Three Papers on How Physics Bears on Philosophy, and How Philosophy Bears on Physics

Erik Curiel
Doctor of Philosophy
Department of Philosophy
Division of the Humanities

JUNE 2005

I respectfully and affectionately dedicate this work to Prof. Howard Stein.

It is a pleasure to acknowledge and thank the following people.

Bob Wald—for giving me a desk in the Relativity Group and kindly treating me like one of his graduate students when he had absolutely no reason to do so other than his interest in supporting all sorts of intellectual investigation, and especially for often proving Bob Geroch wrong in our weekly informal seminars over lunch.

Ian Mueller—for exemplifying the spirit of careful scholarship, and for making me realize that sometimes (not often, but sometimes) studying the secondary literature can be almost as rewarding as reading the original text.

Arthur Adkins—for being a scholar, a gentleman and an inspiration *nonpareil*, displaying a stoicism so natural as to seem no burden or affect, but merely the least he could do in living his εὐδαιμόνικος life, all the time with that twinkle in his eye and not-quite-hidden smile.

All the Erics—for never letting me get away with any bullshit while we were plastered at Jimmy's, or anywhere else for that matter.

Michelle and Hisham—for that rarest of comradeships, intellectual, social and just plain damn fun. And for having taken me to see "Celine and Julie Go Boating".

The whole Taylor clan, and George and Meg in particular—for being the most wonderful second family anyone could ever dare hope for.

JDH—for, well, for more than I can well say.

Mom—for giving me a well furnished mind.

Elpis—for supporting me, believing in me, and reminding me that appreciation of all things physical, spiritual and aesthetic is not an alternative to intellectual contemplation, but its most necessary and inextricable companion.

It almost seems silly to, the debt is so great, but, finally, I thank David Malament for his unwavering, supererogatory support and for teaching by example and inspiration what it means to think clearly, carefully and precisely about all matters, both theoretical and practical, all while being the paragon of a *mensch*; Bob Geroch for betting me beers about the answers to physics problems, and for teaching by example and inspiration what it means to think like a philosopher (though he will not thank me for saying so) while being a mathematician and a physicist without peer; and Howard Stein for never having given me an 'A' on any paper I wrote for him, and for teaching by example and inspiration what it means to think like a mathematician and a physicist while being a philosopher without peer. I do not claim to have absorbed any of these lessons; I claim only their attempt and their appreciation.

A student repays his teacher very poorly by never disproving him.
– Nietzsche, *Beyond Good and Evil*

To make a beginning
      out of particulars
To roll up the sum by defective means...

Rigor of beauty is the quest
But how will you find it when it is locked away in the mind
      beyond all remonstrance?

William Carlos Williams, *Paterson*

When science starts to be interpretive
it is more unscientific even than mysticism.

"Self-Protection", D. H. Lawrence

Ahhh, it's so hard, ya' know, it's so hard to believe in anything anymore, ya' know what I mean? It's like religion, you can't really take it seriously, because it seems so mythological and it seems so arbitrary, and then on the other hand, science is just pure empiricism, and by virtue of its method it excludes metaphysics. And I guess I wouldn't believe in anything if it weren't for my Lucky Astrology Mood-Watch.

Steve Martin, "A Wild and Crazy Guy"

How can we lose when we're so sincere?

Charlie Brown

Contents

Chapter 1

Précis of Dissertation

1.1 The Spirit of the Thing

The dissertation consists of three separate papers. The three could not be more different from each other, as determined using a legion of different measures. The most important is this: the earliest one, on physicalist interpretations of causality, though I stand by the substance of its claims, is jejune in style and presentation; the later two, on, respectively, singular structure in relativistic spacetimes and the regime of applicability of theories of mathematical physics, are relatively more mature. If I had had the opportunity, I would have liked to have substituted another paper for the first, or to have rewritten it; but time did not allow this.

The three do possess two related features of great importance, gestured at by the title of the dissertation. On the one hand, all three could be said to manifest a spirit of mitigated skepticism with regard to the interest and indeed the coherence of many questions dealt with traditionally by philosophers, both in their content and in their framing. On the other hand, they also all manifest a spirit that may be thought at first contrary to the first, though it is not, one of guarded optimism for the advancement of real understanding and comprehension of the issues these traditionally philosophical questions have purported to describe, albeit an optimism qualified by a demand intimated as well by the title of the dissertation: that many questions traditionally dealt with by philosophers require knowledge both detailed and comprehensive of our best physical theories, and that of both their formal and empirical content, if one is to make any substantive progress on them.

The time is long since past when anyone, philosopher or physicist, can hope to address in any serious way questions deeply rooted in the physical states of affairs of the world by sequestering himself in his study and requiring of the physical world that it conform, in nature and substance, to the limitations of his imagination and his "powers" of pure reasoning.[1] This is perhaps recognized by most professional philosophers practicing today, and indeed many workers in many fields at least pay lip-service to the idea that, in investigating certain issues, account must be taken of what physics has to say on the matter, even if only to dismiss it in the end as not relevant (which I will be the first to admit is frequently). I applaud this trend—or perhaps merely what I hope is a trend. We'll see.

1. I cannot resist quoting one of Stephen Crane's poems here:

> A man said to the universe:
> "Sir, I exist!"
> "However," replied the universe,
> "The fact has not created in me
> A sense of obligation."

In any event, I have qualms about how these good intentions are often translated into practice. Many times, when a physical theory or principle is introduced in the context of a philosophical discussion, even when the introductor earnestly and attentively tries to do it justice, the treatment of the theory or principle degenerates into the manipulation of simplistic, purely formal toy models, without any serious attempt to understand the implications of the deeper formal structures of the theory, and usually with no attempt at all to consider how the theory finds its actual application in experimental practice. Too often, the lessons we want to draw from superficial treatment of physical theories founder on one of these two rocks (or on both of them at once, when the ship is big enough). Close examination of the formal structure of a physical theory often leads one to conclude that notions and ideas seemingly so fundamental and natural as to be impervious to empirical falsification are, in fact, not only not fundamental and natural, but are not even coherently formulable in the terms of the theory. Examples abound, such as special relativity's dismissal of an absolute temporal structure, and general relativity's refusal to countenance anything like the classical principle of the conservation of energy. While this lesson has been somewhat learnt by philosophers, I think the great majority of them, even among those who try scrupulously to address physical theory in their investigations, do not give the second factor—consideration of the empirical applicability and standing of the theory—a second thought. This practice is common not only among philosophers. It is common among mathematical physicists as well. [2]

When one considers the fundamentals in the schooling our intuitions have received in our contemplation of well worked out examples of physical theories, which by and large tend to include mathematical structures that strike us as 'simple' and 'natural', as applied to the modeling of physical systems whose sole virtue is their analytical tractability, this ought not escape our notice: most such examples of physical theories brandished by philosophers are demonstrably false (Newtonian mechanics and classical Maxwell theory) or have at the moment insuperable problems of interpretation (quantum mechanics) or experimental accessibility (general relativity). We should beware of relying too much on intutions trained in such schools—especially when one also recalls how much of our contemplation and employment of those theories involves models of systems with physically unrealistic perfect symmetries and vaguely jusified approximations, simplifications and idealizations. My earnest attempt in all three of these papers has been to leaven the loaf somewhat.

These three papers also share, if one likes, a spirit of revolt against a naive and virulent form of neo-Aristotelianism persisting still among many philosophers of science and among many physicists today, exemplifed by the idea that there is such a thing as '*the* causal relation', or '*the* proper definition' of a singularity in general relativity, or '*the* relations' that theoretical and experimental knowledge have to each other in the practice of contemporary physics, and that, if physical theories are nagged and worried enough, they will yield up the οὐσία of the thing, without the investigator's ever having paused to consider whether what she asks of the theory is reasonable or even coherent, when framed in its terms. This is like trying to define *the* way to cook sweet potatos at Thanksgiving, without even pausing to consider whether one is making the dinner for a bunch of yankees or for a bunch of confederates.

I contend rather that, when one plays with physical theory to see what it may or may not have to say on the issues one is interested in, one ought to do it with something like the tentative but eager stance of negotiation assumed by a group of children trying to work out the rules of hide-and-seek to apply to the game about to be played, freely adapting the general principles to suit the particular characters of the field of play, the age and condition of the players, temporal constraints on the length of the game, and so on, while still remaining true to the core tenets of the game (for instance, that most of the children will hide and one, or at most a few of the rest, will try to find them). If

2. I presented several years ago, at one of the biennial meetings of the Philosophy of Science Association, a paper entitled "Against the Excesses of Quantum Gravity: A Plea for Modesty", inveighing solely against mathematical physicists for perpetuating these same bad habits. (In the concluding remarks, I also upbraided philosophers for letting them get away with it uncritically, but that was not the thrust of the paper.)

a kid approached the game in its progress, demanded to join, and further demanded that all must play by *the* rules of hide-and-seek, as he lays them down, without his having considered whether they are appropriate for the context, we all know what would (or should) happen to him. I regret only that each physical theory cannot make these ungracious children, the ones who bully it and worry it and take no care to attend to its own peculiar character, eat dirt.

1.2 The Constraints General Relativity Places on Physicalist Accounts of Causality

The attempt in the 1960's to provide a physically based, causal theory of reference, grounded more or less on the idea of energy propagation and energy transference among physical systems, that would be suitable for use in analyzing fundamental physical theories, provides an excellent case-study of these issues. The founding ideas of the accounts—energy and its propagation—turn out not to be coherently definable in one fundamental theory (general relativity) and, in the fundamental theory in which they are more or less definable (quantum mechanics), turn out, by virtue of the sorts of idealizations and approximations we must exercise in order to employ these notions in actual, experimental proceedings, not to have the characteristics those theories of reference require of them, most notably lacking those of continuity, locality and identifiability. The first paper in the dissertation addresses these issues, examining the way that general relativity precludes us from formulating a rigorous, fundamental notion of energy, and of, indeed, any classically conserved dynamical quantity, in its terms, and so stands a hindrance in the attempt to construct accounts of causality based on the ideas of energetic quantities and phenomena.

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 in classical physics. I should emphasize that I am not claiming that one cannot talk about energy at all within the theory. One can speak of it, but only in certain physically special situations, in which 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 itself.

I must emphasize that none of the arguments in the paper 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 of this dissertation do not pretend to bear on any such accounts. Certain sorts of accounts of causality, perhaps best exemplified by Russell in *The Analysis of Matter*, 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 in the paper. I conclude that, in so far as such accounts pretend to be founded on fundamental physical theories and, moreover, pretend to characterize *the* idea of causality, they cannot do so in the terms they are usually formulated in, those of the propagation and transference of conserved, dynamical quantities.

1.3 The Analysis of Singular Spacetimes

In the second paper of the dissertation, the two issues I raised in the introduction to this summary—the need for both a formal and a practical sense of a physical theory one works with, if one is not to be led astray—are this time found to distribute themselves in problematic ways almost equally among both philosophers and physicists. The issue is that of *the* proper definition of 'spacetime singularity' in general relativity, and the question of whether the prediction of such entities (how so ever exactly one characterizes them) leads to a crisis for the theory. The arguments I present are mostly long and mostly technical, not lending themselves easily to a brief, non-technical synopsis. I will remark here only that I draw two main conclusions in the paper. First, how one characterizes such entities is not something that can be settled in the abstract, once and for all, but will depend inextricably on such factors as the purposes and requirements of the investigation at hand, the aesthetic predilections of the investigator, and other such pragmatic considerations. Second, most, if not all, of the arguments claiming that much of the structure often referred to, in the context of general relativity, as 'singular' points to a deep, even a pathological, inadequacy in the structure of the theory itself turn, in the end, on nothing more than just these sorts of aesthetic and pragmatic considerations—what Bob Geroch evocatively calls 'psychology'—and that those who make such arguments, rather than addressing the possibility that the theory is trying to tell us something about novel phenomena that may manifest themselves in our universe and about how we perhaps could go about trying to observe them, rather demand that the theory must be bad, because they cannot conceive of the universe manifesting such phenomena.

1.4 On the Formal Consistency of Experiment and Theory in Physics

The third paper is the most difficult to summarize, as well as the most difficult to digest. In it, I investigate a series of questions on the complex interplay between the theoretician and the experimentalist required for a mathematical theory to find application in modeling actual experiments and, in turn, for the results of those experiments to have bearing on the shaping and substantiation of a theory. On the one hand, we have the rigorous, exact and often beautiful mathematical structures of theoretical physics for the schematic representation of the possible states and courses of dynamical evolution of physical systems. On the other hand, we have the intuitive, inexact and often profoundly insightful design and manipulation of experimental apparatus in the gathering of empirical data, in conjunction with the initial imposition of a classificatory structure on the mass of otherwise disaggregated and undifferentiated raw data gathered. Somewhere in between these extremes lie the mutual application to and qualification of each by the other.

It is one of the games of the experimentalist to decide what theory to play with, indeed, what parts of what theory to play with, in modeling any particular experimental or observational arrangement, in light of, *inter alia*, the conditions under which the experiment will be performed or the observation made, the degree of accuracy expected or desired of the measurements, *etc.*, and then to deduce from the exact, rigorous structure of that theory, as provided by the theoretician, models of actual experiments so that he may judge whether or not the predictions of those models conform to the inaccurately determined data he gathers from those experiments. It is one of the games of the theoretician to abduct exact, rigorous theories from the inaccurately determined, loosely organized mass of data provided by the experimentalist, and then to articulate the rules of play for those theories, by, *inter alia*, articulating the expected kinds and strengths of couplings the quantities of the theory manifest and under what conditions they are manifested, leaving it to the experimentalist to design in light of this information probes of a sort appropriate to these couplings as manifested under the given conditions. Jointly, the two try to find, in the physical world and in the realm of mathematics, common ground on which their games may be played. No matter what one thinks of the status of these sorts of decisions and articulations in science—whether one thinks they can ultimately be explained and justified in the terms of a rational scientific methodology or whether one thinks they are, in the end, immune to rational analysis and form the incorrigibly asystematic bedrock of science, as it were—it behooves us, at the least, to get clearer on what is being decided and articulated.

I do not examine the actual play of the theoretician and the experimentalist in their attempts to find common, mutually fruitful ground on which to engage each other. I leave those issues, fascinating as they are, to other, more competent hands. Neither do I examine all the different sorts of games in which they engage in their respective practices, rather treating only those played in one small part of the playground shared by the theoretician and the experimentalist, that having to do with the comparison of predicted and observed values of a system as it dynamically evolves. I do not deal explicitly with others, such as predictions that have nothing to do with comparison to observations (for instance, the use of Newtonian gravity in calculating trajectories during the Apollo project's flights to the Moon), or the calculation of fundamental properties of physical systems based on theoretical models (for instance, the use of the quantum theory of solids to calculate the specific heat of a substance).

I examine in the paper only what one may think of as the logical structure of the relations between the practice of the theoretician and that of the experimentalist, and, *a fortiori*, of those between theory and experiment. I do not mean to claim that there is or ought to be a single such structure *sub specie æternitatis*, or indeed that there is any such structure common to different branches of physics, or indeed even one common to a single branch that remains stable and viable over arbitrary periods of time, in different stages of the scientific enterprise. I investigate only whether one can construct such a structure to represent some idealized form of these relations. I am not, in the paper, interested in how exactly the experimentalist and the theoretician may make in practice the transitions to and fro between, on the one hand, inaccurate and finitely determined measurements, and, on the other, the mathematically rigorous initial-value formulation of a system of partial-differential equations, whether their exact methods of doing so may be justified, *etc.* I am rather concerned with the brute fact of its happening, whether there is indeed any way at all of constructing with some rigor and clarity a model of generic methods for doing so. Having such a model in hand would show that there need be no gross logical or methodological inconsistency in the joint practice of the theoretician and the experimentalist (even if there is an inconsistency in the way physicists currently work, which I would not pretend to hazard a guess at). Indeed, it is difficult to see, on the face of it, how one may comprehend these two to be engaged in the same enterprise in the first place, difficult, indeed, to see even whether these two practices are in *any* sense consistent with each other, since it is not even clear what such consistency may or may not consist of. While I seriously doubt that any formal analysis of the relations between theory and practice I or anyone else may propose could answer this question definitively with regard to a real physical theory and its experimental applications, the sort of analysis I attempt to outline, if successful, would perhaps have the virtue of underlining the sorts of considerations one must take account of in judging the consistency of a real theory. This may seem a Quixotic project, at best, on the face of it, but I think I can say a few words in defense of its interest. In defense of its feasibility, I offer the paper.

