s and *B*'s linear momenta after impact equals *A*'s linear momentum before impact, and similarly for their angular momenta defined relative to any fixed point.

Were one inclined, say, to a Humean point of view, one could not draw any *causal* conclusion from this description of the process that did not ultimately depend on the way past experience had habituated one's reasoning faculties. For one who champions the idea that causality is a feature of the physical world, however, the story offers an obvious option: *A*'s striking *B* caused *B* to start into motion in virtue of the transfer of some of *A*'s energy to *B* via the work performed by mutual pressure during the time they were in contact. The crucial supposition is that the energy manifested by *B* as it starts into motion 'came from' *A*: *A* imparts some of *its own* energy to *B* during the process. Were *A* to strike *B* and *B* subsequently to start into motion, though one somehow determined that the kinetic energy *B* had acquired came not from *A* but had been transmitted in occult fashion from the distant body *Q* (and perhaps that the kinetic energy that *A* had lost as it slowed or stopped upon contact with *B* had flitted instantaneously off to the distant body *Z*), or one determined that *B*'s newfound kinetic energy had sprung spontaneously into being at that moment *ex nihilo*, and likewise that *A*'s lost energy simply vanished *ad nihilum*, one might well question whether *A*'s coming into contact with *B* had been the cause of *B*'s starting into motion; perhaps the transfer of energy from *Q* or the spontaneous creation of energy would be better described so. Without the idea that *B*'s newfound energy once had been *A*'s, there seems no way to bind the two into a relationship intimate enough to move beyond the simple conditional that mathematical physics warrants -'If *A* strikes *B* thus, *A* and *B* will move so'- and into the *recherché* realm of the causal.

What, then, schematically stated, are the roles energy must play in order to support causal judgements? The primary task for energy will be to provide warrant for asserting judgments of the form "*C* causes *E*" when physical theory has already affirmed the truth of the proposition "If *C* then *E*." The *C*'s and *E*'s, as shown by the example of this section, will in some cases be different states of the same system, as when the state of *A* above at the time it contacted *B* is said to have been mediately caused by its state at earlier times. In other cases, they will be states of different systems at the same time, as when the state of *B* at the time it started into motion is said to have been caused by the state of *A* at that time. Finally, in still other cases they will be states of different systems at different times, as when the state of *A* at some early time is said to have been the mediate cause of the state of *B* at some later time. In the first type of case, it is the continuous propagation of energy, here in the form of the continuous motion of a ponderable body, that is to warrant the causal claim; in the second, it is the proximate, continuous transfer of energy among dynamical systems, here in the form of the work performed by contact pressure between two ponderable bodies, that is to do so; and in the third, it is a combination of the first two. I therefore turn now to analyze these two types of energetic processes, assuming that the third can be understood as a simple combination of the first two, to determine the properties energy must have to realize these processes in such a way as to support causal claims.

4. The Propagation and Transfer of Energy

Consider first propagation. As best we know, there is no such phenomenon in this world, nor in the realms of the theories of classical physics and relativity, as pure energy propagating, of and by itself, as its own entity -as its own dynamical system. To use a Scholastic idiom for a moment, energy is not a substance to support attributes. Even photons, which may appear so, are not: they have also momentum, angular momentum and spin. One could with the same justice say that the photon was pure momentum, with the attributes of energy, angular momentum and spin, as say that it was pure energy with the attributes of momentum, *etc.* The *photon itself* is the system, and the rest are its attributes. The famous relativistic equation of mass and energy might lead one to think otherwise, but it is beside the point here, for mass itself is only one more possible attribute of a system and does not by itself constitute a substance capable of supporting attributes. Talk about the propagation of energy must always be understood to be shorthand for talk about the propagation of a particular dynamical system, to which the energy is attributed.

In order to conceive of energy as propagating, one must first be able to reidentify it as it occupies different spatiotemporal loci. In classical physics, this means being able to identify a certain hunk of it as belonging to the same system at different locations in space, at different moments of time; in relativistic physics, to the same system at different points of spacetime.[8] The question of reidentifying energy over spatiotemporal intervals, therefore, reduces in part to the question of reidentifying dynamical systems over spatiotemporal intervals. In addition, there must be criteria for deciding when the system 'has the same energy as it had before', in the sense relevant to supporting causal claims.

In both classical and relativistic physics, the identity over a spatiotemporal interval of a dynamical system is constituted by (at least) the continuous occupation of the points of the interval by an entity the same in all (or enough) relevant respects. This is not the vacuous truism it may seem. With very few exceptions, all known fundamental equations of classical and relativistic physics have the character that, given an entity occupying a particular bounded spacetime region, obeying a particular set of equations of motion, with a particular specification of its physical properties, there will be in the absence of interference a continuously varying family of continuous timelike paths (a four-dimensional spatiotemporal 'tube'), each path of the family having one of the points of the original spacetime region as an end-point, such that each point of the tube is occupied by a point of an entity the same, in just those respects normally germane to identification, as the entity in the original spacetime region. This is only to say that, in their guise as first-order ordinary differential equations, classical and relativistic equations of motion all possess well-posed initial-value formulations.[9] It is this aspect of dynamical systems -that they propagate continuously, a guaranteed consequence of the *form* of their equations of motion- that ultimately underwrites our identification of them over spatiotemporal intervals, and so this must be a criterion by which a hunk of energy propagating along with the system, so to speak, is to be identified over time in the respect relevant to the support of causal relations.

What counts as "those respects of the system germane to identification" will vary from type of system to type of system, and even perhaps for the 'same system' (in a naïve sense) as it appears in different sorts of investigations. Were one, for example, able to track a photon traversing cosmological distances through an expanding cosmos such as ours, the fact that its energy continually decreases (the "red-shift effect") would in many circumstances not stop one from asserting that it was in an important sense the 'same' photon that got emitted from a certain star. Otherwise one would appear to rule out the possibility of investigating dynamical systems at cosmological distances from us: if one cannot assert that this photon is in some important respect the 'same' as the one emitted from that star, one will normally have no ground for using any information gleaned from the photon to infer any information about the star. In other cases, changes in a system concomitant with changes in its energy may push one to conclude that the resulting system is not the same as the original system. If one pumps enough energy of the right sort into a hydrogen atom, the electron will escape the central proton and fly off freely. To an organic chemist, the widely separated, relatively independent proton and electron may no longer constitute the same system as the original hydrogen atom, whereas a high-energy particle physicist may consider them to be precisely the same system, only in a very different state than before. Finally, in some instances there may be no way to reidentify a system indefinitely over spatiotemporal intervals: I know of no way to pick out 'part' of an excited atom and rightfully assert that it is identically continuous with the photon it absorbed a moment earlier.

I take it as a fundamental assumption of physical science that "those respects of the system germane to identification" can be specified in any given case, even though there is no universal formula specifying a procedure for doing so in all cases. There is, for instance, no variable in its equations of motion the value of which represents the fact that the system is, say, a hydrogen atom -only variables for position, momentum, *etc.* That what is being modeled is a hydrogen atom is encoded in the formal relations among the variables representing its dynamical quantities and in the values of the intrinsic, kinematic parameters one must fix (mass, spin, *etc.*) to represent it, *i.e.* precisely in the form of its equations of motion and in the canonical geometry of its space of states and the set of vector fields on its space of states representing the system's kinematically allowed dynamical evolutions.[10] Let us call the set consisting of the system's equations of motion, its space of states, and the set of vector fields on its space of states representing its kinematically allowed dynamical evolutions the *dynamical representation* of the system. Then the similarity in form of the dynamical representations of entities at neighboring spacetime regions is in almost all cases a necessary condition for identifying the entities as being spatiotemporally proximate manifestations of one and the same dynamical system. Let this suffice for a discussion of the first condition required for reidentifying a bit of energy over spacetime intervals, that of the reidentifiability of dynamical systems over spacetime intervals.[11]

The last of the three examples cited above, that of the atom that had absorbed a photon a moment earlier, suggests what is required for the second condition, that or those guaranteeing that the energy of the evolving, continually identical system remains the 'same' in the sense required to support causal claims. It is not the case that a hunk of energy can be identified once and for all, no matter what happens to a system that happens to 'contain' it at any given moment, so to speak. A hunk of energy can be identified over time as being the same in the relevant sense *only* so long as its associated system evolves in isolation, with no external interactions;

moreover, one cannot naturally divide up the entire energy content of an isolated system into separate parts in order to keep track of each part over time. One can keep track of and reidentify only the entire quantity of energy associated with an isolated system over time. What is needed then is a criterion for determining when a system is isolated, which is to say, not interacting with its environment.

The analysis of the identity of dynamical systems just offered, as depending on the form of a system's dynamical representation, suggests a definition of 'in isolation'. A system will be said to be *isolated* during a stretch of its dynamical evolution if its actual equations of motion imply conservation of all classically conserved quantities during that stretch. Stipulating that the system be isolated, though, does not by itself suffice for concluding that the energy associated with the system at each instant it is isolated is in some significant sense the 'same'. The forms of the equations of motion of isolated systems both in classical physics and in special relativity certainly imply that isolated systems will have a definite quantity of energy at every instant, as represented in a fixed global coordinate system, but they do not state that the energy of a given system is the 'same' in any way at any two instants other than perhaps being quantitatively the same. Identifying the energy of an isolated system in both classical physics and special relativity as the 'same' in more than the mere quantitative sense during the period it is isolated can itself be justified only by the fact that the quantity of the system's energy at each instant is in fact the same: the energy of an isolated system is conserved. This brute fact is supposed to justify the further thought that energy can be neither created nor destroyed, and so, *a fortiori*, is the 'same' at each instant in some deeper way.

With so much behind us, we need not take long discussing the second of the processes required for supporting causal claims, energy transfer from one system to another -that, when two (or more) systems interact, the gains and losses of energy of each system during the interaction can in a significant way be matched up with each other, *e.g.* the energy gained by *A* was *transferred* from *B*, or the energy lost by *A* was *transferred* in part to *B*, in part to *C* and in part to *D*, *etc.* The same sort of analysis as worked for propagation will apply here as well, so I shall only adumbrate it. In interactions represented in classical physics and in special relativity, it is a consequence of the form of the fundamental equations of motion that energy is conserved. This brute fact provides the warrant for asserting that, in such interactions, the energy that one system loses is the 'same' as the energy the other system gains, at least in the way required to support the desired sorts of accounts of causality. The fact that it is always the same *amount* of energy that is gained and lost by interacting systems is supposed to preclude the idea that energy can be created or destroyed, and so warrant the inference that energy is actually *transferred* between interacting systems, as required.