Without a doubt, one can learn an extraordinary amount about a physical theory by examining only its structure in isolation from the conditions required for its use in modeling phenomena, as is most often done in philosophical discussion of a technical nature about physical theories in particular, and about the character of our understanding of the physical world in general. I argue, however, that comprehensive understanding of a physical theory will elude us unless we examine as well the procedures whereby it is employed in the laboratory, and, moreover, that comprehension of the nature of such knowledge as we may have of the physical world will similarly elude us without a serious attempt to understand both the theoretical and the practical characters of that knowledge. In particular, the question I plan to address is not how one gets to a system of exact partial-differential equations from inaccurate data; nor is it how one gets from exact solutions of partial-differential equations to predictions that may or may not accord with actually observed, inaccurate data (though this latter is touched upon *en passant* to some degree). It is rather a question of the consistency of, perhaps the continuity between, the two—a question, if you like, of whether the theoretician and the experimentalist can be understood as being engaged in the *same enterprise*, that of modeling and comprehending the physical world, in complementary, indeed mutually supportive, ways. The answer I propose is constructive—a proposal for a more or less formal, explicit method of representing the connection between the stocks in trade of the two that remains true to the character of these two stocks. Another way of putting the point: philosophers, when having tried to understand the relation between theory and experiment, tend to have been vexed by the problem of how a theory gets into (and out of!) the laboratory, often framed in terms of the putatively inevitable "theory-ladenness" of observations; I am concerned with what one may call the converse problem, that of getting the laboratory into the theory, and the joint problem, as it were, whether the theory and the laboratory admit at least in part a consistent, common model. Along the way, I present an argument, in large part constituted by the body of the construction itself, that the initial-value formulation of the partial-differential equations of a theory provides the most natural theater in which this sort of investigation can play itself out.

I focus the discussion around the idea of the *regime of applicability* of a physical theory. From a purely extensive point of view, a regime of a physical theory, roughly speaking, consists of the class of all physical systems *cum* environments that the theory is adequate and appropriate for the modeling of, along with a mathematical structure used to construct models of these systems, and a set of experimental techniques used for probing the systems in a way amenable to modeling in the terms of that structure. It can be represented by, at a minimum:

1. a set of variables representing physical quantities not directly treated by the theory but whose values in a given neighborhood are relevant to the issue of the theory's applicability to a particular physical system in that neighborhood, along with a set of algebraic and differential expressions formulated in terms of these variables, representing the constraints these ambient, environmental quantities must satisfy in order for physical systems of the given type to be susceptible to treatment by the theory when they appear in such environments
2. a set of algebraic and differential expressions formulated in terms of the variables and constants appearing in the theory's system of partial-differential equations, representing the constraints the values of the quantities represented by those constants and variables must satisfy in order for the system bearing those quantities to be amenable to treatment by the theory; these expressions may include as well terms from the set of variables representing relevant environmental quantities
3. a set of algebraic expressions formulated in terms of variables representing the measure of spatiotemporal intervals, constraining the character of the spatiotemporal regions requisite for well-defined observations of the system's quantities to be performed in; these expressions may include terms from the set of variables representing relevant environment quantities, as well as from the set of variables and constants appearing in the theory's system of partial-differential equations
4. a set of methods for calculating the ranges of inaccuracy inevitably accruing to measurements of the values of the system's quantities treated by the theory, depending on the sorts of experimental techniques used for probing the system, the environmental conditions under which the probing is performed, and the state of the system itself (including the stage of dynamical evolution it manifests) at the time of the probing—these methods may include, *e.g.*, a set of algebraic and differential expressions formulated in terms of the variables and

constants appearing in the theory's system of partial-differential equations, the variables representing the relevant environmental factors, and the variables representing the measure of spatiotemporal intervals

5. a set of methods for calculating the ranges of admissible deviance of the predictions of the theory on the one hand from actual measurements made of particular systems modeled by the theory on the other, depending on the sorts of experimental techniques used for probing the system, the environmental conditions under which the probing is performed, and the state of the system itself (including the stage of dynamical evolution it manifests) at the time of the probing—these methods may include, *e.g.*, a set of algebraic and differential conditions formulated in terms of the variables and constants appearing in the theory's system of partial-differential equations, the variables representing the relevant environmental factors, and the variables representing the measure of spatiotemporal intervals

The idea of a regime is perhaps best illustrated by way of an example. For the theory comprising the classical Navier-Stokes equations to model adequately a particular body of fluid, for instance, elements of its regime may include these conditions and posits:

1. the ambient electromagnetic field cannot be so strong as to ionize the fluid completely
2. the gradient of the fluid's temperature cannot be too steep near equilibrium
3. only thermometric systems one centimeter in length or longer are to be used to measure the fluid's temperature, and the reading will be taken only after having waited a few seconds for the systems to have settled down to equilibrium
4. the chosen observational techniques to be applied, under the given environmental conditions and in light of the current state of the fluid, yield data with a range of inaccuracy of $\pm 1\%$, with a degree of confidence of 95%
5. a deviance of less than 3% of the predicted from the observed dynamic evolution of the system's temperature, taking into account the range of inaccuracy in measurement, is within the admissible range of experimental error for the chosen experimental techniques under the given environmental conditions, in light of the current state of the fluid

I do not pretend to offer in the paper a definitive analysis of the concept of a regime or indeed of any of its constituents. I rather sketch one possible way one may construct a (moderately) precise and rigorous model of the concept, with the aim of illuminating the sorts of questions one would have to answer in order to provide a more definitive analysis. The hope is that such a model and correlative demonstration may serve as a contructive proof of the formal consistency of the practice of the experimentalist and the practice of the theoretician in physics, at least in so far as one accepts the viability of the sort of formal model I construct, indeed, as a construction of the common playground, as it were, of the two, playing with the toys and rides and games of which we may pose precise questions of a technical nature about the interplay between theory and experiment, and attempt to answer such questions.

I conclude that, not only are theory and experiment consonant with each other, they are mutually inextricable—not, however, as equals. Theory plays Boswell to the subtle and tragic clown of experiment's Johnson.

The entire dissertation, if you will, may be considered an exercise in approximation and idealization in the philosophy of physics.

Chapter 2

The Constraints General Relativity Places on Physicalist Accounts of Causality

> 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 κτῆμα ἐς ἀεί,—*semper eadem*—is the *name* of it.
>
> Charles Sanders Peirce
> *Reasoning and the Logic of Things*

2.1 Introduction

It is well known, at least among those who think about these things, that energy is not a fundamental quantity in the theory of general relativity. 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. This 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 [1] and special relativity, and even then only by employing explicitly approximative and idealizing techniques that are not part of the theory *per se*—the Einstein field equation does not of its own nature exhibit simplifying modifications to itself appropriate for the weak gravitational field regime, much less justify the use of such modifications in any particular circumstances. Consequently talk about energy reflects nothing fundamental about the theory itself.

In this paper, I discuss how this fact may place restrictions on the sorts of accounts of causality that can be considered fundamental *vis-à-vis* general relativity. Some accounts of causality purport

1. 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.

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 of this paper do not pretend to bear on any such accounts. Certain sorts of accounts of causality, perhaps best exemplified by (Russell 1927), 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, in so far as this concept represents actual structure of the physical world as modelled by our best physical theories. 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 will refer to all such accounts of causality as *transfer accounts*, gesturing at the fact that these sorts of accounts take the transfer (and the propagation) of energy to play the most important role in constituting the causal relation. 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.

My arguments, though likely affronting to some dearly held contemporary intuitions, should not, I think, be terribly surprising. The ideas of propagation and transfer, as we understand them, have not been associated with the notion of causality commonly held by the intelligentisia for long, certainly not for more than 400 years, since the time of Galileo.[2] Before the scientific revolution, such ideas were no part of generally held conceptions of causality among those who contemplated such matters. 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.[3] 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 do not think we should feel any reluctance to jettison some or even all of the contemporary notions of and intuitions about causality, in so far as they may apply at a fundamental level. At least, we should feel less reluctance to give up lessons learned early and often enough by the child to be termed 'intuitions' by the adult, than we should feel to ignore what our best physical theories seem to be trying to tell us about the character of the world. Philosophers, when having investigated questions of this sort, have tended to focus their attention on the lessons to be drawn from quantum mechanics. There are natural reasons for this, perhaps the most important being the great difficulty of the technical machinery one must master in order to study general relativity in any depth and with any breadth, as compared to that required for the examination of (non-relativistic) quantum mechanics. When they do turn their attention to general relativity, it too often happens that philosophers approach it with a raft of traditional concepts, questions and issues already in hand and ask how general relativity bears on them, as in the case of the time-worn debate

2. I do not think that the medieval doctrine of the identity of cause and effect involves a notion of continuity even vaguely analogous to ours. The character of the concepts in the terms of which medieval philosophers formulated their ideas and arguments differ so much from those used today as to make meaningful comparison impossible.

3. I fault Hume's misunderstanding of Newtonian theory for blinding much of the philosophical world to this fact, even down to this day. See (Stein 1994) for a brief discussion.

on substantivalism versus relationalism, rather than studying general relativity both broadly and in detail to see what questions the theory itself naturally suggests we attend to and what concepts and structure it offers up as the most natural in the terms of which to formulate these questions. Many traditional questions and issues do not seem to get even so much as a secure footing in general relativity when the time and care is taken to ascertain what sort of structures one needs to have in place to ask the frame them sensibly; and many questions longing to be asked remain wall-flowers. [4]

This paper occupies a funny middle ground between what I have painted as virtue and vice: it asks no wall-flowers to dance, and it does focus on questions and concepts originally mooted in the theater of classical physics; it does nevertheless attempt to show why this strategy can be misleading, by demonstrating the disharmony of the arguing that the classical concepts used to frame the questions do not fit within the structure of general relativity without much shoe-horning and discomfort. This paper does not pretend to be a mathematical demonstration of a certain philosophical point. 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, whereas general relativity does not share the features of those theories that made the suggestion so plausible. [5]

I begin in §2.2 by discussing the motivations behind transfer accounts of causality. In §2.3 I attempt to make precise in a schematic way the relations between energy and causality that transfer accounts presuppose. I will not attempt to give an exact analysis of energy as represented in any particular physical theory; rather I will assume some general propositions about energy that are part of the widely accepted folklore about physics, just to get the ball rolling. Next, in §2.4 I examine in detail the properties energy must have in order to play the roles demanded of it by transfer accounts. I will argue that, if the relations introduced in §2.3 to hold between energy and causality, then one must be able to formulate something like the classical principle of the conservation of energy. A slight generalization of this argument will show that an appropriate conservation law is the necessary condition for any stuff's standing in such a relation to causality, from momentum to matter to electromagnetic field intensity—if one wants the relation of cause to effect to be mediated by the propagation and transfer of some stuff from the one to the other, then that stuff must be capable neither of being created nor of being destroyed. Finally, in §2.6 I turn to the theory of general relativity. In the theater of a generic general relativistic spacetime, not only can such relations not hold between causality and energy as transfer accounts require, but such relations cannot hold between causality and *any* stuff representable in the theory. The structure of general relativity, in so far as it precludes the requisite sorts of conservation principles, militates against this type of account of causal relations. I conclude with a brief description of the sorts of conceptions of causality that general relativity is not overtly hostile to.

Before beginning these arguments, I should say a few words about quantum mechanics, and why I will steadfastly ignore it after these few words. The aim of this paper is twofold. The first goal

---

4. Many of Russell's remarks in *The Analysis of Matter* (1927) and Eddington's in his *Mathematical Theory of Relativity* (1921) show a great sensitivity to general relativity's demand for new questions and new issues, and not a stale re-hashing of centuries old debate. Many contemporary philosophers, including (Stein npub), (Malament 1977), (Earman 1995) and (Norton 1985), have what I think are similar views. I believe, among contemporary physicists, (Penrose 1968) and (Geroch 1973) do as well.

5. Many among both philosophers and physicists have retrojected the idea of locality (widely accepted in the physics community only since the turn of the century) 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 will 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.

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 any specific attributes such a grander theory ought to manifest—the grandiquolently speculative claims of superstring theorists notwithstanding—there is extremely good reason to think that 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 $\frac{1}{r^2}$ 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 which 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 acquire 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. The greatest obstacle, as already remarked, 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 no good way to adjudicate among a swarm of competing schemes for, *e.g.*, providing approximate solutions to the general relativistic equations of motion for small, dense bodies; producing closed, analytical solutions to these equations is out of the question.[6] We also have only a superficial understanding of how to catalogue solutions to the Einstein field equation according to physically relevant, generic features they may share,[7] I think that neither the arguments nor conclusions of this paper constitute a real advance in our comprehension of general relativity, along these or any other lines. I rather hope only that it may make a small part of what we do comprehend available to a broader audience.

These difficulties by themselves would perhaps not suffice as a reason to ignore the tenets of quantum mechanics in so far as they bear on accounts of causality, but I think they do when taken together with the character of what it is we do know about quantum mechanics: 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 an understanding of how to model phenomena in general relativity. Foundational questions in quantum mechanics at the moment are so turbid and disputed as to allow a defense of almost no clear, precise propositions about causality, except for the obvious: proceed with extreme caution. Perhaps this paper may make the case by example that need need not be the case in general relativity.

6. See *e.g.* (Quinn and Wald 1999) for a survey of the issue. To be more precise, the problem is usually posed as finding solutions to appropriately simplified (*i.e.* 'idealized'—a nobler term) equations of motion.

7. See (Curiel 1999) for a brief discussion of this point, in the context of examining proposed classifications of singlur structure in relativistic spacetimes.

2.2 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;[8] it was also generally recognized to have 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.[9] 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, 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.[10] Salmon, inspired by (Reichenbach 1956), comes close to identifying causal processes as exactly those that transmit energy.[11] (Russell 1927), (Russell 1948) and (Reichenbach 1956) predicate their conclusion that transfer of energy is intimately connected to causal relations upon a thorough analysis of scientific theory, including the working through in detail of examples of energy propagation and transfer in various physical situations.

On the face of it, the idea that energy and causality, as these notions may arise in or be suggested by physical theory, 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 (ignoring the constraints imposed by the second law of thermodynamics)—by itself suggests such a relation: a cause is, roughly speaking, something that produces a change in something else, and changes 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: that between any entity A and any other entity B such that A causes B there is always a third entity Q such that A causes Q and Q causes B.[12] A related principle holds that, if A causes B, where A and B are spatiotemporally separated, then the causal efficacy "travelled" from A to B via Q and *not* via Z, where Q is "spatiotemporally between" A and B and Z is not. The thought is that one can draw a more or less narrow tube through spacetime both such that the loci of all links in the causal chain lie inside it, and such that it respects the null-cone structure of spacetime ("the causal efficacy propagates more slowly than light"). All propositions of this sort I will group together under the rubric *the principle of causal continuity*. I am convinced that most of the attraction of transfer accounts of causality arise from this idea of causal continuity and its "explanation" by reference

8. *Cf.* (Helmholtz 1853, *passim*).

9. 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.

It is also a suggestive—and highly obscure—fact that one of the other primary developers of the principle of the conservation of energy, Mayer, relied heavily on arguments based on the scholastic doctrine of the continuity and indestructibility of causes, as encapsulated in the apothegm *causa æquat effectum* ("causes equal effects")—*cf.* (Mayer 1842, *passim*).

10. See (Quine 1973, esp. p. 5).

11. See (Salmon 1984, esp. p. 146).

12. In fact, 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 and unnecessary technicalities. I trust the sense of what I intend is clear without them.

to the continuity of energetic processes.[13] Last, but certainly not least, the linking of energy and causality would appear to provide a straightforward solution to one of the most vexed problems in this vexed area of inquiry, that of winnowing out the true causal processes and chains from sequences of correlated junk—the true causal chains are those along which energy propagates in the proper fashion. This alone would make it worth fighting for.

I think Russell and Reichenbach were right to try to ground such views of causality upon a thorough analysis of fundamental physical theory, including the working through of several concrete examples, for such views presuppose several substantive theses about the structure of the physical world (or at least about the structure of our best physical theories): that, for example, a dynamical quantity called 'energy' exists and has at least 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 at least brief periods of time, but also so that one can determine, in at least a certain type of case, 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'; that when a system is 'interacting with its environment' it is possible, in at least some cases, to identify the precise bits of the environment it is interacting with, which is to say the source or sink of any energy it is acquiring or losing during the interaction; and so on. None of these seem on their face particularly contentious. A conception of causality, however, that requires such theses and that has the stated goal of being fundamental—in the sense that it finds its motivation or justification in fundamental physical theory—just because it presupposes such substantive theses about the physical character of the world, requires for its justification a thorough investigation of the best going current physical theory or theories, if not to ground it satisfactorily at least to demonstrate that such assumptions as these, and whatever others the analysis of such an account of causality turns up, do not overtly conflict with any of the precepts of our best current science. One would for instance have to ascertain 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—that the theory supports the existence of energy and predicts that it will be exchanged in certain situations does not guarantee *ipso facto* that energy will be exchanged when and how a particular such would account of causality requires. 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 our best theories, in the same way as classical particle mechanics is thought to require no interpretation.