In sum, a necessary condition for characterizing both the propagation and the transfer of energy so as to be of use to transfer accounts of causality is that dynamical systems satisfy the principle of the conservation of energy. It is, moreover, not just any representation of the conservation principles that is required by these accounts, but primarily the integral form. The idea that some stuff is being transferred or is propagating refers of necessity to more than one spacetime point, that or those from which and that or those to which, and whereas the differential form of the conservation principle refers only to one spacetime point, it is the integral form that represents the fact that the stuff is conserved over finite spatiotemporal volumes -the form Bob Geroch once in conversation felicitously referred to as "the mathematical representation of thinghood," in so far as it suggests traditionally essential features of substance, in particular its identifiability over time and its permanence, which is to say the impossibility of its creation *ex nihilo* and of its destruction *ad nihilum.* Not only do both the differential and the integral forms of the conservation principle hold both in classical physics and in special relativity, but they are in fact logically equivalent to each other, so all seems in place for the possibility that energetic processes as represented by these theories can be used to ground the desired causal claims.

5. Energy in General Relativity

General relativity does not support any mathematical structure with which to construct relations similar enough to classical conservation principles to wear the name gracefully, at least so far as the sorts of accounts of causality I am considering are concerned. The precise statement is that the only 'conservation law' one can formulate in a generic general relativistic spacetime is a differential *covariant* conservation law. No two physicists, not to mention philosophers, seem to agree on what exactly the import of being covariant is for an equation in general relativity.[12] For my purposes, it suffices to say (which I think is not contentious) that a necessary part of what makes the differential covariant conservation law in general relativity *covariant* is the fact that the 'differential' in this law comes from the covariant derivative operator naturally associated with the ambient spacetime metric. Consequently, in a generic general relativistic spacetime there is no privileged, physically significant way to cast the differential covariant conservation law into the form of an ordinary partial differential equation or set of such equations, and so such a law cannot in general be transformed into an *integral* conservation law. There simply are no integral conservation laws in the generic general relativistic spacetime.

That the covariant differential conservation law does not imply an integral conservation law follows from these two facts: first, that the fundamental energetic quantity in general relativity (as in special relativity) is not a scalar function on spacetime but is rather a two-index covariant tensor field, the stress-energy tensor T_{ab}; second, that generic spacetimes in general relativity have no preferred class of 'time-frames', as Minkowski space in special relativity has, *viz.* those defined by classes of worldlines of inertial observers all at rest with respect to each other.

Were the differential covariant conservation law to yield an integral conservation law in general relativity, this integral law would have to pertain to a scalar quantity, such as energy or a particular component of linear momentum. Integral conservation laws keep track of gross quantities of a stuff over finite regions, and in general relativity it does not make sense to ask of a vectorial or tensorial quantity "How much total is in this finite region?"; more precisely, the 'integral of a vectorial or tensorial quantity' is not a well-defined mathematical entity in general relativity. Over regions of spacetime one can integrate only scalar quantities.[13] This is a consequence of the fact that the generic general relativistic spacetime has intrinsic curvature.[14] Consequently, that the fundamental energetic quantity in general relativity is a tensor -the stress-energy tensor- and not a scalar already bodes ill for the possibility of formulating integral conservation laws.

That it is a tensor, however, cannot by itself preclude the formulation of an integral conservation law, for the fundamental energetic quantity in special relativity is also a tensor, not a scalar, and yet there are integral conservation laws therein. These integral conservation laws in special relativity are founded on the circumstance that in Minkowski spacetime, the spacetime of special relativity, there are families of timelike curves representing inertial observers all at rest with respect to each other, which have the extraordinarily nice technical property of being simultaneously geodesics and the integral curves of a Killing field. Minkowski spacetime is essentially the only one possessing this feature; this fact is related to the idea that Minkowski spacetime is essentially the unique spacetime with no 'gravitational field'.

In the generic general relativistic spacetime, not only will there not be a family of timelike curves that are all simultaneously geodesics and the integral curves of a Killing field, but there will not even be a family of timelike curves that are simply the integral curves of a Killing field. This last structure by itself in fact suffices for formulating an integral conservation law. Now, the presence of intrinsic curvature in the spacetime manifold does not by itself imply that there cannot be timelike Killing fields: there are solutions to the Einstein field equation that represent curved spacetimes with timelike Killing fields, and in these spacetimes integral conservation laws can be defined. That the generic general relativistic spacetime possesses intrinsic curvature, though, does make it extremely difficult for it to have Killing fields, timelike or not. The reason behind this, intuitively speaking, is as follows. If an observer were to embody an integral curve of a timelike Killing field, she would record an extraordinary fact: the metrical structure of spacetime, in a sense that can be made precise, would appear to her not to change in the slightest as her proper time elapsed. At every moment of her proper time, spacetime would appear essentially the same as at every other moment. For this reason, timelike Killing fields are said to represent 'time-translation symmetries'. This property undergirds the Killing field's capacity to facilitate the formulation of integral conservation laws -they provide a physically significant, preferred temporal background, so to speak, against which one can track the gross quantity of energy in a given spatial volume as it 'evolves' with respect to the metrical structure of the spacetime, which thanks to the symmetry implied by the presence of the Killing field can be taken as 'constant over time' in a certain sense.

The spacetimes in which Killing fields occur, however, are highly special and unphysical. Special, because a generically curved spacetime will not manifest such extraordinary symmetry, as one ought to expect: think of a 'generically curved' surface -perhaps a sheet of rubber that is distended and stretched at random- and it will manifest any sort of symmetry whatsoever only under rare circumstances, not to mention manifesting a perfect, global symmetry such as is embodied in a Killing field. Unphysical, because such spacetimes are unstable against arbitrarily small inhomogeneities: the smallest speck of dust the tiniest bit out of place in only one spot in the entire spacetime precludes the existence of a Killing field. It is only in such unphysically dainty spacetimes, by dint of the daintiness itself, that one can define a quantity that behaves enough like energy even to be tempted to call it that -it goes without saying that the actual spacetime we inhabit possesses no Killing fields. Before one gives in to the temptation to think of the scalar field used in such an integral law, which the timelike Killing field allows one to construct, as the total energy density at a point, two important caveats must be registered. First, this 'energy density', though constructed relative to a particular family of timelike geodesics ('inertial observers'), will not be the energy density that any actual observer instantiating one of the geodesics would measure using any standard experiment for measuring energy density; this would occur only if the integral curves of the timelike Killing field were also geodesics, which extraordinary confluence occurs only in Minkowski spacetime. Second, it does not include any contribution due to 'gravitational energy', since this is not localizable in general relativity.

The second caveat brings us to the heart of the matter: the impossibility of defining in general relativity a mathematical object that represents a local energetic quantity specifically associated with the 'gravitational field' -one cannot ascribe a local energy density to gravity, or really any localized energetic quantity to it at all with any degree of rigor.[15] It is no accident that this occurs precisely in spacetimes with nontrivial curvature, for it is precisely spacetime curvature itself that encodes all the dynamical information traditionally ascribed to the gravitational field. The upshot is that any other localized energetic quantity will always suffer an irremediable defect: in so far as it will not represent *all* the available 'energy' in a given region, one will never be able to formulate an integral conservation law for it.

That one cannot attribute a definite localized energy to the gravitational field in general relativity is on its face a puzzling result, for it is not difficult to convince oneself that one can extract energy from the gravitational field -after all, energy is continually transferred from the moon's orbit to the oceans through the work done in the rising and ebbing of the tides.[16] The principle of energy conservation, moreover, seems one of physics' most dearly held principles. Its consequences produce the predictions that have confirmed our most fundamental quantum theories to mind-boggling degrees of accuracy. Engineers employ it constantly in designing the contraptions that, by and large successfully, house, feed, transport and entertain us. So what gives? The answer is that general relativity tells us that, rigorously speaking, there is no such quantity, but that in certain sorts of approximations and idealizations one can recover a quantity that is naturally identified as energy. When the background curvature of a spacetime region is 'small', one may treat the region as being for all practical purposes flat, with the consequence that there will be an 'approximate timelike Killing field', and one may proceed to define energy as one does in the presence of a true timelike Killing field. That this approximation holds good in the region of the solar system explains how the idea of energy can be so useful to us, and appear so fundamental, when in fact one of our two fundamental theories says it is not. This procedure is in fact doubly approximative, in that there is no precise definition of an 'approximate Killing field' -in practice, physicists wing it on a case by case basis, and this is appropriate for their tasks. For we, though, who investigate the would-be *fundamental* features of the world as represented by general relativity, approximations and idealizations, no matter how well justified in certain experimental calculations and practical endeavors, have no relevance.

Even though the idea of 'energy', classically so dependent on the conservation principle for its definition, in one sense disappears in general relativity, it does not do so completely. It would be more accurate to say that the idea of 'energy' *alters* in the transition from classical physics to special relativity, and again in that from special to general relativity. In the first place, although this is not often explicitly recognized, in classical physics there are actually two separate conceptions of energy, each with its own distinct proper mathematical representation: energy as the capacity to do work (closely related to the idea of potential energy), properly represented by a 1-form on the space of states of a classical dynamical system, the 'work 1-form', *i.e.* a linear mapping from vectors tangent to the space of states ('rate of change of the state of the system') to real numbers; and energy as the generator of the dynamical evolution of a system (closely related to the idea of kinetic energy), properly represented as a scalar field on the space of states in conjunction with a linear mapping from scalar fields to a certain set of vector fields on the space of states, those representing the kinematically possible dynamical evolutions of the system. When these objects satisfy certain conditions, then one can formulate the usual conservation laws, which quantitatively relate the two conceptions of energy by equating the total amount of energy (in the sense of the generator of evolution) gained or lost by a system to the amount of work performed on or by it during an interaction.