2.3 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—Hume's game of billiards, say. The usual story says that when such a body in motion, A, strikes such a body at

13. It is astonishing how widespread such an assumption of causal continuity is in the philosophical literature, especially when one considers how recently the physics adopted analogous principles. Under the influence of Newton's theory of gravity, action-at-a-distance theories reigned until the late 19th century, and even today such comprehension of quantum mechanics as we have does nothing to encourage such views. Philosophers as disparate in temperament and aim as (Ducasse 1926), (Russell 1927), (Russell 1948), (Reichenbach 1956), (Quine 1973), (Mackie 1980), (Bunge 1979), (Salmon 1984), (Lewis 1986), (Mellor 1995) and many others too numerous to mention have all invoked in discussions of causality, more and less crucially, 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 cannot, be questioned.

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 are floating in 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. A few quadrillion photons may be hitting them, a few quintillion neutrinos passing through them, but it is an excellent approximation to treat them as free bodies; in particular, they themselves are slight enough so that any gravitational force they may exert on each other or on themselves may be safely ignored. Assume we are observing the system in a laboratory co-moving with B, so it appears to be at rest (recall that this example is taking place in Newtonian spacetime, so this assumption is viable). By hypothesis, before the collision A propagates with a uniform velocity 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 moves along 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 seems that this is only a necessary and not a sufficient condition. We would likely want to say, under many circumstances that a ball swung in a circle at a constant rate on the end of a string interacts with the string and with whatever holds the string, depending on the nature of the investigation, our purposes in studying the system, the sorta of machinery we prefer to employ to model such systems, *etc.* In this case, we may predicate the claim that the ball interacts with the string on the fact that, as formulated using Newton's second law of motion, in order to solve the ball's equation of motion, one must know the force the string exerts on the ball. [14] Still, because this force acts everywhere perpendicularly to the ball's acceleration—the string performs no work on the ball, and vice-versa—the ball exchanges no energy with the string. [15]

B's center of mass does not start into motion immediately at the instant of contact, just as the state of motion of A's center of mass does not alter immediately at the instant of contact—this would happen only if the two bodies were perfectly rigid, but no bodies in fact are. [16] The shock waves the impact generates in the bodies take some small but finite time to propagate and reverberate through the respective bodies, ultimately setting the center of mass of each into roughly uniform motion apart. More precisely (though not much more), when the two bodies come into contact the bit of B's surface that gets hit is deformed ever so slightly inward under the force exerted on it through the rapid deceleration of the bit of A's surface hitting it. This deformed part of B exerts a force as it bends inward on bits of B further inside, as those bits inside resist altering their relative position in virtue of their stable cohesion, and so on, all throughout B (and similarly for A). Mechanical work is defined as a force acting through a given distance in the same direction as the force points, so in the first instant the bits of A in contact with B do work on B; in the second instant the bits of B deformed inward do work on bits of B further in towards its center, and so on. By hypothesis again, the bodies are rigid enough so that it is an excellent approximation to set to zero the time between the instant of contact and the instants at which the states of motion of their respective centers of mass alter; to set to zero the amount of kinetic energy converted by the impact into gross internal vibratory motion of each body as it flies away; and to set to zero the amount of kinetic energy that gets transformed into thermal energy in both A and B through the random thermal fluctuations excited in their constituent particles during the time the waves

14. This is not true, for instance, in the Hamiltonian formulation of the ball's equation of motion.

15. We are ignoring such niceties as the change in the string's thermal and internal electrostatic energy induced by its being stretched.

16. Indeed, in relativity theory one cannot even formulate a notion of absolute rigidity. One has only so-called Born rigidity (see, *e.g.*, (Pauli 1921)) as the best approximation to the idea in classical physics.

of deformation pass through them immediately after contact. The net result is that A's direction of motion changes 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 vector sum of A's and B's linear momenta after impact equals A's linear momentum before impact, and similarly for angular momentum defined relative to any arbitrarily fixed point.

Were one inclined, say, to a point of view similar to Hume's, 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; if one followed Al-Ghazzali, on the other hand, one would conclude that the will of god served as the only guide keeping the behavior of these phenomena, as manifested at each instant, in conformity with the regularity we ascribe to causation, though in fact we would think that no such relation obtained between any different events or bodies *per se*. In either case, one could say nothing about any possible relation among dynamical systems, intrinsic to them alone, not depending in any way on the way god's will intercedes at every moment to impose structure on the world, and not depending on the history and constitution of one's own mental faculties, in virtue of which one (state of a) dynamical system might justifiably be said to be the cause of (a state of) another. For one who champions the idea that causality is a feature of the physical world, however, not reliant on the continual intercession of god, and independent of the history and constitution of our mental faculties, 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 absolutely crucial supposition here is that the energy manifested by B as it starts into motion "came from" A, in some sense or other: 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 were able somehow to determine that the kinetic energy B suddenly had acquired came not from A but had been transmitted in some immediate, 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 instantaneously been transferred to the distant body Z), or even that B's newfound kinetic energy had simply sprung spontaneously into being at that moment *ex nihilo*, and likewise that A's lost energy simply vanished *ad nihilum*, one would naturally conclude that A's coming into contact with B had not been the cause of B's starting into motion, but rather the transfer of energy from Q or the spontaneous creation of energy had been.

No sane person, I wager, could bring himself to swallow the Gargantuan global conspiracy required for these sorts of events always to be occurring in accord with the principle of the conservation of energy; for, with no physical theory to account for exactly which other bodies Q and Z might be involved in the transferrals of energy and why, or to pick out the conditions under which energy might vanish and appear spontaneously, a global conspiracy is the only explanation. The total conservation of energy during the process, predicted theoretically and verified experimentally, serves, it is claimed, to justify the idea that B's newfound energy once had been A's, and it is this tangible link between the two that is to provide the ground for the causal claims one wants to make about this process: that A's propagation along that path, a manifestation of its kinetic energy, caused it to come into contact with B; that the impact of A and B caused both B to start into motion *and* A's state of motion to change;[17] that A's state of motion at earlier times was the mediate cause of B's state of motion at later times. 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

17. Note the peculiarity of this instant in the proceedings, that one is forced to make two apparently distinct causal claims about it rather than only one. This reflects the fact that the impact of A and B mediates an *inter*action of the two systems, and not merely the unidirectional action of one "active" system on another "passive" one, so to speak. Philosophers have often ignored this interactive aspect of physical phenomena, which has led to much confusing and confused ink spilled on the question of the 'directionality' of causality. See (Stein npub) for a more thorough discussion of closely related topics.

conditional that mathematical physics provides—'If A strikes B thus, A and B will move so'—and into the *recherché* realm of the causal.

What roles, schematically stated, must energy play in order to support causal judgements? Its primary job is to provide warrant for asserting judgments of the form "C causes E" when the solution to the equations of motion 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 (and so will necessarily obtain at different times), as when the state of $\mathcal{A}$ above at the time it contacted $\mathcal{B}$ 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 $\mathcal{B}$ at the time it started into motion is said to have been caused by the state of $\mathcal{A}$ at that time. Finally, in still other cases they will be states of different systems at different times, as when the state of $\mathcal{A}$ at some early time is said to have been the mediate cause of the state of $\mathcal{B}$ 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 transfer of energy, 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 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.

2.4 The Propagation and Transfer of Energy

Consider first propagation. There is no such thing 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 to be 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*, and, in any of these cases, it would not be just at all. The photon itself is the system, and the rest are its attributes.[18] The famous relativistic equation of mass and energy might lead one to think otherwise, but it is beside the point, for mass itself is only one more possible attribute of a system and does not by itself constitute a substance capable of supporting attributes. 'This material thing here, with a certain mass, has this momentum' is a meaningful scientific proposition, but not 'this quantity of mass here has this momentum' (excluding the colloquial use of 'mass' to refer to a material system). 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, therefore, one must assume a physical system that propagates, to which the energy is attributed. In order for the propagation of this energy, as carried by the system, to support the desired kinds of causal claim, we must know what it means, in classical physics, to identify a physical system, at one point of space at one moment of time, as being the same in a substantial sense as a physical system at a different point of space at a different moment of time, and, in relativistic physics, to do the same for systems at different points of spacetime. (From here on, for brevity's sake I will speak generally of points of 'spacetime', in the contexts of both classical and relativistic physics, distinguishing the two cases only when a crucial point rests on the difference in spatiotemporal structure between the two.) Besides knowing that the system bearing the energy is the same over time, in some substantive and relevant sense, moreover, one must also

18. One should take this "substance/attribute" talk with a pinch of salt. Even in classical physics, for instance, it is likely more proper to represent quantities such as momentum as relations rather than as attributes—a relation, say, between a dynamical system and an orthonormal tetrad on a region of spacetime. The scholastic jargon I hope only drives home more forcefully the point I am trying to make.

know that the system is of an appropriate sort to exchange energy with the effected system, both of the amount and with the rate of exchange required, which means that one must have a classification of physical systems based on some substantive and relevant differentiæ on the basis of which one can found such judgments. In somewhat more technical terms, one must know, for a given kind of physical system, the types of coupling with other systems (physical interaction) it can manifest, with what other sorts of systems these couplings can occur, and what the strength of such couplings may be in all the possible cases (how much energy will be exchanged, at what rate, in which direction, *etc.*). In our example of the balls, for instance, we are confident that the quintillions of neutrinos passing through B at the moment it leapt into motion had no role in this effect. We know this not because the neutrinos do not carry enough total energy to do this—in fact, they do—but rather because the neutrinos do not interact with, do not couple with, ponderable bodies in a way that has as part of its character that kind of exchange of energy. We know this, furthermore, because of the knowledge we have of the type of system a neutrino is, and, in particular, the knowledge we have of the admissible couplings, and the strengths of those couplings, neutrinos manifest with other types of systems.

In sum, in order to use energetic quantities and processes as the foundation of an analysis of causal relations, we require at a minimum the following capacities:

1. to identify a physical system as belonging to particular category of physical systems
2. to reidentify an individual system over time, in the course of its dynamic evolution
3. to quantify and measure to the required degrees of accuracy and exactitude the values of the physical quantities borne by the system we hold responsible for the causal efficacy in any given case

With this knowledge in hand, the idea is, we can rely on the quantitative agreement in the total amount of energy as it distributes itself among different systems in the course of particular interactions, as guaranteed by the principle of the conservation of energy, to give sense to claims of the form 'the transferral of energy *from A to B* at the moment of their contact caused B to start into motion'.

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. To see why, fix a given physical entity. Its state at a "particular instant of time" is represented by a bounded region of a spacelike hypersurface ("the region of space the body occupies at that instant"), with a particular attribution of values, in that bounded region, for the physical quantities it bears. [19] A set of equations of motion determines the appearance and behavior of the entity at later times. With very few exceptions, all known sets of such fundamental equations of classical and relativistic physics have the following character. Starting from the given values of the entitity's properties in the initial region of spacetime, the equations will have a unique solution for at least some finite time into the future. Based on this solution, one can construct a continuously varying family of continuous, mutually disjoint, timelike paths, such that

19. We are glossing over many questions of interest and difficulty, that are beyond the scope of this paper. For example, it is not clear to me that one can consistently and coherently frame a precise notion that captures what we seem to be gesturing at with the idea of "all the physical properties borne by a physical system". Consider, for example, a body of viscous fluid. What are "all its physical properties"? Those such as its temperature, gross fluid velocity, its hydrostatic pressure, its state of shear and internal stress, its coefficient of thermal conductivity, and all the rest one needs to model the theory using the Navier-Stokes equations? Must one also include the acceleration field of all its constituent molecules? The quantum state of each of the atoms in each of its molecules? The SU(3) representation of all the strong-force interactions occurring in the nucleus of each of those atoms? I do not think that these are questions that can be answered in the abstract. One must rather have the context of a particular sort of physical investigation, employing a specific set of theoretical and experimental tools, in order to make the fundamentally pragmatic decision about the way one will model the system, including the physical quantities one treats as attached to it during the time one models it. See (Curiel 2011) for a discussion of these issues.

1. the paths collectively form a four-dimensional, solid, spatiotemporal tube, through every point in the interior of which passes exactly one path in the family
2. each path in the family has a point in the initial bounded region as its endpoint, and every point in the bounded region is the endpoint of exactly one of the paths in the family
3. if one parametrizes the curves by proper time,[20] and one fixes some small enough positive real number δ, then the collection of points consisting of the point on each path a time δ later than that of the initial region forms a bounded region of a spacelike hypersurface ("the region of space the body occupies at that later instant"), and the physical system occupying that region is of the same type as the system in the initial region, in the sense that one can characterize its state using the same physical quantities, and the same set of equations of motion determine its appearance and behavior at later times

This is just to say that, in their guise as differential equations, the system's equations of motion possess well-posed initial-value formulations.[21] It is this predictably homogeneous aspect of dynamical systems, a guaranteed consequence of the *form* of their equations of motion, that ultimately underwrites our identification of them over spatiotemporal intervals.

This last remark, on the fundamental role played by the form of the equations of motion, demands an explication, by way of a brief detour. Fix, for the purposes of this detour, a physical system and the system of partial-differential equations it obeys. An initial-value formulation of the system consists of an attribution of values to all the dynamically relevant quantities of the system—at least, "all" with respect to the theory modeling the system using the partial-differential equations at issue—in such a way that the partial-differential equations have a unique solution for some finite interval of time, and a solution, moreover, that is stable in the sense that small (in a technical sense) perturbations in the initial conditions yield small (in a technical sense) differences in the induced solutions. For example, one could with equal justice say that both the mean kinetic energy of the molecules of a Navier-Stokes fluid and its gross, fluid temperature are, in some sense, "dynamically relevant quantities" it possesses. If one is using a thermodynamical theory to model the fluid, however, say one comprising the classical Navier-Stokes equations, one will treat the temperature and not the mean kinetic energy of the molecules, whereas using a statistical theory to model the fluid, *à la* Maxwell-Boltzmann, one will treat the mean kinetic energy of the molecules and not the temperature.[22]

Now, the equations of dynamical evolution most often thought of as "fundamental" by the physicist are of a particular mathematical type[23]—they are quasi-linear and hyperbolic.[24] Hyperbolic equations have three particularly nice, inter-related properties for the representation of the dynamical evolution of physical systems not shared by other types of partial-differential equations:

20. This means that any observer whose worldline instantiates such a path will, over the course of any segment of the curve, record an interval of proper time as having passed numerically equal to the metrical length of the segment.

21. See (Geroch 1996) for a thorough discussion.

22. See the third paper in this dissertation, "On the Formal Consistency of Theory and Experiment in Physics", for an expanded discussion of this point.

23. I am skeptical of the idea of "fundamental" theories in physics. Consider quantum field-theory. It can not solve in closed form the dynamical equations representing the evolution of arguably even the simplest micro-system, the isolated hydrogen atom. It rather relies on perturbative expansions, and thus requires the system to be not too far from equilibrium of one sort or another. Quantum field-theory in general, moreover, can not handle phenomena occurring in regions of spacetime in which the curvature is too large. The Standard Model breaks down in regimes far above the Planck scale. Not even quantum field-theory formulated on curved-spacetime backgrounds can deal rigorously with phenomena under such conditions. I know of no theory of quantum gravity mature enough for it even to attempt the claim that it can be thought of as fundamental. And when, as I sincerely hope, one of these theories does mature and gain primacy, it will have no more warrant for proclaiming itself the ultimate bedrock than any of its predecessors.

24. See, *e.g.*, (Sommerfeld 1964) for a discussion of this type of equation, as well as the parabolic and elliptic types, with regard to how their respective properties bear on the modeling of physical systems. I discuss in the third paper of this dissertation the privileged role hyperbolic equations appear to play.

1. their characteristic "wave-fronts" propagate at strictly bounded speeds (the maximal speed depending on the particulars of the equations at issue)
2. their solutions are not necessarily analytic fields
3. discontinuities in initial data propagate in a continuous fashion through the solution to an initial-value formulation using that initial data

The first implies, among other things, that one can guarantee that the system modeled by the equations propagates at speeds less than that of light, as demanded by relativity. The second means, roughly speaking, that arbitrarily distant systems can not "influence" its evolution. [25]

The third requires a more involved explanation. Consider a metal rod at a uniform temperature, 10^oC higher than that of the ambient environment, say, air. If one represents the temperature of the entire system, the rod plus its environment, by a scalar field, in the context of a grossly thermodynamical theory, then that scalar field will have a discontinuity on the two-dimensional, spatial surface determined by the boundary of the rod—it leaps (or falls, depending on which way one is going, so to speak) by 10 degrees at all points on that boundary. If one models the subsequent thermal evolution of the system using a non-hyperbolic partial-differential equation, say, Newton's law of cooling, which is parabolic, starting from these initial conditions, then at *any* finite time after the initial instant the solution representing the joint state of the system, rod *cum* environment, will be an analytic field. The discontinuity in the initial data has been smoothed out and the values of the temperature and its derivatives at any point of the system equally determine the value of the temperature and its derivatives at all other points of the system, instantaneously. If one uses a hyperbolization of Newton's equation, [26] formed, for example, by adding to its lefthand side terms purporting to represent "relaxation effects", perhaps in the person of higher-order moments of the distribution function, then, in the absence of other interferences, at any finite time after the initial moment, the solution will exhibit a discontinuity in the value of the temperature, still at the spatial boundary of the rod. The measure of the discontinuity, moreover, the scalar field on the boundary representing the jump in value of the temperature, will, in general, itself be a continuous and continuously varying field. I'll refer to this sort of case as the *propagation of a discontuinity in value*.