In special relativity, there is fundamentally only one energetic quantity, the stress-energy tensor. The relativistic equation of mass and energy requires a mathematical structure that will keep track of the fact that energy flux has momentum ("$E = mc^2$") and that momenta contribute to energy flux -which is all neatly encoded in the person of a two-index symmetrical tensor, *viz.* the stress-energy tensor. One can derive from it analogues to the objects representing the two conceptions of energy in classical physics by fixing an inertial coordinate system and decomposing the stress-energy tensor into its energetic, linear momental and angular momental components. Whereas in classical physics there were two fundamentally distinct conceptions of energy, united only by the conservation laws and this only contingently, special relativity teaches us that there is only one underlying quantity, stress-energy, with some, but not all, of the characteristics of energetic quantities in classical physics. Notably: energy density is still `quadratic in velocities' -it requires two tangent vectors to extract a scalar from the stress-energy tensor; and integral conservation laws of a certain sort can still be formulated in special relativity, so the gross energetic quantities displayed in such laws can be related in physically significant ways to localized scalar energetic quantities. Finally, in the shift to general relativity one retains much of the structure of stress-energy in special relativity, except this key aspect only: there are in general no integral conservation laws, and correspondingly there are in general no well-defined scalar energetic quantities of physical significance. In many situations of practical and theoretical interest, however, one can formulate approximate integral conservation laws and the correlative scalar quantities with many of the properties such structures have in special relativity.

Although the idea of energy and the sorts of fundamental energetic quantities extant do shift dramatically as one progresses up the ladder of theory, they do not alter beyond recognition, and in fact there are fundamental continuities in the relevant mathematical structures, as I have tried to emphasize, continuities which reflect the fact that the newer, better theories must represent those aspects of physical phenomena represented with great exactness by the older theories they supercede. I think this is absolutely important to realize points similar to it often get overlooked in contemporary philosophical discussions of 'paradigm shifts', and the like. General relativity demands revision of the classical conceptions of 'energy' and 'conservation' but not a complete jettisoning of them. I yet believe that we have not come fully to grips with the revisions it urges on us.

6. Causality after General Relativity

So where does all this leave us? It seems clear to me that it leaves us with no way to represent transfer accounts of causality within the fundamental structure of general relativity. Almost every aspect of general relativity, in fact, militates against this conception of causality, and indeed against any conception of causality reliant on a strong principle of causal continuity. One can predict with great certainty the regularity of certain relations among energetic sorts of quantities in general relativity, but this by itself will not suffice to support the types of causal claims advocates of transfer accounts of causality want to make. In the absence of integral conservation laws of the proper sort, there is no reason to take such regularity and its sure prediction as expressing anything more than the bare mathematical assertions that they are -"If energetic quantities of a certain sort, in a certain amount, are here, in this situation, then energetic quantities of a certain sort will be there, in a certain amount." This sort of proposition is exactly as causal -or not- as one thinks unadorned mathematical physics is. If one *does* think mathematical physics by itself is causal enough, then there is no need to muck about with anthropocentric and turbid notions such as 'propagation' and 'transfer', which are no part of mathematical physics *per se*.

For the determined advocate of transfer accounts, a reasonable first reaction to my arguments might be to give up the old characterization of propagation and look for a new one better suited to general relativity's *mise en scène*. Since general relativity does not allow the rigorous definition of a localizable physical quantity that has the essential features of energy as it appears in classical mechanics, however, nor of any other classically conserved quantities, it is not clear what one would have propagate even could one devise a new definition of it. As a last ditch attempt to salvage the notion of propagation, one might be tempted simply to take particles themselves as what, by propagating, support a transfer account of causality -one cannot rigorously ascribe energy, mass or momentum to a particle, but its mere continuous, self-identical existence along its path through spacetime ought to count as a perfectly good case of propagation, and surely such propagation can underwrite the sorts of causal claims people want to make. This looks to be perhaps a promising avenue until one realizes that it will never work. Strictly speaking, one cannot formulate the Einstein field equation in a mathematically sensible way for point-particle sources, as one can, say, for the Maxwell equations in special relativity.[17] Even were this technical hurdle surmounted, a more serious problem confronts this proposal: point particles are only idealized entities, useful in certain sorts of approximations; nothing in nature answers exactly to the idea. To the best of our knowledge, there are only extended bodies and fields, perhaps only fields. Consequently point-particles cannot be utilized to ground an account of causality with pretensions of being fundamental.

Finally, neither extended bodies nor fields will yield by themselves any way to define propagation in general relativity: there is no way to single out any particular curve in a spacetime region occupied by a spatially extended object or field in such a way as to give one reason to claim that anything of significance propagates along that curve. Quantities such as energy and momentum serve this function in classical mechanics, but one cannot call on them here. One might try to use the 'propagation' of an entire extended body to try to underwrite the desired causal claims, but my analysis in §4 above of the continuing self-identity of dynamical systems shows that, in doing this, one no longer is relying only on the fundamental structure of general relativity. The equations of motion of dynamical systems do not contain 'identity variables' -dividing the world up into discrete, extended bodies is *not* a part of fundamental physics as captured by general relativity, but is rather tied up with our preferred way of doing physics, what Bob Geroch evocatively calls 'psychology'. It should also be emphasized that such a causal conceit is artificial to the point of being ludicrous when one considers *fields* rather than extended bodies -it is difficult to know how to make sense of the idea of a discrete, bounded 'chunk' of field propagating *en bloc*. I actually should want to say that another lesson general relativity urges on us is that the distinction between fields and ponderable bodies is not a fundamental one -that gravitational energy cannot be localized, in conjunction with the relativistic equation of energy and mass, renders Quixotic any attempt to distinguish with any degree of precision spacetime regions void of ponderable matter from those occupied by it- but that is a sermon for another time.

General relativity does not by itself suggest entities or quantities that one will want to characterize as 'propagating', no matter how one defines it. The very different structure of spacetime in the theory from that of spacetime in classical physics and in special relativity does not naturally suggest any sort of transfer account of causality, nor does it easily admit one. The only reason I can imagine for trying to force one to fit into the framework of general relativity is because one approached the theory with a set of classical notions and questions already in hand to address, and did not rather begin by asking general relativity what the important notions and questions ought to be in its new framework -which is no better than a neo-Aristotelian's asking Newton how to define the 'φυσιζ' of a body in his gravitational theory.

If one renounces transfer accounts of causality, as I see it there remain only two general sorts of accounts of causality that can be made consonant with general relativity. One may postulate an account in which one or more discrete, localized entities that are spatiotemporally separated from each other 'cause' a distinct entity, the 'effect', itself spatiotemporally separated from all the 'causes'. Just as with any attempt to hold on to propagation, however, such an account would in no way arise from the structure of general relativity itself, but would rather have to be forcibly superimposed on its structure, under the guidance I suppose of purely metaphysical urges. Otherwise, there is the initial-value formulation of mathematical physics, my preference for the best one can do in representing causality in general relativity. Of course one also has the option of not giving an account of causality at all, and simply going about one's business with the physical theory -this may be my favorite of the options. Whatever sort of account one will give looks to come perilously close to saying that one thing *follows* upon another, perhaps dimming the explanatory glow that many philosophers have found in causality.

Philosophers involved in projects ranging from arguments for the physical basis of the direction of time (*e.g.* Reichenbach (1956[1991])), to the origin of linguistic reference (*e.g.* Putnam (1973)), to analysis of perception (*e.g.* Russell (1927[1954])), to accounts of physical measurement (*e.g.* Hacking (1983)) and defenses of realism (*e.g.* Boyd (1983)), have in the past freely availed themselves of causal notions relying on the ideas of causal continuity and the propagation and transfer of stuff. I do not think that they ought to have access to such notions for free any longer. That certain concepts do not accurately mirror the structure of the world at a fundamental level does not *ipso facto* preclude them from useful service in many areas of intellectual endeavor, but it does demand that such use be scrutinized and justified. It would, for instance, be a strange though possible theory of linguistic reference that broke down in the vicinity of black holes -surely something would have to be said about why this ought to be so. To repeat: I do not desire to extrude from the philosophical and scientific armory all notions of causality that depend on ideas of propagation, transfer and classically conserved quantities, much less to banish them from everyday discourse about everyday objects, but I think my argument does demand from any philosopher who wishes to invoke such a notion in his arguments an account of why he is justified in doing so, why his topic calls for that sort of notion, in light of the fact that there are strong grounds for believing such a notion cannot be, fundamentally speaking, true. In particular, any account of causality richer than the initial value formulation of mathematical physics that is supposed to arise naturally from an analysis of physical theory ought to be treated with suspicion.

My argument, though affronting to some dear contemporary intuitions, ought not be too surprising on reflection. The ideas of propagation and transfer have not been intimate with the notion of causality commonly held by the intelligentsia for long, certainly not for more than 400 years, since the time of Galileo. Before the scientific revolution, such ideas were no part of generally held conceptions of causality among those who contemplated such matters. Not even Aristotle's 'efficient cause' resembles in any way any of the ideas associated with transfer accounts; and one will search long and hard in Hellenistic, Roman and medieval philosophy and still turn up nothing that suggests such a picture. It was only in response to the development of classical mechanics that these ideas began to ingratiate themselves widely, and to acquire the honorific 'intuition', with its attendant privilege -future argumentation and theorizing had to conform now to these principles, but they themselves did not stoop to be questioned. The wide acceptance of action-at-a-distance pictures of gravitation and electromagnetism in the eighteenth and nineteenth centuries even suggests that these ideas were not regarded as intuitions until quite recently, perhaps beginning only with the general acceptance of Maxwell's electromagnetic theory, after Hertz's landmark experiments in the 1870's proved the existence of electromagnetic radiation.[18] Such intuitions would then have hit their stride with the acceptance of relativity theory in the early decades of this century, the more troublesome aspects of quantum theory being conveniently overlooked. I fault Hume's eloquent misunderstanding of Newtonian theory for blinding much of the philosophical world to the import of this history, even down to this day; but that again is a story for another time.

Acknowledgements. I thank David Malament and Howard Stein for impeccable advising on and penetrating criticism of my doctoral thesis, from the third section of which this paper was harvested, and for many stimulating conversations on these and related topics. I also thank Robert Geroch for many stimulating, edifying conversations. If it were not too cheeky, I would thank him for being an unerring oracle on all matters physical and mathematical. And if I were to do that, I could not in good conscience fail to thank Howard Stein for being an unerring oracle on all matters philosophical –but I cannot yet bring myself to do this.

Notes

1 Curiel (1999).

2 When I say that a physical theory or system is 'classical', I mean that it finds its proper representation in the spacetime of pre-relativity physics and that it does not fall under the purview of quantum mechanics.