This system bears another, more subtle, and perhaps more important, possible discontinuity, closely related to the one just discussed: a discontinuity in the matter and the form, as it were, of the partial-differential equations used to model the two parts of the system, rod and air, in the way most appropriate for the requirements of the investigation at hand. Consider first a discontinuity in the matter only. In this case, equations the same in form—both being instances of Newton's law of cooling—model the air and the rod respectively. The two equations differ in the values of the constant, kinematic quantities characterizing the two systems, in this case the thermal conductivity of each, and the value of this quantity once again jumps discontinuously at the boundary of the rod. One may find it more appropriate, in the event, to use the Navier-Stokes equations to model the thermal evolution of the air, while still using Newton's equation to model the thermal evolution of the rod. In this case, not only will the system possess a kinematical discontinuity—the "matter" of the partial-differential equations changes discontinuously—but it will have a dynamical discontinuity

25. It is worth remarking that it follows from this fact—that only hyperbolic equations can have non-analytic solutions—that, contrary to what is often opined by both philosophers and physicists (see, *e.g.*, (Russell 1927)—though it is at least excusable in his case, in so far as the classification of partial-differential equations in this way, along with the clarification of their properties, was only then being settled at around the same time as he was writing, by, notably, (Hadamard 1923)), the condition that physical systems obey a "principle of causal locality" is *not* expressed by the fact that equations of motion are partial-differential equations. It is rather expressed in the fact that equations of motion tend to be hyperbolic partial-differential equations. If they are not hyperbolic—for instance, if they are parabolic, as are the classical Navier-Stokes equations—then all bets are off about the "locality" of interactions.

26. See (Geroch 1996) for a discussion of hyperbolizations.

as well—the partial-differential equations change in form across the boundary. This discontinuity also propagates in a continuous fashion.[27] Consider, in constrast, a discontinuity in the value of the temperature of the rod itself. At a given moment, say, the left half of the rod is 10^{o}C hotter than the right, and this difference manifests itself as a discontinuity in the value of the temperature in the interior of the rod as one passes through a lateral surface moving from the one side to the other. In this case, the system manifests neither a kinematical nor a dynamical discontinuity—it is continuous in both matter and form, suffering a discontinuity only in the value of one of its attributes.

We are finally in a position to offer a tentative characterization of the temporally continuous identification of a physical system as being the same, in those respects normally germane to physical identification as determined by the requirements of the investigation at hand. The interior of a four-dimensional spatiotemporal tube represents the course of dynamical evolution of a single, continuously identical physical system if

1. the partial-differential equations modeling the dynamical evolution of the physical fields one's theory ascribes to systems of that type suffer no dynamical or kinematical discontinuity within the tube
2. the values of these fields on any spacelike slice through the tube constitute a well set initial-value formulation of these equations
3. the tube is maximal in the sense that no tube containing it satisfies these conditions

I must stress that this characterization pretends to offer only necessary, not sufficient, conditions for the temporally continuous identity of physical systems. One must account for many pragmatic factors as well, if one wants to classify the system as being of a single, well recognized, continuous type over the course of its evolution, such as 'photon' or 'Hydrogen atom' or 'pendulum', for many types of systems that we think of as different in many ways obey, in some theories treating them, partial-differential equations identical in form. The three cases just listed, for instance, can all be modeled as (superpositions of) simple harmonic oscillators.

Note that what counts as "those respects of the system germane to identification" will vary from type of system to type of system, and even for the same system as it appears in different sorts of investigations, in so far as different theories will be used to model the dynamical evolution of the system in different 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[28] (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 very well push one to conclude that the resulting system is not the same as the original system. If one pumps enough energy into a Hydrogen atom, for instance, eventually 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 the atom absorbed a moment earlier.

27. One can make precise these ideas about discontinuities in the form and the matter of partial-differential equations by treating them as quasi-linear operators over appropriate function spaces.

28. It is delicate to state this a precise proposition in the context of general relativity, but it can be done. See (Wald 1984, §5.3, pp. 101–4), for example.

I take it as a fundamental assumption of physical science that "those respects of the system germane to identification" can be defined 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 modelled 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.[29] 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.[30]

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 such parts 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 at an instant' if its actual equations of motion at the instant imply conservation of all classically conserved quantities. The system will be said to be 'isolated during a spatiotemporal interval' if it is isolated at every instant of that interval. 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. Again, there is no variable in the equations of motion that labels a particular bit of energy and whose time-derivative tracks its evolution. Identifying the energy of an isolated system (in both classical physics and special relativity) as the 'same' during the period it is isolated is justified by the fact that the energy of an isolated system is conserved: the quantity of the system's energy at each instant is the same. This brute fact is supposed to justify the thought that energy can be neither created nor destroyed, and so, *a fortiori*, is the 'same' at each instant in some physically significant way. One must not take

29. This point is related to the remark of (Stein npub, §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.

30. I cannot stress enough that this discussion as it stands is far from adequate. Spatial and temporal constraints do not allow a proper treatment of the issue. I hope to publish one in the near future.

'identifying a hunk of energy associated with a system' in too strong a sense, even for periods when that system is isolated. An isolated, excited atom, for example, may at some point emit energy in the guise of a photon, but there is no sense that can be attached to the question, *which* particular bit of the atom's initial energy was emitted in the form of the photon. This is the force of a remark by (Maxwell 1877, ch. VI, §109, p. 90), which might otherwise be taken to controvert my discussion here: "We cannot identify a particular portion of energy, or trace it through its transformations. It has no individual existence, such as that which we attribute to particular portions of matter." What one *can* posit is the following proposition, required by accounts of causality that wish to invoke energetic processes to support their causal claims: that a certain unchanging quantity of energy is identifiably 'attached' in a significant way to a particular dynamical system evolving in isolation, while it so evolves, so long as the equations of motion of the system imply the principle of the conservation of energy, as do those of isolated systems in classical physics and special relativity.

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 $\mathcal{A}$ was transferred from $\mathcal{B}$, or the energy lost by $\mathcal{A}$ was transferred in part to $\mathcal{B}$, in part to $\mathcal{C}$ and in part to $\mathcal{D}$, *etc.* The same sort of analysis as worked for propagation will apply here as well. Consider the interaction of two dynamical systems. Barring Newtonian gravity, the fundamental interactions of systems in classical and relativistic theories share this feature: two dynamical systems, whether both are classical or both are relativistic, are represented as interacting with each other only when, roughly speaking, at least some bit of one is spatiotemporally continuous with some bit of the other. Except in examples such as that of the ball swung on a string at a constant rate in a circle, in which the magnitude of none of the velocities of the systems changes during the interaction, the interactions of systems in classical and relativistic theories also share this feature: during the interaction, the energetic quantities of both systems will jointly alter in a regular, predictable way. Both of these features, again, are consequences of the form of the equations of motion the systems obey during their interaction. Using the equations of motion to represent an interaction during which energy is conserved, *e.g.*, one can predict that a decrement of energy in one system will be exactly and simultaneously counter-balanced by an increment of energy in a system in some way spatiotemporally continuous with the first. This regularity and predictability of changes in energy in adjoining systems, which itself is guaranteed by the form of the equations of motion, *partially* warrants the claim one needs to support causal relations in the desired way: that the systems have *exchanged* energy—that the one has lost and the other has gained, in virtue of the fact that they were interacting, the 'same' energy. This regularity and predictability by themselves justify only the statement that energy changes in interacting systems are correlated with each other. To warrant the further claim, that energy is exchanged, a criterion is still needed for determining when energy actually 'passes' from one system to the other—when one system has *acted on* the other. Otherwise one is still in the realm of the initial value formulation of ordinary differential equations and the simple conditional propositions it entails, but not in the purportedly richer realm of causality.

Now as remarked above, one cannot tag hunks of energy as one can hunks of cheese, and so one cannot identify the energy that this system lost with the energy that that one gained in the same way one could if one were talking about cheese. The way it is actually done in scientific practice relies on the fact that, 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. If one considers the physical concatenation of the interacting systems itself to be a single system, and this combined system is isolated during the interval in which the original systems interact, then at any given instant the total amount of energy of the combined system is the same as it was at the beginning of the interaction. As a corollary, the rate at which one of the original systems gains energy during the interaction must be the additive inverse of the corresponding rate of the others, where a negative rate of gain represents loss. This brute fact provides the warrant for asserting that, in an interaction of the type considered here, 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. If energy were not conserved in interactions, then the natural conclusion would be that some of the total energy of the two systems had been either created or destroyed during the interaction. If energy were the sort of thing capable of being created or destroyed, then in any given case in which the equations of motion asserted, say, that one system lost energy and another one gained it, there would be no more reason to infer that there had been *transfer* of energy from the one to the other than to infer that some of the 'original' energy of one simply had vanished and some 'entirely new' energy had simply appeared in the other one, with no other relation between the two at all—certainly no relation rich enough to support a causal claim that is supposed to assert more than would a bare conditional statement of the form: "If energy gained here, then energy lost there." 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 exchange 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 [31]—the form of law Bob Geroch in conversation referred felicitously to as "the mathematical representation of thinghood," in so far as it encodes traditionally essential features of substance such as its identifiability over time and its permanence, which is to say the impossibility of its creation *ex nihilo* and of its destruction *ad nihilum*.

There are actually two separate formulations of the principle of the conservation of energy, the differential and the integral. The one keeps track, at individual spacetime points, of the continuous *flow* of energy into and out of the immediate neighborhood, while the other compares gross quantities of the stuff in different bounded spacetime regions in timelike relation to each other. The differential conservation law guarantees that there are no sources (or sinks) of energy in the sense of the existence of a point into (out of) which more energy flows than flows out of (into) it. The integral conservation law guarantees that, if a dynamical system gains or loses energy, then that energy loss or gain 'registers' as a non-zero total energy flux through the bounding surface of some spatiotemporal volume completely containing the system, in just the proper amount to balance the amount of energy the system gained or lost. Another way to think of the integral fomrulation is that it precludes the existence of finite or gross, not necessarily localized energy sinks and sources.

The integral form of the principle is the crucial one for sustaining the idea that energy *propagates* and is *transferred*, and not merely that some appears here and some disappears there. Given a dynamical system possibly isolated for an interval, first draw a four-dimensional spacetime tube closely around the system throughout the interval. Applying the integral form of the law over this tube allows one to keep track of the system's energy for the whole interval, which includes determining whether the system gained or lost any energy during the interval and, if so, the specifics of the gain or loss. In so far as one has succeeded in reducing causal propositions in the first place to ones about energy propagation and transfer, this device further allows one to affirm propositions such as "The causal process evolved via *this* continuous chain of events, and not via any of those," rather than merely saying "The sum of events in this spacetime region here determined, or was

31. I emphasize that the satisfaction of such conservation principles is *only* a necessary condition for being able to define propagation and exchange of conserved quantities as needed for transfer accounts of causality. Howard Stein has argued convincingly in a private communication to me that the propagation of energy and of other classically conserved quantities do not always track *prima facie* facie causal relations in, *inter alia*, classical electromagnetic theorythe propagation of causally relevant information. The propagation of optical information in a diffracted optical field, for instance, does not necessarily follow the flowlines of the flux of any classically conserved quantity in classical electromagnetism.

causally responsible in some way for, the event or sum of events in that region there."

In classical physics and in special relativity, these two formulations of the conservation principle are essentially equivalent to each other—both hold in general and each implies the 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.

2.5 Energy in General Relativity

General relativity does not naturally support any sort of 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.[32] 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. If one accepts the argument I made in §2.4 above, that integral conservation laws are a *sine qua non* of defining propagation and transfer of energy (or of any classically conserved quantity), in so far as one wants to have these ideas support rich causal claims, it follows immediately that one cannot formulate transfer accounts of causality in a 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}, which can be thought of as a linear map from ordered pairs of vectors on spacetime to real numbers; second, that generic spacetimes in general relativity have no preferred class of "frames of reference", 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. The lack of these two structures collude to hinder the formulation of an integral conservation law.

In special relativity, one can use is also a tensor, not a scalar, and yet one still can formulate integral conservation laws perfectly well there. By applying the ambient stress-energy tensor in turn to each of the canonical timelike killing fields on Minkowski spacetime, one constructs a canonical set of scalar fields naturally thought of as energy densities, and formulates a distinct integral conservation law for each scalar field in the set. In order to construct a scalar from a stress-energy tensor in general relativity, two vectors are needed (recall that the stress-energy tensor is a linear map from ordered pairs of vectors to real numbers—to scalars). The vector tangent to the worldline of an observer provides an obvious and natural candidate, and indeed it is a fundamental fact about general relativity that the ambient total energy density experienced by any given observer at a given spacetime point is precisely the real number one gets by applying the stress-energy tensor at that point to the ordered pair each component of which is the vector tangent to the observer's worldline at that point. So far, this is precisely the same procedure followed in special relativity to formulate the integral law of energy conservation. To formulate an integral law in general relativity, therefore,

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

it would seem that all that has to be done is to pick an appropriate family of observers, *viz.* a family of timelike curves, so that by applying the stress-energy tensor to their respective tangent vectors at every point in the region one will get a scalar field representing the total energy-density of all matter in the region as experienced by those observers, which can be used to formulate the integral law. In special relativity, families of timelike curves representing inertial observers all at rest with respect to each other play this role, and because such worldlines in Minkowski space have a few extraordinarily nice technical properties—summed up in the proposition that they are simultaneously geodesics and the integral curves of a Killing field—it turns out that the scalar field one gets in this way does yield a perfectly good integral conservation law. One oddity of the situation is that one will get not one but an uncountable cardinality of different integral conservation laws, one for each 'preferred time-frame' defined by a family of co-moving inertial observers. There is nothing inconsistent about this—each family of observers will experience energy, *et al.*, as being conserved in their own time-frame, and will be able to predict that all other families of co-moving inertial observers will have the same experience, though no two families will agree on the values of the ambient energy density and flux. In fact, this is true in classical physics as well: one will get different integral expressions for the conservation principle for each different Galilean inertial coordinate system as well.

The crucial difference between special relativity and general relativity, between, that is, Minkowski spacetime and the generic general relativistic spacetime, is that, 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. In fact, this last structure by itself suffices for formulating an integral conservation law. To formulate an integral conservation law in a given region of a general relativistic spacetime that has a timelike Killing field, one would pick a family of timelike geodesics filling the region ('inertial observers'—these can always be found); then, to construct the appropriate scalar field, one would apply the stress-energy tensor at each point in the region to the pair of vectors one component of which was the tangent vector to the geodesic from the fixed family passing through the point and the other component of which was the Killing field vector at that point (because the stress-energy tensor is symmetric, it would not matter how one ordered the two vectors); finally one would use this scalar field to formulate the integral conservation law. The properties of the timelike Killing field ensure that one will be able to formulate an integral conservation law for the resulting scalar, in analogy to those of special relativity. The scalar field that results can be thought of in a certain sense as the total energy density at a point, with two important caveats. 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. Second, it does not include any contribution due to 'gravitational energy', since this is not localizable in general relativity. [33]

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 time passed. 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 yield integral conservation laws—they provide

33. If the Killing field is spacelike, then one may get conservation laws for quantities analogous to linear and angular momenta as they appear in special relativity.

a physically significant 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.[34]

Even if general relativity does not allow the formulation of such conservation laws, and so does not allow the defining of a quantity like energy as it appears in special relativity, one may still wonder whether energetic quantities useful to transfer accounts can be defined in other ways. They cannot. The fundamental structure of general relativity by itself does not provide the appropriate setting for any localizable energetic quantity to be rigorously defined in any way analagous to how such quantities are defined in either classical physics or in special relativity. They just are not fundamental components of the theory as they is in classical physics and special relativity. I will not enter here into the technical details of this result; I will only remark that the heart of the matter lies in 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.[35]

On the face of it, this is an extremely 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.[36] 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

34. Another class of special spacetimes, the so-called *asymptotically flat* ones, also admit two precisely defined energy-like quantities. One is most naturally interpreted as the total energy contained in the spacetime 'at a single instant of time', *i.e.* in a single spacelike hypersurface (see (Arnowitt, Deser, and Misner 1962)); the other represents the total amount of energy 'radiated off to infinity' at any given time (see (Bondi, van der Burg, and Metzner 1962) and (Sachs 1962)). These are both global quantities, akin to the total energy of a dynamical system in classical mechanics, but with one very strange feature: they have no local analogues—there is no scalar or vector field in such spacetimes that one integrates up to get this 'total energy' (see (Curiel 1996)). Consequently, while these quantities are of great interest for purposes of calculation, they do little good for the advocates of transfer accounts of causality.

35. The precise statement is that one cannot define a 'stress-energy tensor for the gravitational field': the only two-index covariant, symmetric 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.