3 Many among both philosophers and physicists have retrojected the idea of locality into all facets of classical physics, and have used the idea as a foundation in their analyses of causality even when they consider only classical phenomena. Since this intuition lies behind the sorts of accounts of causality I critique, but mostly because it makes classical theory more similar to relativistic theory, which we have good reason to think expresses a deeper understanding of the physical world, I shall treat classical systems as interacting locally, and classical physics in general as conforming to a principle of local action. Of course, this requires that I bracket Newtonian gravitational theory entirely, since it is a non-local theory down to the bone. See Curiel (1999) for a discussion of energy and causality in Newtonian gravitational theory.

4 *Cf.* Helmholtz (1847[1853], *passim*).

5 Though Mill refers to it as the 'Principle of the Conservation of Force'; *cf.* Mill (1874, preface, p. viii). He concludes in the added section (book III, chapter V, §10) that in fact nothing requires alteration in his original account of causality *per se*, though it does afford the possibility of an interesting elaboration of it, *viz.* by providing a criterion to winnow in certain situations true causal chains from spurious correlations, by tracking energy flow.

6 This statement says only that causal chains are *dense*, not necessarily continuous. A precise statement of the necessary and sufficient conditions for the continuity of causal chains would involve elaborate, unnecessary technicalities.

7 It is astonishing how widespread such an assumption of causal continuity is in the philosophical literature, especially when one considers how recent has been the adoption of such views in the physics community (wherein action-at-a-distance theories reigned until the late 19th century), and when one considers our woefully inadequate comprehension of quantum mechanics, which itself does nothing to encourage such views in any event. Philosophers as disparate in temperament and aim as Russell (1927[1954] and 1948[1956]), Reichenbach (1956[1991]), Quine (1973), Mackie (1974), Salmon (1984), Lewis (1986), Mellor (1995) and many others too numerous to mention have all invoked in discussions of causality a principle of causal continuity. Perhaps most striking of all is that, among this whole lot, only Russell, in both works, discusses the possible grounds for holding such a principle and the consequences of its falsity, if it should turn out so. Every other philosopher takes it as an *a priori* principle from which conclusions about causality are to be drawn, but which itself need not, perhaps *can*not, be questioned.

8 From here on, for brevity's sake I shall speak generally of points of 'spacetime', in the contexts of both classical and relativistic physics, distinguishing the two cases only if a crucial point rests on the difference in spatiotemporal structure between the two.

9 See Geroch (1996) for a thorough discussion.

10 This point is related to the remark of Stein (unpub., §VI, p. 15), to the effect that the fundamental forces of physical theory are most aptly analogized not with the Aristotelian efficient cause, but rather with the Aristotelian formal cause. For an account of the distinction between dynamical quantities -those for whose evolutions one solves the equations of motion, *e.g.* energy and momentum- and intrinsic quantities -more or less those used for identifying the system in practice, *e.g.* mass and charge- see Curiel (1998).

11 I cannot stress enough that this discussion as it stands is far from adequate. Space constraints do not allow a proper treatment of the issue, for an attempt at which see Curiel (1999).

12 See Norton (1993) for a thorough review of the topic.

13 This is not exactly correct as stated -my intent is to give the flavor of certain technical features of general relativity without being roo technical and without making any false statements, but the two desiderata conflict here. There is a fine theory of integration on arbitrary differentiable manifolds for differential forms, a class of rensor-like objects, and this integration theory can be carried over intact to pseudo-Riemannian manifolds such as general relativistic spacetimes, even ones of arbitrary curvature. This fact does not conflict with my arguments. The desired sort of conservation laws relate the integral of something over a bounded 4-volume to the integral of something over its boundary; the calculus of differential forms requires a 4-form to integrate over the 4-volume and a 3-form to integrate over its boundary, which in the presence of a fixed pseudo-Riemannian metric are essentially equivalent to a scalar field and a tangent vector field respectively.

14 More precisely, the pseudo-Riemannian metric of a generic spacetime has a non-vanishing Riemann tensor.

15 The precise statement is that one cannot define a 'stress-energy tensor for the gravitational field': the only two-index covariant, symmetric, divergence-free tensors that are concomitants of the Riemann curvature tensor are linear combinations of the Einstein tensor and the metric, and these are not viable candidates for a gravitational stress-energy tensor. See Curiel (1999) for details. Similar results about the indefinability of gravitational energy hold in Newtonian graviry, though the situation is somewhat better there in that, in special situations, one can define a local energy density for the gravitational field, which one can never do in general relativity. Again, see Curiel (1999) for details. I felt it necessary to bring the heavy machinery of general relativity to bear against transfer accounts of causality and not rest content with the example of Newtonian gravity primarily because general relativity is the fundamental physical theory, not Newtonian gravity, and I am interested in constraining accounts of causality that have some pretense of being fundamental.

16 See Bondi (1962) for a more detailed argument that one can extract energy from the gravitational field in Newtonian theory, and Geroch (1973) for such an argument in general relativity.

17 The technical reason for this is that point-particles would have to be represented by a mathematical object known as a distribution, which is essentially linear; the Einstein field equation, being non-linear, has no well-posed distributional formulation.

18 Though Maxwell (1965) himself, ever open minded and philosophically skeptical in the best sense of the phrase, left the possibility open that action-at-a-distance theories might prove to be better than local theories.

BIBLIOGRAPHY

Bondi, H.: 1962, 'On the physical characteristics of gravitational waves', in A. Lichnerowicz, A. Tonnelat (eds.): *Les Théories Relativistes de la Gravitation*, Paris, Centre National de la Recherche Scientifique, Proceedings of a conference held at Royaumont in June, 1959, *Colloques Internationaux* Number 91, pp. 129-135.

Boyd, R.: 1983, 'On the current status of scientific realism', *Erkenntnis* 19, 45-90.

Curiel, E.: 1998, The Conceptual Structure of Classical Mechanics as Deduced from the Mathematical Structures of Hamiltonian and Lagrangian Mechanics. Unpublished manuscript.

Curiel, E.: 1999, *Foundational Issues in the General Theory of Relativity*, Ph.D. thesis, University of Chicago. Under preparation for defense in the spring of 2000.

Geroch, R.: 1973, 'Energy extraction', *Annals of the New York Academy of Sciences* 224, 108-117.

Geroch, R.: 1996, 'Partial differential equations of physics', in G. Hall, J. Pulham (eds.): *General Relativity*, Proceedings of the 46th Scottish Universities Summer School in Physics, Aberdeen, July 1995, Aberdeen, Scotland, Scottish Universities Summer School in Physics, pp. 19-60.

Hacking, I.: 1983, *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*, Cambridge: Cambridge University Press.

Helmholtz, H.: 1847[1853], 'On the conservation of force; a physical memoir', in J. Tyndall, W. Francis (eds.): *Scientific Memoirs from the Transactions of Foreign Academies of Science and Foreign Journals*, London, Taylor & Francis, pp. 114-62. Originally read before the Physical Society of Berlin (in German) on July 23, 1847. Translator unknown.

Lewis, D.: 1986, 'Causal explanation', in *Philosophical Papers* 2, 214-241. Oxford, Oxford University Press.

Mackie, J.: 1974, *The Cement of the Universe*, Oxford, Oxford University Press.

Maxwell, J.C.: 1965, 'On Action at a Distance', in W. Niven (ed.): *The Collected Papers of James Clerk Maxwell*, vol. 2, New York, Dover Publications, Inc., pp. 311-323. Originally published in the *Proceedings of the Royal Institution of Great Britain*, vol. VII.

Mellor, D.: 1995, *The Facts of Causation*, London, Routledge, International Library of Philosophy.

Mill, J.S.: 1874, *A System of Logic, Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation*, eighth ed., New York, Harper & Brothers Publishers.

Norton, J.: 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[1992], *Reasoning and the Logic of Things*, Cambridge, MA, Harvard University Press. The Cambridge Conferences Lectures, delivered by Peirce between February and March of 1898, in Cambridge, MA. Edited by K. Ketner.

Putnam, H.: 1973, 'Explanation and reference', in *Conceptual Change*, eds. G. Pearce and P. Maynard, Dordrecht, Kluwer Academic Publishers, pp. 199-221.

Quine, W.: 1973, *Roots of Reference*, LaSalle, IL, Open Court Press.

Reichenbach, H.: 1956[1991], *The Direction of Time*, Berkeley, University of California Press. Edited posthumously by Maria Reichenbach. Reprinted in 1991 with a foreword by H. Putnam.

Russell, B.: 1927[1954], *The Analysis of Matter*, New York, Dover Press. A 1954 reprint of the original 1927 edition.

Russell, B.: 1948[1956], *Human Knowledge: Its Scope and Limits*, George Allen and Unwin, Ltd. Third impression, printed in 1956.

Salmon, W.: 1984, *Scientific Explanation and the Causal Structure of the World*, Princeton, Princeton University Press.

Stein, H.: unpub., How does physics bear upon metaphysics; and why did Plato hold that philosophy cannot be written down? Unpublished manuscript.

Erik Curiel is a Fellow in the Humanities at Stanford University. He is completing his PhD dissertation on the Foundations of General Relativity, from the University of Chicago. He has a number of research projects underway, ranging from the conceptual structure of classical mechanics to an analysis of the two pleasure arguments in Plato's *Republic*.

MUST THE MICROCAUSALITY CONDITION BE INTERPRETED CAUSALLY? BEYOND REDUCTION AND MATTERS OF FACT

Jordi CAT*

* Conceptual Foundations of Science, University of Chicago, 1126 E. 59th St., Chicago IL 60637, USA. Email: Jordicat@midway.uchicago.edu

BIBLID [0495-4548 (2000) 15: 37; p. 59-85]

ABSTRACT: The 'microcausality' condition in quantum field theory is typically presented and justified on the basis of general principles of physical causality. I explore in detail a number of alternative causal interpretations of this condition. I conclude that none is fully satisfactory, independent of further and controversial assumptions about the object and scope of quantum field theories. In particular the stronger causal readings require a fully reductionist and fundamentalist attitude to quantum field theory. I argue, in a deflationary spirit, for a reading of the 'microcausality' condition as merely a boundary condition, inspired by Relativity, that different possible formulations of quantum field theory must obey.

Keywords: Causality, Quantum Field Theory, Microcausality Condition.

CONTENTS

1. Introduction: The problem

In this paper I explore the physical interpretation of the principle known in Relativistic Quantum Field Theory as the 'microcausality condition' (MC). The principle establishes that operators representing physical quantities at space-time points separated by finite space-like distances must commute. In 1928 P. Jordan and W. Pauli presented -as part of a covariant formulation of Dirac's quantum electrodynamics, the theory of the interac-