Similar results about the indefinability of gravitational energy hold in Newtonian gravity, though the situation is somewhat better there in that, in special situations, one can define an energy density for the gravitational field, which one can never do in general relativity. 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.

36. 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.

sorts of approximations 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 actually 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, no matter how good and 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 conservation laws for its definition, in one sense disappears in general relativity, it does not do so completely. I think 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 time-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 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 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 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, 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 scalar energetic quantities, which are always well-defined. 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, as I have tried to emphasize. 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. I think it will be helpful in making this point more clear to take a very brief look at some historical material. To my great surprise, the hero among the early proponents of general relativity (of those I have read with some care—I make no claim to have perused a large fraction of the significant historical literature in this area), at least with regard to having made a beginning of formulating a coherent and I think largely proper conception of the role of the stress-energy tensor and conservation principles in the theory, is Eddington. He has of the time perhaps the most sophisticated treatment of the role the stress-energy tensor of matter plays in general relativity, of the way it is introduced in the theory, and of the proper view to have of classical conservation laws and why they are not fundamental to the theory. [37] I take the liberty of quoting him at length:

> . . . [W]e have. . . spoken of [the differential covariant conservation law] as the law of conservation of energy and momentum, because, although it is not formally a law of conservation [since it has no integral formulation], it expresses exactly the phenomena which classical mechanics attributes to conservation. . . . ¶ As soon as the principle of conservation of energy was grasped, the physicist practically made it his definition of energy, so that energy was that something which obeyed the law of conservation. He followed the practice of the pure mathematician, defining energy by the properties he wished it to have, instead of describing how he measured it. This procedure has turned out to be rather unlucky in the light of the new developments. . . . We find that [the stress-energy tensor] is not in all cases formally conserved, but it obeys the law that its [covariant] divergence vanishes; and from our new point of view this is a simpler and more significant property than strict conservation.(Eddington 1923, §59, pp. 135–6)

I want to focus on the conclusion of the quotation, allowing the rest to speak (rather more eloquently than I could) for itself: Eddington claims that *general relativity itself* teaches us that the differential covariant conservation law is the principle it is proper to expect to govern the behavior of whatever energetic quantities there may be in the world, and not a classical conservation law that could be transformed into an integral conservation law. He does not make an explicit argument explaining why this equation is the one that best captures, in the mathematical language with which one represents the relativistic world, the physicists' practice of making energy and momentum measurements and finding that to an extraordinarily high degree of accuracy on the surface of the Earth certain classical conservation principles hold, but it is easy enough to sketch out what I believe he had in mind.

What general relativity, "our new point of view," demands to be taken into account is that for the purpose of quantitative calculation we represent all such measurements of energy and momentum in particular coordinate systems, none privileged over the others, and that all such measurements can be approximated as occurring over an infinitesimal region of spacetime, on the supposition that they are not occurring in regions of extremely high curvature. Consequently a *covariant* law must be formulated that expresses the fact that, infinitesimally, these conservation principles are observed to hold when expressed in any coordinate system: all observers, no matter their state of motion will agree that these principles, properly formulated, do hold. The differential covariant conservation law precisely encodes all this information. This is why it deserves the honorific 'conservation law' even though it yields no integral equation in general: in a certain sense, the covariant law becomes a classical differential conservation law in the limit of the infinitesimally small. Thence the deep continuity between the general relativistic energetic quantities and classical conceptions of energy—they represent very nearly the same class of physical processes and operations, and are used to

37. *Cf.* (Eddington 1923, esp. §§53-4 and §59, pp. 116–22 and pp. 134–7). (Russell 1927, ch. 9, pp. 84–95), one of the first philosophers to examine relativity theory with a high degree of mathematical sophistication, also took especial delight in Eddington's analysis. We are in no position today to condescend to Eddington, either—part of the original impetus behind this dissertation was a fruitless search for a philosophically illuminating contemporary examination of energy and conservation laws in general relativity.

model the same experiments and to make predictions about them, predictions that in many cases are well-nigh indistinguishable among the various theories; moreover, the mathematical structure of the former can be shown to 'contain' the mathematical structure of the latter, in the sense that the structure of the latter falls out of that of the former under certain natural approximations, using certain natural manipulations. The vanishing of the covariant divergence of the stress-energy tensor is "simpler and more significant" in general relativity than a classical conservation law would be for the simple reason that it is a well-posed statement utilizing only structure intrinsic to the theory, *viz.* the stress-energy tensor and the affine connection of spacetime.

In the original derivations of the field equation that bears his name, Einstein repeatedly relied on the principle of energy conservation in arguments that motivated and even 'proved' many essential propositions.[38] Most of the reasons he gives for relying on this principle do not look so strong once one considers its status in the complete theory. In this respect, energy conservation is analogous to Mach's principle and to the principle of equivalence—a guiding intuitive principle that helped inspire the construction of the theory, but whose classical formulation does not seem quite to hold in the final theory itself. This fact by itself, however, does not invalidate his arguments: the deep practical, empirical and theoretical continuities among energetic quantities in the various theories I have emphasized show why Einstein's arguments are so good, why they are so successful, for in an important way we are still talking about the same underlying structural features of the physical world. Only now the concepts we use to investigate and understand them have evolved; they have not discontinuously metamorphosed. Correspondingly, the physical theory we use to represent and model these underlying structural features of the world has also changed. The sorts of calculations and derivations one can employ a stress-energy tensor in have changed, for example; the mathematical framework within which one specifies a stress-energy tensor has changed; judgments about the propriety of certain sorts of approximations—what to include in calculations, what to ignore—have changed. And so on. This new physical theory, though, both on its own and in its intricate connections to past theories and their better understood concomitant concepts, provides the key to understanding the new concepts it has introduced.[39]

General relativity demands revision of the classical conceptions of 'energy' and 'conservation'. I believe that we have not yet fully come to grips with the revisions, and perhaps abandonments, it urges on us.

2.6 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. 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."

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

38. See for instance (Einstein 1916, *passim*) and (Einstein 1984, *passim*).

39. The too short discussion of this paragraph was inspired in large part by the discussions of similar matters in (Stein 1989) and (Stein npub).

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.[40] 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 surely 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 equations in a mathematically sensible way for point-particle sources, as one can, say, for the Maxwell equations.[41] 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 §2.4 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. For any dynamical system, the equations of motion by themselves do not contain an 'identity variable'—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 conceit is extremely artificial 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*.[42]

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 in the first place already with a set of classical notions and questions to address, and did not rather ask general relativity what the important notions and questions ought to be in its new framework.

If one renounces transfer accounts of causality, as I see it there remain only two general sorts of

40. One could perhaps try to use entropy and entropy-flux to define propagation and causal continuity, since entropy can be rigorously defined and treated, at the macroscopic level at least, in general relativity (see, for instance, (Tolman 1934, §§119–20)). So far as I know, whether one can give a rigorous treatment of entropy at the atomic and sub-atomic level in general relativity is not known, and presumably must await advances in quantum field theory on curved spacetimes and quantum theories of gravity. See (Wald 1999) for a recent survey of this problem.

41. 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. Recently, (Colombeau 1992) has developed a theory of so-called new generalized functions that can be viewed as a non-linear generalization of distributions. Although in this new framework one can make sense of a much wider class of metrics than one could in the past, one still cannot rigorously construct a metric representing a point-particle. See (Vickers and Wilson 1998) for a recent survey of applications of Colombeau's theory to general relativity.

42. 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, but that is a sermon for another time.

accounts of causality that could be grounded in physical theory. 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.[43] Whatever sort of account one will give looks to come perilously close merely to saying that one thing *follows* upon another.

I have been concerned with accounts of causality that aspire to be fundamental, to reflect the actual, basic structure of the physical world as best we know it. I would not desire to preclude from the philosophical and scientific armory all notions of causality that depend on ideas of propagation and classically conserved quantities, much less to banish them from everyday discourse about every-day objects, but I think my argument does demand from any philosopher who wishes to invoke such a notion in his arguments an accounting 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 inital value formulation of mathematical physics that is supposed to arise naturally from an analysis of physical theory ought to be treated with suspicion.

Philosophers involved in projects ranging from arguments for the physical basis of the direction of time (*e.g.* (Reichenbach 1956)), to the origin of linguistic reference (*e.g.* (Putnam 1975)), to analysis of perception (*e.g.* (Russell 1927)), to accounts of physical measurement (*e.g.* (Hacking 1983)) and defenses of realism (*e.g.* (Boyd 1991), (Hacking 1983) and (Shimony 1993)), no longer get access to such accounts of causality for free. 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. 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. Though this lesson about the circumscriptions on uses of certain causal notions perhaps could have been drawn from quantum mechanics alone,[44] such an argument would not have been nearly so clean and straightforward as that from general relativity, given the hotbed of dispute surrounding any interpretational theses forwarded about quantum phenomena.[45]

43. 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.

44. (van Fraassen 1989) attempts a related project.

45. I thank David Malament and Howard Stein for impeccable advising on and penetrating criticism of my doctoral thesis, from the third chapter 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 topics physical and mathematical.

Chapter 3

The Analysis of Singular Spacetimes

> The mind of man, by nature a monist, cannot accept *two* nothings; he knows there has been *one* nothing, his biological inexistence in the infinite past, for his memory is utterly blank, and *that* nothingness, being, as it were, past, is not too hard to endure. But a second nothingness—which perhaps might not be so hard to bear either—is logically unacceptable.
>
> V. Nabokov, *Ada*

ABSTRACT

Much controversy surrounds the question of what ought to be the proper definition of 'singularity' in general relativity, and the question of whether the prediction of such entities leads to a crisis for the theory. I argue that a definition in terms of curve incompleteness is adequate, and in particular that the idea that singularities correspond to 'missing points' has insurmountable problems. I conclude that singularities *per se* pose no serious problem for the theory, but their analysis does bring into focus several problems of interpretation at the foundation of the theory often ignored in the philosophical literature.

3.1 Introduction

I suspect that, for many, talk of a singularity in the context of general relativity conjures up the image of something like a rent in the fabric of spacetime. [1] Perhaps unbounded curvature from the self-gravitational collapse of a massive body tore the fabric, or perhaps the cloth was simply ill-woven from the start, but in any case the idea of a flaw in the fabric of spacetime naturally accompanies the word 'singularity'. This metaphor, evocative as it may be, is perhaps misleading: a web of cloth exists in space and time, and one naturally would rely (implicitly, at least) upon this fact were one to define what one meant in saying the cloth were rent. For instance, one might define a cloth to have a hole if one could thread a string through the cloth, tie the ends of the string together and

1. By 'spacetime', I will always mean a smooth, 4-dimensional, connected, paracompact manifold endowed with a fixed, smooth metric of Lorentz signature.

have the string touching disjoint components of the edge of the cloth.[2] When thinking of spacetime, though, one does not have the luxury of imagining it embedded in any physically meaningful way in a larger space with respect to which one can try to define what one means by saying there is a hole.

One can think about holes in cloth in (at least) two ways: no cloth has been removed, but parts of the cloth have simply been separated from each other (torn) for a length; a bit of the cloth has actually been excised and removed. In the former case, all the points (bits) of the cloth are still there but the topology has changed, whereas in the latter case there are actually points (bits) that once were part of the cloth now missing from it. I wager that people usually conceive of singular spacetimes in a way analogous to the latter idea when they think in a vague, intuitive way: there are points missing from spacetime. For example, in thinking about the self-gravitational collapse of a massive body, one might imagine the "point" in which all the matter in the body becomes eventually concentrated. In a normal collapse, the curvature of spacetime will in some sense become unboundedly large as one approches this "point", so, again loosely, one will not be able to define the spacetime metric at that "point"—and now one sees why I have been enclosing 'point' in scare-quotes, for spacetime comprises solely points of a manifold with a pseudo-Riemannian metric of Lorentz signature defined thereat. The "point" to which all the matter collapsed is missing from the spacetime.[3]

On a manifold endowed with a positive-definite Riemannian metric, one can give a precise characterization, according quite well with our intuitions, of what it is for there to be missing points. Turn the manifold into a pointwise-metric space (*i.e.*, one possessing a true distance-function on the space of ordered pairs of its points) via the usual construction: define the distance between any two points to be the infimum of the lengths, with respect to the Riemannian metric, of all smooth curves connecting them. The manifold has no missing points if and only if it is Cauchy complete with respect to the constructed pointwise-metric. Intuitively speaking, if a sequence of points begins to accumulate, there ought to be a place at which they actually do accumulate. If there are missing points, one may take the Cauchy completion of the manifold its guise as a pointwise-metric space to fill in the gaps, as it were.

On a manifold with a pseudo-Riemannian metric of Lorentz signature, such as a spacetime in general relativity, there is no natural way to construct a true pointwise-metric measuring the distance between points of the manifold, so one cannot employ this technique to test whether a spacetime has missing points. By the Hopf-Rinow-de Rham theorem, the manifold in the Riemannian case is Cauchy complete with respect to the constructed pointwise-metric if and only if it is geodesically complete with respect to the Riemannian metric.[4] This naturally suggests that we define a spacetime to have missing points if and only if it is geodesically incomplete with respect to the spacetime pseudo-Riemannian metric. Now one faces a severe problem, which lies at the heart of the difficulty in giving a precise and intuitively satisfying definition of singular structure as a point missing from spacetime: there is no natural way to take a Cauchy-like completion of a spacetime manifold having incomplete geodesics, in order to give substance to the idea that there really "are" points that in some sense ought to have been included in the spacetime in the first place.[5] In the Riemannian

2. More rigorously, this amounts to saying a 2-dimensional compact topological manifold has a hole if and only if it has a boundary not homeomorphic to $\mathbb{S}^1$. Thus the torus does not have a hole, since it has no boundary; neither the spherical shell with a small cap excised nor the Möbius strip has a hole, since the boundary of each is homeomorphic to $\mathbb{S}^1$ (the spherical shell with a cap excised is homeomorphic to the planar disk); the finite cylinder has a hole, since its boundary is homeomorphic to the disjoint union of $\mathbb{S}^1$ with itself (the finite cylinder is homeomorphic to a planar annulus). I thank James Geddes for illuminating discussion on this question.

3. More precisely, a point of a spacetime manifold ought to be considered a point of spacetime itself if and only if, on the bundle of pseudo-Riemannian metrics over the manifold, the cross-section representing the spacetime's metric is well defined in the fiber over the point in question.

4. See (Spivak 1979a, ch. 9) for a precise statement and proof of the theorem, and for more information on the constructed pointwise-metric and the Cauchy completion of a manifold endowed with a Riemannian metric.

5. The scare-quotes now come from the fact that it is not clear in the slightest what sense may accrue to the

case, roughly speaking, one constructs the missing points by taking equivalence classes of incomplete curves that get arbitrarily close to one another as measured by the constructed pointwise-metric. In the pseudo-Riemannian case there is no natural way to measure how close two curves come to one another, so, *a fortiori*, there is no natural way to define missing points as the equivalence classes of incomplete curves that come arbitrarily close to each other.[6]

The usual tack taken at this point in the physics literature is simply to bracket the question of missing points and define a spacetime to be singular if and only if it contains incomplete, inextendible curves of a certain specified type, and the spacetime manifold itself satisfies a few collateral conditions. The commonly accepted schema for fixing a rigorous definition of a singular spacetime, then, is:

> A spacetime ($\mathcal{M}$, g_{ab}) satisfying ____ is *singular* if and only if there exists a curve γ incomplete in the sense that ____.[7]

Such a conception of singular structure actually has a lot to say for itself, as capturing the idea that singular structure is somehow physically *outré*, even if one is not able to hook it up cleanly to an idea of missing points. As (Hawking and Ellis 1973, p. 258) put it,

> Timelike geodesic incompleteness has an immediate physical significance in that it presents the possibility that there could be freely moving observers or particles whose histories did not exist after (or before) a finite interval of proper time. This would appear to be an even more objectionable feature than infinite curvature and so it seems appropriate to regard such a space as singular.

The current paper has several concrete aims: to investigate particular ways that have been proposed to fill in the blanks of the schematic definition with an eye to determining whether they capture the spirit of the idea that an incomplete curve corresponds to singular structure; to examine the relation between curvature pathology and singular structure so defined; to argue that the idea of missing points ought not be central in thought about singular structure; and to argue that the reasons most often given for condemning singular structure as unphysical do not withstand scrutiny. It also has one overarching, more nebulous aim: to try to give a sense of the philosophical riches still waiting to be mined from thorough investigation of the foundations of general relativity—which is to say, a sense of how little of this theory we even now comprehend, and how much we stand in need of that comprehension if we wish to understand the world.

---

attachment of an existential quantifier to "points that are possibly spacetime points but in the event are not". We touch on this issue in §3.5. I will dispense with them from hereon, the point having been made.

6. Cauchy completeness of a Riemannian manifold with respect to the constructed distance-function happens also to be equivalent to the following condition: every bounded (with respect to the constructed distance-function) subset of the Riemannian manifold is relatively compact (Kobayashi and Nomizu 1963, p. 172). So far as I know, no one has tried to parlay this equivalence into a definition of 'missing points' in the pseudo-Riemannian case. Again, since there is no distance-function in the pseudo-Riemannian case, there is no natural candidate for what ought to count as 'bounded subsets' of the manifold. A first stab might be: points are missing from the manifold if and only if, for some p and q in the manifold, $J^+(p) \cap J^-(q)$ is not relatively compact, where $J^+(p)$ ($J^-(p)$) is the causal future (past) of the point p. If this could be made to work, it would have the great virtue of "localizing" the missing point—when asked, "where is the point missing from?", one could point to the salient $J^+(p) \cap J^-(q)$, and say, "from that region". The obvious problem with this candidate is that it fails to categorize Schwarzschild spacetime as having a missing point, whereas one might have thought that Schwarzschild was the paradigm of a spacetime with a missing point, *viz.*, the 'point' into which all the matter from a body undergoing self-gravitational collapse squeezes itself. In fact, the first stab fails to categorize any globally hyperbolic spacetime as having missing points, since all globally hyperbolic spacetimes by definition satisfy the proposed condition (Wald 1984, p. 209). Even though a fairly obvious first candidate fails, it still might be interesting to explore whether one could propose a reasonable analogue of 'bounded subset' for the pseudo-Riemannian case and use this to define missing points. Of course, because of the known examples of compact, geodesically incomplete spacetimes (*cf.* (Misner 1963)), one should expect that any such characterization based on the relative compactness of 'bounded' subsets would be bound to differ in what it counts as singular from the traditional characterization in terms of geodesic completeness.

7. See, for example, (Hawking and Ellis 1973, pp. 256–61), (Wald 1984, pp. 212–6), (Clarke 1993, p. 10), and (Joshi 1993, pp. 161–2).

3.2 Curve Incompleteness

The path-breaking work of the mid-1960's demonstrating the existence of singular structure in generic solutions to the Einstein field equation invoked timelike or null geodesic incompleteness as a sufficient condition for classifying a spacetime as singular, in so far as timelike and null geodesics represent possible world-lines of particles and observers and, *prima facie*, it appears physically suspect for an observer or a particle to be allowed to pop in or out of existence right in the middle of spacetime, so to speak.[8] There was, however, no consensus on what ought to count as a necessary condition. In particular, workers at the time were unclear on the role played by curvature pathology in singular structure. For example, Hawking, in his very early work, distinguished between the mere incompleteness of the spacetime manifold (as characterized by the existence of incomplete geodesics) and what he referred to as a "physical singularity" apparently meaning a spacetime region wherein, in one of a number of technical senses, the magnitude of the curvature grows without bound: "Penrose has shown that either a physical singularity must occur or space-time is incomplete if there is a closed trapped surface. . . ."[9] Context makes clear that Hawking relates the existence of a trapped surface with the existence of pathology in the behavior of the curvature. It is worth remarking that, based on a careful reading of (Penrose 1965), to which Hawking here refers, it is not at all clear that Penrose himself would have endorsed this statement of his result. In an apparent lightning-fast sequence of changes of mind that strikingly illustrates the uncertain and fluid nature of the idea of singular structure in the field at the time, in April, 1966, Hawking proposed using the prediction of singular structure (which, note, meant only the existence of incomplete timelike or null geodesics) as a possible test of the validity of general relativity,[10] whereas by February of the very next year he concludes that the singularity theorems proved up to that point "probably" indicate not that singular structure actually occurs in the universe but rather that general relativity breaks down in the strong field regime![11]

The field was ripe for a little sober reflection, happily provided by (Geroch 1968b), who gave the first extended discussion of the difficulty of framing a satisfactory definition of a singular spacetime.[12] Geroch's discussion begins in earnest with a Galilean dialogue, a form, as Earman notes, nicely suited for displaying the unsettled state of the topic.[13] After concluding that one can use neither the physical components of the Riemann curvature tensor nor any of the scalar-curvature invariants to define precisely what one means by, and construct necessary and sufficient conditions for, saying a spacetime contains regions wherein the curvature grows without bound in a physically accessible manner,[14] the discussants in the dialogue settle on simple geodesic incompleteness as the criterion

8. *Cf.* (Penrose 1965), (Hawking 1965), (Geroch 1966), (Hawking 1966a), (Hawking 1966b), (Hawking 1966d), and (Hawking 1967).

9. (Hawking 1965, p. 689).

10. (Hawking 1966b, p. 511).

11. (Hawking 1967, p. 189). The dates referred to in the text (as opposed to those of the citations proper) are those on which the journal recieved the papers for review, as indicated in the published versions. In the event, the second viewpoint seems to represent Hawking's settled opinion on the matter—*cf.* (Hawking and Ellis 1973, §10.2) and (Hawking and Penrose 1996, p. 20).

I stress that I do not take this oscillation of Hawking's from position to position as an act worthy of derogation, far from it. Rather, he seems to me to have been engaged in the practice of a good scientist: entertaining all the decent possibilities presenting themselves so as to test them by use in his investigations, in order to see which bear fruit and which do not.

12. (Kundt 1963), (Misner 1963) and (Hawking 1967), among others, had already broached in a cursory manner several of the topics Geroch discussed in this paper.

13. (Earman 1995, p. 27). Only Sagredo and Salviati discuss the issue, with no word from Simplicio. I can speculate only that the issue was too difficult for Simplicio's limited capacities.

14. The physical components of the Riemann tensor are its components relative to any pseudo-orthonormal tetrad; roughly speaking, a scalar-curvature invariant is a scalar 'function' of the metric, the Riemann tensor and its covariant derivatives that is preserved under diffeomorphisms of the spacetime. See (Ehlers and Kundt 1962) for details. I will discuss in §3.3 why none of these suffice for constructing necessary conditions.

for singular structure, conceding that the definition is perhaps overly inclusive, but better to brand 10 innocents than to allow one guilty man unmarked. The possible innocents include spacetimes all of whose timelike and null geodesics are complete, but possess incomplete spacelike geodesics (null and timelike complete and spacelike incomplete, for short). Spacelike incompleteness (in the absence of the other two types of incompleteness) sets off no serious alarms, or so thought commonly goes, for an incomplete spacelike geodesic seems to represent structure of the spacetime not physically accessible to any observer in a direct way. [15] Moreover, not only does geodesic incompleteness lock up a few possible innocents but, as Geroch proceeds to show, it almost certainly fails to nab a few clever guilty parties, for a spacetime can be geodesically complete and yet possess an incomplete timelike curve of bounded total acceleration—that is to say, an inextendible curve traversable by a rocket expending only a finite amount of fuel, along which an observer could experience only a finite amount of proper time.

Because of these problems, null and timelike geodesic incompleteness continued to be used as a sufficient condition for declaring a spacetime singular, but was (and still is) considered inadequate as a definition. [16] To analyze the structure of non-geodetic curves in the search for a necessary condition, we require a method for characterizing their completeness. The following appears tempting at first glance: an inextendible curve is incomplete just in case it has finite proper length. Even if one puts aside for the moment the fact that every null curve has zero proper length, one still faces the following problem with any approach based on proper time: every spacetime, including Minkowski space, has inextendible timelike curves of finite total proper length, *viz.*, those of unbounded total acceleration that go zooming off to infinity, so to speak, asymptotically approaching the speed of light. Surely one does not want to classify Minkowski space as incomplete, and anyhow, if an observer is able to reach infinity, as it were, even in a finite amount of time, the prevailing sentiment in the physics community at large seems to be that such structure ought not qualify as singular. [17] One wants a method of winnowing such acceptably finite curves from unacceptable ones.

(Schmidt 1971) appears to have been the first to propose using so-called generalized affine parameters to define the completeness of general curves. Let $\mathcal{M}$ be an n-dimensional manifold with an affine connection, $\gamma(t)$ a curve through $p = \gamma(0)$, and $\{\overset{i}{\xi}{}^a(0)\}_{i=1\dots n}$ a basis for the tangent space at p. One can now write $\gamma^a(0)$, the vector tangent to γ at p, as a linear combination of the chosen basis with coefficients $\gamma_i(0)$:

$$\gamma^a(0) = \sum_{i=1}^{n} \gamma_i(0)\, \overset{i}{\xi}{}^a(0).$$

If one parallel-transports the chosen basis along $\gamma(t)$, one gets a similar expression at every point on $\gamma(t)$. The *generalized affine parameter* $\theta(t)$ of $\gamma(t)$ associated with this basis is defined by:

$$\theta(t) \equiv \int_0^t \left(\sum_{i=1}^{n} (\gamma_i(t'))^2 \right)^{\frac{1}{2}} dt'.$$

15. See, *e.g.*, (Synge 1960, ch. I, §14, pp. 24–6) for a discussion of the physical content the measurement of spacelike intervals in general relativity may possess. In a similar vein, one may also consult (Geroch 1981). It would be of some interest to investigate whether one could parlay discussions such as these two into arguments for the "physicality" of incomplete spacelike geodesics.

16. Hawking defines a singular spacetime as one which is timelike or null geodesically incomplete in (Hawking and Penrose 1996, p. 15), but I believe this is not meant as a serious attempt at a strict definition, merely an easy criterion to work with in light of the fact that all known singularity theorems prove the existence of incomplete timelike or null geodesics. It would be of interest, again, to investigate the question whether there exists a set of conditions that "physically reasonable" spacetimes ought to satisfy, having as a consequence the existence of an incomplete spacelike geodesic.

17. I think this sentiment represents a hypocrisy on the part of the community, as I will discuss briefly just below and in more detail in §3.6.

In effect, one treats the parallel-transported basis of vectors as though they were the orthonormal basis of a Riemannian metric and then defines the 'length' of $\gamma(t)$ accordingly. The generalized affine parameter of a curve does not depend on the basis chosen in one crucial respect: whether or not the generalized affine parameter of the curve increases without bound. Furthermore, any curve of unbounded proper length automatically has an unbounded generalized affine parameter, but not vice-versa—any inextendible timelike curve of unbounded total acceleration and finite total proper time in Minkowski space, for example, has an unbounded generalized affine parameter. A spacetime in which every inextendible curve has an unbounded generalized affine parameter will be referred to as *b-complete.*[18] This sort of completeness promises to distinguish precisely what wanted distinguishing, and works just as well for null as for timelike or spacelike curves. Thus, one has what (Earman 1995, p. 36) refers to as the "semioffical view": a spacetime is said to be singular if and only if it is b-incomplete.[19] This definition is more general than geodesic completeness, in that it implies, but is not implied by, the latter, as Geroch's example demonstrates.

It is difficult to think of a more comprehensive criterion of completeness than *b*-completeness, and I suspect its popularity arises from this fact,[20] but that it sits comfortably with some of the intuitions that drove the search for a definition of singular structure in the first place is not so clear on reflection. That it counts some timelike curves of total finite proper time as complete (*viz.*, some of those of unbounded total acceleration) is perhaps its most unsettling feature, if one of the intuitions driving the search for a definition of singular structure is the impropriety of having observers or particles who can exist for only a finite period of time. It is also a cumbersome and technically awkward criterion to deploy in practice. In fact, perhaps the most damning fact about *b*-completeness is that, so far as I know, it is never used in the statement or demonstration of any results of true physical interest. All the singularity theorems, for instance, demonstrate only the existence of null or timelike geodesics, and are formulated only in those terms. For the moment, I will waive these qualms and accept *b*-incompleteness as the definition of singular structure—when I refer to 'incomplete curves', unless I explicitly state otherwise I will mean *b*-incomplete, inextendible curves. I will return to some of these questions below in §3.6.

3.3 Explosive Curvature Growth along Incomplete Curves

While curve incompleteness seems to capture one aspect of the intuitive picture of singular structure, it completely ignores a different aspect, curvature pathology. One may measure the growth and diminution of the magnitude of spacetime curvature in various ways, but it turns out that the unbounded growth of curvature according to any of these measures is neither necessary nor sufficient for the existence of incomplete, inextendible curves. To get an idea of the independence of the existence of incomplete curves from the presence of curvature pathology, consider the striking ease with which examples of a spacetime with everywhere vanishing Riemann tensor and incomplete geodesics can be constructed: excise from 2-dimensional Minkowski space a closed set in the shape of

18. '*b*' for 'bundle': with this construction one tacitly defines a natural (basis-dependent) Riemannian metric on the bundle of frames of the spacetime manifold to define curve completeness.

19. Strictly speaking, this is not the standardly accepted definition, since I have not mentioned anything about the maximality of the spacetime in question, whether, that is, it can be embedded in (thought of as merely a part of) a larger spacetime in such a way as to make previously incomplete, inextendible curves extendible. I will take up this issue in §3.6.

20. (Schmidt 1971) claims that *b*-completeness is the "natural" generalization for pseudo-Riemannian metrics of completeness with respect to a Riemannian metric, in so far as it is equivalent to completeness with respect to a Riemannian metric when used on a Riemannian manifold: geodesic completeness with respect to a Riemannian metric is logically equivalent to *b*-completeness as defined by its affine connection. I do not know what 'natural' signifies in this context, in so far as the criterion in the Riemannian case may be formulated without the use of components of geometric objects in an arbitrary coordinate system, but Schmidt's method cannot be formulated without it.

an echidna. This example may strike one as cheating, since one has only to restore the excised set to restore geodesic completeness (or, in fancier terms, one has only to isometrically embed our mutilated spacetime by the natural inclusion map back into Minkowski spacetime to restore completeness). So a slightly more sophisticated example: for some $0 < \phi_0 < \frac{\pi}{2}$, excise from Minkowski space, represented in polar coordinates, the wedge consisting of all points with azimuthal coordinate $0 < \phi < \phi_0$; identify the corresponding points on the hyperplanes $\phi = 0$ and $\phi = \phi_0$. By a suitable redefinition of the coordinate neighborhoods of the points on $\phi = 0$, the resulting space can be given the manifold structure of $\mathbb{R}^4$, and the Minkowski metric can be smoothly extended to the points at $\phi = 0$, $r > 0$. It cannot be smoothly extended to the points $r = 0$, however, and so these points must be excised from the spacetime. The Riemann tensor of this spacetime vanishes everywhere, but any geodesic that previously passed through the line $r = 0$ will now be incomplete; there is, moreover, no other spacetime into which this spacetime can be embedded and in which the metric can be smoothly extended.[21] This sort of structure is known as a 'conical singularity', since the singular structure has many of the same characteristics as that accruing to the two-dimensional real plane with a wedge removed and the edges pasted together, so as to form a cone.

Perhaps this example will also strike the reader as too artificial, too contrived, to have any physical relevance.[22] I believe that on a matter such as the global topological structure of spacetime, about which we have so very little hard data and so little prospect of gathering any for the foreseeable future, one should be wary of ignoring certain sorts of examples on the ground that they appear 'artificial'. This judgment has its roots in the schooling our intuitions have received in our contemplation of well worked out examples of physical theories, which by and large tend to include mathematical structures that strike us as 'simple' and 'natural'. This ought not escape our notice: most such examples of physical theories are demonstrably false (Newtonian mechanics and classical Maxwell theory) or have at the moment insuperable problems of interpretation (quantum mechanics) or experimental accessibility (general relativity). We should beware of relying too much on intutions trained in such schools—especially when one also recalls how much of our contemplation of those theories involves models of systems with physically unrealistic perfect symmetries and vaguely jusified approximations, simplifications and idealizations. It may turn out, for all we know, that spacetime instantiates just such topological structure as $\mathbb{R}^4$ with a closed set excised (assuming, for the moment, that we can make sense in a physically substantive and cogent way of the idea of the global topological structure of spacetime). Perhaps the most important point to notice, though, is that "$\mathbb{R}^4$ with certain closed sets excised" is a *misleading* description of such a manifold. It suggests that we built that manifold from a more fundamental one, *viz.* $\mathbb{R}^4$. But that manifold *simply is* a manifold all on its own, with no intrinsic reference to $\mathbb{R}^4$, or indeed any other manifold. Because of certain facts about how we practice mathematics, the most convenient presentation of that manifold happens to be "$\mathbb{R}^4$ with certain closed sets excised". One could as legitimately present $\mathbb{R}^4$ as that manifold glued together with certain other manifolds-with-boundary. There are no good grounds I can see for suspecting that the universe heeds our preferred methods for organizing mathematical structures.

In event, I am happy to report that I do not need to rely on these constructions and considerations to demonstrate, for those unmoved by my sermon, that curvature pathology has no necessary connection to the existence of incomplete curves. More acceptable examples present themselves. The two most commonly used methods of measuring the growth of curvature intensity are the behavior of scalar-curvature invariants along some particular curve through the region of interest, and the behavior of the physical components of the Riemann tensor as measured by a frame parallel-propagated

21. This example is from (Wald 1984, p. 214). See (Ellis and Schmidt 1977, pp. 921–3) for further discussion of this sort of singular structure.

22. (Ellis and Schmidt 1977, p. 932) exemplify this sort of simplicity-chauvinism: "We know lots of examples of [flat singular spacetimes], all constructed by cutting and gluing together decent space-times; and because of this construction, we know that these examples are not physically relevant." See §3.5 for further remarks on this issue.

along some particular curve through the region of interest (if any of the physical components grow without bound in such a frame on a particular curve, or oscillate endlessly without settling down to a fixed, limiting value, then they will do so in all such frame-fields on that curve).[23] In accordance with customary usage, we will refer to the existence of an incomplete curve along which the physical components of the Riemann tensor in a parallel-propagated frame do not approach a finite, limiting value as *p.p.-singular structure*, and we will refer to the same of some scalar-curvature invariant along an incomplete curve as *s.p.-singular structure* ('s.p.' for 'scalar polynomial'). We will call the existence of an incomplete curve along which the physical components of the Riemann tensor in parallel-propagated frames and all its scalar invariants converge to finite values *quasi-regular singular structure*.[24] Note that curvature pathology on these definitions occurs not only if some feature of the curvature grows without bound along an incomplete curve, but also if it oscillates indefinitely (even if only within finite bounds), never settling down to a limiting value.

I believe there are two primary motivations for using a parallel-propagated frame in the terms of which to express the components of the Riemann tensor. First, one naturally expects the presence of curvature pathology to show itself, at the least, in misbehavior of the tidal forces an observer would experience along his or her worldline.[25] The intensity of tidal force, as measured in any frame, is directly proportional to the components of the Riemann tensor in that frame, as one can see by inspection of the equation of geodesic deviation. In a back-of-the-envelope sort of way, the unbounded growth of the components of the Riemann tensor in a parallel-propagated frame would seem to indicate that an observer traversing that curve would experience unbounded tidal forces as well. Second, (Clarke 1973) demonstrated that an incomplete curve in a singular spacetime has a local extension if and only if the relevant incomplete curve constitutes quasi-regular singular structure. A local extension is an isometric embedding of an open subset containing the incomplete curve from the spacetime manifold into another spacetime in which the (image of the) curve can be extended. Local extensions can exist even when the singular spacetime as a whole is not embeddable as a proper open submanifold into a larger spacetime in which the (images of the) incomplete curves can be extended.[26] Many take the existence of local extensions to indicate that nothing *local*, such as curvature pathology (narrowly construed), goes wrong in quasi-regular singular spacetime, but rather some global structure impedes the extension of spacetime.

The motivation for using the behavior of scalar-curvature invariants as a criterion for the existence of curvature pathology is somewhat more straightforward. First, all the points made with regard to the components of the Riemann tensor in parallel-propagated frames hold as well for scalar invariants. Even better, though, a scalar-curvature invariant at a point does not depend on what curve through that point or what frame on a curve through that point one uses to probe the point: it is, as the name suggests, invariant. Unbounded growth of a scalar-curvature invariant, moreover, is logically equivalent to the unbounded growth of the components of the Riemann tensor as measured in *every* frame-field along the curve, parallel-propagated or not.

S.p.-singular structure implies, but is not implied by, p.p.-singular structure. In fact, all scalar-curvature invariants can be zero and yet the Riemann tensor not be equal to zero, as in plane

23. A *frame* is a pseudo-orthonormal complete set of basis vectors for the tangent plane over a point of a manifold. A *frame-field* is an assignment of frames to points in some specified region, *e.g.*, along a curve or in an open set.

24. Quasi-regular singular structure is perhaps the most psychologically disturbing, since it can be absolutely inobservable until one runs into it, so to speak, creating a hair-raising hazard for spacetime navigation.

25. Tidal force is generated by the differential in intensity of the gravitational field, so to speak, at neighboring points of spacetime. For example, when I stand, my head is farther from the center of the Earth than my feet, so it feels a (practically negligible) smaller pull downward than my feet. For a graphic illustration of the effects of tidal forces on observers in strong gravitational fields, see the description in (Misner, Thorne, and Wheeler 1973, §32.6) of what would happen to a person standing on the surface of a collapsing star—not for the faint of heart, or weak of stomach.

26. (Ellis and Schmidt 1977, pp. 928–9).

gravitational wave spacetimes. [27] Spacetimes with colliding, thick gravitational waves provide examples of p.p.-singular structure in regions where all scalar-curvature invariants are well behaved; more strikingly, spacetimes containing colliding sandwich plane gravitational waves can exhibit p.p.-singular structure and yet all scalar-curvature invariants remain identically zero. [28] Finally, there are spacetimes containing colliding plane gravitational wave having incomplete curves in regions of a spacetime in which the Riemann tensor itself vanishes identically. These, the claim goes, provide examples of the existence of quasi-regular singular structure less artificial than that of the conical singularity above. [29] Thus the existence of incomplete curves does not *ipso facto* necessitate any sort of curvature pathology as conventionally quantified. That the misbehavior of the physical components of the Riemann tensor in a parallel-propagated frame or of a scalar-curvature invariant in the limit as one traverses a curve does not suffice to ensure that the curve be b-incomplete follows from examples of spacetimes produced by (Sussmann 1988) in which scalar-curvature invariants diverge asymptotically along complete timelike and null geodesics.

Though there is no necessary connection of any sort between the existence of incomplete curves and curvature pathology as quantified in the standard ways sketched above, (Ellis and Schmidt 1977) used b-completeness as a criterion to construct a classification of singular spacetimes according to the behavior of the curvature along the incomplete curves, as quantified in those standard ways. The classification has a binary, branching structure: first, an incomplete curve is said to constitute *essential singular structure* if there is no larger spacetime into which the singular spacetime can be embedded as a proper open submanifold, in which the (image of the) incomplete curve is extendible; otherwise it is said to be *inessential*. Essential singular structure is then sub-divided into quasi-regular and p.p.-singular structure; finally, p.p.-singular structure is subdivided into s.p.-singular and non-s.p.-singular structure. The thought behind the putative importance of the classification scheme seems to be as follows. Very little is known about singular structure at the present time, in part due to the difficulty of the mathematics involved in analyzing singular structure rigorously and in part due to the vanishingly small amount of experimental access we can get to singular structure in the foreseeable future. Nevertheless, the singularity theorems indicate that the spacetime we actually inhabit is singular, so it behooves us to try to understand such structure as much as possible. Classifying singular structure appears to be a way for us to organize and begin to get a grip on such a daunting task, and the scheme proposed by Ellis and Schmidt does seem to have many desirable features, such as clarity and simplicity. (Earman 1995, pp. 37, 43–4) goes so far as to proclaim one of the most seminal virtues of the definition of singular structure in terms of b-completeness that it allows for a classification of this sort.

To be appropriate for such a task, I submit, the mathematically different species of singular structure ought to exhibit sorts of physical behavior *prima facie* different from each other in a physically significant way, as near as one can judge that sort of thing with the crude tools at our disposal; otherwise it will be difficult to see the physical relevance of this so far purely mathematical classification. As already noted, in a spacetime with s.p.-singular structure, the Riemann tensor components will behave badly as expressed in *any* frame-field along the relevant incomplete curve, and, moreover, will do so in general along any curve close enough, as it were, to the incomplete curve. [30] The tidal forces a body will suffer as along its worldline are naturally measured in a spacelike

27. (Penrose 1960, p. 189).

28. (Konkowski and Helliwell 1992).

29. There is something odd about this claim—though it is accepted without question in the physics and in the philosophy literature—especially in comparison with the contrary claim concerning flat spacetimes with conical singularities: no observations we have made or can make in the foreseeable future rule out the possibility of the existence of conical singularities in the actual spacetime in which we reside, but innumerable observations we have already made demonstrate with a hard finality that our spacetime cannot contain radiation of *any* sort in a form approximating to that of plane-waves, not even in the most extravagant and inaccurate of approximations. Our spacetime is too lumpy. Oughtn't this make spacetimes containing plane-waves "more unphysical" than those containing conical singularities?

30. More precisely, in general there will exist an open neighborhood of the incomplete curve such that every curve

3-frame fixed rigidly in the body, orthogonal to the timelike unit vector tangent to the curve, used to fill out the full 4-frame. Based on what has already been said, one might expect that the state of motion of the observer along the curve, whether the observer is slowing down and speeding up somewhat, or spinning on his or her axis, would have no effect on how the observer experiences the curvature pathology: when a scalar-curvature invariant grows without bound along a curve, after all, the tidal forces as measured in *any* frame along the curve also will grow without bound. Interestingly enough, however, the state of motion of the observer as it traverses an incomplete curve, in the person of so-called inertial effects, can be decisive in determining the physical response of an object to the curvature pathology. Whether the object is spinning on its axis or not, for example, or accelerating slightly in the direction of motion, may determine whether the object gets crushed to zero volume along an s.p.-singular curve or whether it survives (roughly) intact all the way along the curve.[31]

The effect of the observer's state of motion on his or her experience of tidal forces can be even more pronounced in the case of p.p.-singular structure that is not s.p.-singular, which is precisely the existence of an incomplete curve along which there is a frame-field (necessarily not parallel-propagated) relative to which the components of the Riemann tensor approach definite, finite limiting values along the curve.[32] In such a case, the frame-field in which the physical components of the Riemann tensor stably approach a limit is related to any parallel-propagated frame-field by a Lorentz transformation that, in an appropriate sense, behaves pathologically in the limit along the curve. For a non-geodetic curve, the proper mode of transport along a curve of a frame rigidly fixed in the body of an object traversing that curve is not parallel-propagation but Fermi-transport.[33] A Fermi-transported frame is related to a parallel-propagated frame by a continuously varying Lorentz transform. It can happen, therefore, that an observer cruising along a p.p.-singular curve that is not s.p.-singular would experience unbounded tidal forces and so be torn apart while another observer, in a certain technical sense approaching the same limiting point as the first observer, accelerating and decelerating in just the proper way, would experience perfectly well behaved tidal force, though he would approach as near as one likes to the other poor fellow in the midst of being ripped to shreds. Again, certain gravitational plane wave spacetimes provide good examples of this phenomenon: an observer travelling along the incomplete timelike geodesic constituting the singular structure would experience unbounded tidal acceleration, whereas *any* observer travelling arbitrarily close by would not.[34]

Things can get stranger still. An incomplete geodesic contained entirely within a compact subset of a spacetime, with accumulation point p, satisfying a certain genericity condition, necessarily constitutes p.p.-singular structure, so that an observer freely falling along such a curve would be torn apart by unbounded tidal forces; it can easily be arranged in such circumstances, though, that a separate observer, who actually travels through p, will experience perfectly well behaved tidal forces.[35] Here we have an example of an observer being ripped apart by unbounded tidal forces right in the middle of spacetime, as it were, while other observers cruising peacefully by could reach out to touch him or her in solace during the final throes of agony. This discussion points to a startling conclusion: curvature pathology, as standardly quantified, is not in any physical sense a well defined property of a region of spacetime *simpliciter*; rather, whether or not phenomena both

completely contained in the open neighborhood has Riemann components that are as badly-behaved as one likes in all frames along the curve. The 'in general' hedges against the case where the scalar curvature invariant oscillates wildly along the incomplete curve; in this case, it may be possible for nearby curves to weave cleverly around the incomplete curve in such a way as to avoid the peaks of oscillation, and so have well behaved Riemann tensor components. No hard results are known either way in such cases.

31. (Ellis and Schmidt 1977, p. 944–7).
32. (Ellis and Schmidt 1977, p. 939).
33. (Hawking and Ellis 1973, pp. 80–1).
34. (Ellis and Schmidt 1977, p. 937).
35. (Hawking and Ellis 1973, pp. 290–2).

physically pathological itself and attributable directly to pathology in the behavior of the Riemann tensor manifest themselves may sensitively depend on the sorts of devices, including their states of motion, with which one probes the region of interest! These matters are far more subtle and complicated than many, including (Earman 1995), would lead one to believe.

(Ellis and Schmidt 1977, p. 918) say, *vis-à-vis* their classificatory scheme (the canonical one):

> It is not claimed here that the singularities discussed are *likely* to occur in physically realistic situations, but rather that only when we understand which singularities can occur (a) in general space-times, and (b) in space-times with the field equations satisfied for particular matter content, can we hope to discuss fruitfully their occurrence, equations of motion, and so on.

I do not mean to argue with the motivation for their classificatory scheme, but they beg a serious question with their 'which' in the phrase "when we understand which singularities can occur": clearly the correlative demonstratives of this relative interrogative refer to the different species of their classification, but why ought one think that their classification picks out physically relevant differences among all possible singular structures? This question becomes more poignant when one reflects on the fact that curvature pathologies provide the differentiæ for their speciation, and, as I have attempted to show, curvature pathology as customarily quantified is not a straightforward concept with clear and unambiguous physical content. I believe there is far more work to be done straightening out the physical consequences of the existence of singular structure. The mathematics has outrun the physics, but still masquerades as such.

Taub is the only person I know who shares in print my apprehension about the content of the canonical classification scheme: [36]

> I have difficulty understanding the usefulness of the classification scheme of singularities proposed... by Ellis and Schmidt.... I think that the important work on singularities now being done would become much more important if it turned toward learning how to deal with the physics associated with singularities.... [37]

He appears to be saying that one ought to concentrate first on trying to work out the behavior associated with various singular structures we are more or less familiar with in a clear and unambiguous way, and only then should one feel confident enough to begin classifying singular structures, based on that clear physical knowledge, not on a purely mathematical scheme that becomes murky as soon as one tries to think about it in physical terms. I heartily concur. [38]

3.4 Missing Points

We now have a precise definition of a singular spacetime, and some ideas about what such structure implies and does not imply about the curvature of spacetime, but, as Earman notes, "it is not true to an idea that is arguably a touchstone of singularities in relativistic spacetimes: spacetime singularities correspond to missing points." [39] For those who would argue missing points ought to be such a touchstone, Earman sketches what seems to me the most promising position, that, though

36. Though Geroch has told me in conversation that he does not see the use of the classification scheme either, on the grounds that he cannot see what physical content it has.

37. (Taub 1979).

38. A physically unambiguous sense of curvature pathology occurs in, *e.g.*, the FRWL (Friedmann-Roberston-Walker-Lemaître) metrics, wherein physical quantities such as the mass-density of ponderable matter grow without bound along incomplete curves and thus scalar-curvature invariants correlatively grow without bound as well. This sort of idea is developed nicely in a not very well known paper (to judge by its citation record) by (Thorpe 1977). I think it would be of interest to see whether a classification scheme based on some of Thorpe's ideas could be constructed and compared to the canonical classification.

39. (Earman 1995, p. 40).

the idea of missing points and that of curve incompleteness lead to *prima facie* different concepts of singular structure, they are extensionally equivalent in all physically reasonable singular spacetimes, and so the two concepts are for all practical purposes in agreement. [40] I will argue with this: missing points ought not be a touchstone of discussion of singular structure in relativistic spacetimes.

Missing points, could they be defined, would correspond to a boundary for a singular spacetime—actual points of an extended spacetime at which incomplete curves would terminate. [41] My argument therefore will alternate between speaking of missing points and speaking of boundary points, with no difference of sense intended. In many cases of physical interest, such as the FRWL and Schwarzschild metrics, one can attach boundary points by hand, so to speak, by visual inspection of the metric expressed in an appropriate global coordinate system, though different coordinate systems can lead to different topological structure for the boundary points. If one is to have a general notion of missing points, corresponding to the existence of incomplete curves, determined by nothing more than the metric structure of spacetime, clearly what is wanted is a method of attaching boundary points that does not depend on the choice of coordinate system, and which, moreover, can be used for any singular spacetime, not just the ones with 'simple' metric structure and global topology.

Before I begin examining the primary attempts to define boundary points for singular spacetimes, [42] it is well to note two oddities of the situation. In the case of a manifold with a Riemannian metric, Cauchy completion provides a well defined notion of missing points, and, by the Hopf-Rinow theorem, no points are missing from the manifold if and only if all geodesics of the Riemannian metric are complete (see footnote 4). We have already seen that any definition of missing points for a spacetime may—perhaps ought—not satisfy this condition: a spacetime can be geodesically complete yet still be b-incomplete, as Geroch's example illustrates. This already suggests that, even were one able to come up with a satisfactory definition of missing points in the context of Lorentzian metrics, it may not be extensionally equivalent to the existence of incomplete curves of the physically relevant type. The second, and more striking, circumstance strengthens this suspicion: compact spacetimes can contain incomplete, inextendible geodesics, as shown by a simple example due to (Misner 1963). In a sense that can be made precise, compact sets, from a topological point of view, "contain every point they could possibly be expected to contain", [43] one consequence of which is that a compact manifold cannot be embedded as an open submanifold of any other manifold, a necessary pre-requisite for attaching a boundary to a singular spacetime—a manifold-with-boundary minus its boundary is embeddable by the identity map as an open submanifold into itself, and, in the case when the manifold-with-boundary has a (pseudo-)Riemannian metric, the embedding can be made an isometry from the manifold *cum* metric on the full manifold to the interior of that manifold endowed with the natural restriction of the full metric. We ought not expect then that any definition of a boundary for singular spacetimes will cover every possible kind of singular structure, unless we are willing to swallow *outré* topological structure. [44]

---

40. (Earman 1995, p. 42). Earman continues on to say that even were one to grant this claim, the concept of singular structure as based on the idea of missing points is still conceptually distinct from that based on incomplete curves, and deserves in its own right to be examined. In reading the rest of the book one wishes that Earman had sketched a little more what he had in mind here. In particular, in later chapters, where Earman maps out issues and problems associated with the existence of singular structure, he never clarifies which conception of singular structure he is working with, and how opting for one or the other of the two conceptions would alter the character of the issue or problem at hand.

41. Strictly speaking, such a space would not be a manifold in the usual sense of the term, but a manifold with boundary. See (Spivak 1979a) for a discussion of manifolds with boundary.

42. I will not consider in this paper the 'ideal-point' boundary construction of (Geroch, Kronheimer, and Penrose 1972), as it requires the singular spacetime to be past- and future-distinguishing, a fairly strong causality condition. I intend to sidestep all questions about the physical plausibility or necessity of such conditions.

43. See (Geroch 1985, §30) for a discussion of this precise sense.

44. As an aside, this discussion highlights the fact that the interplay between metric and topological structure in the case of manifolds endowed with pseudo-Riemannian metrics is a far more delicate matter than it is, in general, acknowledged to be, and, in any event, *far* more delicate than in the case of a Riemannian manifold. This point

(Schmidt 1971) produced the most well known boundary construction for singular spacetimes, the so-called b-boundary based on the b-completeness criterion. An affine connection on a manifold allows one to define in a natural way a family of Riemannian metrics on the frame bundle over that manifold equivalent in the sense that they yield the same topology for the bundle-manifold, the natural topology of the frame bundle. It follows that the bundle-manifold is Cauchy complete with respect to one of these metrics if and only if it is so with respect to all. Schmidt showed, moreover, that the bundle-manifold is Cauchy complete in this sense if and only if the spacetime manifold is itself b-complete. To complete a singular spacetime in this scheme, then, one lifts all the incomplete curves from the spacetime manifold to the frame bundle, takes the Cauchy completion of the frame bundle with respect to one of the family of natural, topological, Riemannian metrics, and "projects down" the constructed boundary from the frame bundle to form a boundary for spacetime.

The relativity community at first embraced Schmidt's contruction with enthusiasm, to judge by the remarks in chapter 8 of Hawking and Ellis's canonical work *The Large Scale Structure of Space-Time*. Shortly thereafter, however, Bosshard and Johnson separately showed that the b-boundary had undesirable properties in the most physically relevant spacetimes known, the FRWL spacetimes, which to a quite high degree of approximation accurately model the large scale structure of the actual universe, and the Schwarzschild spacetimes, which represent the neighborhood of spherically symmetric isolated bodies, such as stars. [45] For closed FRWL spacetimes, the b-boundary consists of a single point (the same for the big bang as for the big crunch) that is not Hausdorff-separated from any point in the interior of the spacetime. Not only does one reach the same point, then, by travelling either forward or backward in time, but that point is, in a certain sense, arbitrarily near every single spacetime event! Similarly, the b-boundary of a Schwarzschild spacetime consists of a single point not Hausdorff-separated from any interior point of the spacetime. This certainly will not do for the advocates of missing points. [46]

A second (albeit temporally prior) method of constructing a boundary for singular spacetimes due to (Geroch 1968a) fares much better with physically relevant spacetimes. [47] In this construction, the so-called *g-boundary*, geodesic incompleteness rather than b-incompleteness defines singular structure, and one defines a boundary point to be an equivalence class of incomplete geodesics under the equivalence relation 'approach arbitrarily close to each other' (in a certain technical sense). The set of boundary points can be given a topology and, in many cases of physical interest, can even be given a differentiable and metric structure, so that one can locally analyze the structure

---

does not seem to command the attention I think it ought to, in the philosophical as well as the physical literature; indeed, the matter often seems to be handled with a surprisingly cavalier attitude. It is not uncommon, for example, for a derivation of the Schwarzschild solution to yield, as the solution, the standard Schwarzschild coordinates, after which the deriver, almost always without comment, takes the topology of the spacetime to be $\mathbf{R}^2 \times \mathbf{S}^2$ as a matter of course. The same often happens with derivations of the FRWL spacetimes as well. Of course, a presentation of the metric in a particular coordinate-system does *not* determine the global topological structure of the manifold. It determines only *local* geometry. It makes as much (or as little) sense to take the topology of a spacetime whose metric can be represented in the Schwarzschild coordinates to be $\mathbf{R}^4$ as it does to take it to be $\mathbf{R}^2 \times \mathbf{S}^2$. Compare this state of affairs with that holding for a Riemannian manifold. The induced pointwise-metric (distance function) on the manifold defines a natural topology on the manifold in the standard way: one demands that the family of open balls of all radii centered on all points form a sub-basis for a topology. Conversely, a topological manifold selects a preferred family of Riemannian metrics one may impose on it, those that yield the already given topology.

45. *Cf.* (Bosshard 1976), (Johnson 1977), (Bosshard 1979) and (Johnson 1979).

46. The reactions to these problems vary widely. (Clarke 1993), for instance, still embraces the b-boundary contruction, and defines a singularity to be a point on the b-boundary of a singular spacetime (§3.4). He barely mentions these problems, noting only in passing that the topological structure of the singular spacetime with boundary can be "very strange,"(p. 40) which I do not think qualifies as an adequate address of the issue. (Wald 1984), on the other hand, does not like the b-boundary construction precisely because of these problems (*cf.* pp. 213–4), and (Joshi 1993) does not even mention the possibility of attaching boundaries to singular spacetimes, speaking only of incomplete curves.

47. (Hawking 1966c) apparently proposed a similar construction, but, as the essay was never published (I learned of it from the bibliography of (Hawking and Ellis 1973)), I have not been able to get a hold of it for examination.

of spacetime at a 'singularity' rather than mess around with troublesome limits along incomplete curves.[48] The g-boundary construction, moreover, yields the boundaries one might have expected on physical grounds in spacetimes of particular physical interest: the g-boundary of a Schwarzschild spacetime is a spacelike 3-surface, topologically $\mathbb{S}^2 \times \mathbb{R}$, and that of a closed FRWL spacetime is the disjoint union of two spacelike $\mathbb{S}^3$'s. Pathological topology rears its head here as well, though, in the case of Taub-NUT spacetime:[49] the g-boundary of this spacetime contains a point that again is not Hausdorff-separated from any point in the interior of the spacetime.

The advocate of missing points who wants to hold on to the g-boundary may at this point retort that Taub-NUT spacetime hardly constitutes a physically relevant spacetime for other reasons, namely that it violates strong causality, which is to say that it contains causal curves that come arbitrarily close to intersecting themselves. While I do not think this reply carries much weight,[50] I have a better example at hand. (Geroch, Can-bin, and Wald 1982) construct a geodesically incomplete spacetime with no causal pathology for which a very large class of boundary constructions, including the b- and the g-boundary, will yield pathological topology in the completed spacetime. The conditions that a boundary construction must satisfy to fall prey to this example are quite weak: each incomplete geodesic of a singular spacetime must terminate at some boundary point; and, in a certain technical sense, the boundary points corresponding to incomplete geodesics that are 'close together' must also be 'close together'. The advocate of missing points may point out that the example appears artificial and contrived, with closed sets excised here and conformal factors plastered on there, and in short has no physical relevance. I would reply with the lesson of my sermon from §3.3, and a remark that (Geroch, Can-bin, and Wald 1982, p. 435) make: "The purpose of [a boundary] construction, after all, is merely to clarify the discussion of various physical issues involving singular space-times: general relativity as it stands is fully viable with no precise notion of 'singular points.'" When we contemplate potential phenomena that we have little or no observational access to, I submit that the standards for what can count as a *physical* account of a situation ought to be priggishly severe, if we are not unwittingly to degenerate into pure mathematical discourse.[51] A boundary-construction that yields topological pathology, and contains no precise criteria for what ought to count as a 'physically relevant' spacetime, does nothing to clarify discussion of the physical issues involved in analyzing singular spacetimes.

The abstract-boundary construction, or *a-boundary*, proposed by (Scott and Szekeres 1994) appears at first glance to have the most promise for those wanting a natural, workable definition of missing points for singular spacetimes.[52] It also nicely exemplifies a feature of all missing point constructions I know of or can easily imagine, their dependence on a prior characterization of incomplete curves. For these two reasons, I will consider it in a little more detail than the previous two. An *envelopment* of a manifold $\mathcal{M}$ is an ordered pair $(\mathcal{N}, \phi)$ consisting of a manifold $\mathcal{N}$ and an embedding ϕ into $\mathcal{N}$ of $\mathcal{M}$ as a proper open submanifold of the same dimension.[53] Scott and

48. In certain contrived examples, there is an ambiguity in choice of topology for the g-boundary, but I will waive this concern for the sake of argument. I have bigger fish to fry.

If one suspects that this use of 'contrived' represents a hypocrisy now on my part—well, it may and it mayn't. To quote Geach quoting Whitman, "Do I contradict myself? Very well, I contradict myself. I am large, I contain multitudes."

49. *Cf.* (Hawking and Ellis 1973, §5.3) for a thorough account of Taub-NUT spacetime.

50. (Earman 1995, chs. 6–7) explains better than I could why a violation of strong causality *simpliciter* does not constitute an argument for the unphysicality of a spacetime.

51. Geroch stressed this point to me in a conversation in which he also dismissed the adequacy of his own g-boundary construction *merely because* it gave unphysical results in the admittedly contrived (!) example of (Geroch, Can-bin, and Wald 1982). It gives very nice results in almost all other known types of examples.

52. Whether the a-boundary construction satisfies the conditions of (Geroch, Can-bin, and Wald 1982), and so necessarily leads to pathological topology for certain spacetimes, is not clear, for as of yet Scott and Szekeres have not defined a topology on the relevant entities of their construction. From the structure of the construction, I suspect that any topology one could more or less naturally define for it would satisfy Geroch, Can-bin, and Wald's conditions.

53. When it can cause no confusion, I will often identify $\mathcal{M}$ with its image under the envelopment mapping.

Szekeres propose that singular structure always arises by the deletion of points from an envelopment of a singular manifold. Given an envelopment ($\mathcal{N}$, ϕ) of $\mathcal{M}$, a subset of its topological boundary in $\mathcal{N}$ will be called a *boundary set*. Now, as it clearly is possible to envelop a given manifold in many ways (if the manifold has any envelopment at all), one does not want to consider merely boundary sets of manifolds under particular envelopments, but rather equivalence classes of boundary sets under some appropriate equivalence relation. To this end, Scott and Szekeres propose the following:

Definition 3.1 *A boundary set B of $\mathcal{M}$ in an envelopment ($\mathcal{N}$, ϕ) is said to* cover *the boundary set B' of $\mathcal{M}$ in an envelopment ($\mathcal{N}'$, ϕ') if for every open neighborhood U' in $\mathcal{N}'$ of B' there exists an open neighborhood U in $\mathcal{N}$ of B such that*

$$\phi \circ \phi'^{-1}[U' \cap \phi'[\mathcal{M}]] \subset U.$$

A boundary set B may cover another boundary set B' while B' does not cover B. One easily sees, however, that defining B and B' to be equivalent if they mutually cover each other does in fact yield an equivalence relation; the equivalence class of the boundary set B under this relation will be written '$[B]$' and called an *abstract boundary set*. An equivalence class that contains a singleton as a representative member will be called an *abstract boundary point*. The collection of all abstract boundary points is the *abstract* or *a-boundary*, written '$\mathcal{B}[\mathcal{M}]$'.

Although $\mathcal{B}[\mathcal{M}]$ by itself is defined without reference to any particular geometrical structure on $\mathcal{M}$, such as a pseudo-Riemannian metric or an affine connection, which Scott and Szekeres take to be one of its cardinal virtues, to define singular structure they must select a class of curves $\mathcal{C}$ on $\mathcal{M}$ satisfying what they call the bounded-parameter property: roughly speaking, the curves in $\mathcal{C}$ must cover the manifold and must be such that the parameter along any of the curves grows without bound if and only if it grows without bound along every "nice" reparametrization of the curve. The class of geodesics on a manifold with affine connection and the class of C^1 curves parametrized by generalized affine parameter on a manifold with affine connection provide two examples of classes of curves satisfying the bounded-parameter property. The idea is that curves in $\mathcal{C}$ will be used to probe the boundary to distinguish points 'at infinity' from points that can be reached in a finite parameter interval and hence are candidate singular points. The details of the construction and definitions hereon out become quite complicated, so I will sketch only the most salient points.

First, for a candidate singular spacetime $\mathcal{M}$, Scott and Szekeres wish to remove from consideration all abstract boundary points that have a representative singleton boundary point in some envelopment through which, in a certain technical sense, the spacetime metric can be smoothly extended. In this case, the thought is, the original spacetime simply had not been made as 'large' as it reasonably could have. Such points will be called *regular*, and need not apply as potential singular points. Next, one fixes the class of curves $\mathcal{C}$, and defines the $\mathcal{C}$-*boundary* to be the class of a-boundary points that have, in some envelopment, a singleton representative that is the limit point of a curve in $\mathcal{C}$; such points are also referred to as *approachable*. All other a-boundary points are *unapproachable*. It is straightforward to show that the property of being approachable or unapproachable is invariant under the defining a-boundary equivalence relation, but one must keep in mind that it depends entirely on the class $\mathcal{C}$ chosen. Given an envelopment $\mathcal{N}$ of $\mathcal{M}$, a non-regular boundary-point that is not the limit point of any curve of bounded parameter in $\mathcal{C}$ will be called a *point at infinity*; if, moreover, it cannot be covered by any regular boundary set of another envelopment, it will be called an *essential* point at infinity. This property is clearly invariant under the a-boundary equivalence relation, and so one speaks of a-boundary points at infinity. A non-regular boundary point p that is the limit point of some curve in $\mathcal{C}$ of bounded parameter will be called a *singular point*. If there exists a non-singular boundary set of another envelopment that covers p, then it is said to be *removable*; otherwise it is *essential*. Again, this property is invariant under the a-boundary equivalence relation, so one says that $[p]$ is an essentially singular a-boundary point. These, finally, are the missing points Scott and Szekeres aimed to construct.

The most obvious problem facing the a-boundary approach is its physical significance, or lack thereof. First off, a 'point' of the a-boundary is not a point in any usual sense of the term: an individual boundary point of one envelopment of a manifold can always be made to cover an uncountable number of boundary points in another envelopment. It is the case that, given any envelopment, the representative boundary set of an a-boundary point in that envelopment must be compact, but it is not even true that every compact boundary set is a representative of some a-boundary point, nor does the a-boundary point equivalence relation preserve connectedness and simple-connectedness—ought one think of a candidate singularity as a single point or as a non-simply connected, non-connected compact set? Then there is the unapproachability of some a-boundary points: it can happen, for instance, that regular a-boundary points of a pseudo-Riemannian manifold are not approachable by any geodesic of the metric. The existence of such extraneous points makes one wonder about the physical relevance of those boundary points that are approachable by curves in the spacetime. It is not also not clear what relevance the 'covering' relation they define has to anything physical: for a given $\mathcal{C}$, $\mathcal{C}$-boundary sets may cover unapproachable boundary sets; non-regular unapprochable boundary sets may cover approachable regular boundary sets; essential boundary points at infinity may cover anything except singular boundary sets and may be covered by anything except regular points; essential singular points may cover any kind of boundary set. Given the promiscuity of possible covering relations, I believe an argument is needed why this definition captures any physically relevant information, an argument they do not provide.

Neither do Scott and Szekeres broach a technical point that raises a serious difficulty for their approach at the very initial stages: some spacetimes, such as Taub spacetime, have two incomplete curves such that the spacetime can be extended so as to make either one or the other curve extendible, but no extension of the spacetime exists that makes both curves simultaneously extendible.[54] On Scott and Szekeres's account, both of these curves run into regular boundary points, and so neither will be counted as possible singularities, even though there is no actual envelopment of the spacetime in which both curves are simultaneously extendible.

Finally, on this view, incomplete curves wholly contained in compact regions of spacetime cannot count as singular structure, trivially so since compact manifolds cannot be embedded as proper open submanifolds of another manifold. Scott and Szekeres not only gamely swallow this consequence, but actually claim that it is a "*sine qua non* of any successful theory of singularities,"[55] and cite (Shepley and Ryan 1978) as evidence for this claim.[56] This is not only a contentious view, at best, which they do not bother to argue for, and not only seems to run counter to the spirit of most considerations forwarded in discussions of singular structure, which revolve around incomplete curves, but seems seriously to conflict with their own stated criterion for selecting those points of the a-boundary that will be singular points, *viz.*, limit points of curves of bounded parameter, *i.e.*, curves that are, in some sense or other with (presumed) physical significance, incomplete.

This last point brings out my final consideration against the idea of missing points as touchstones in the investigation of singular spacetimes: the definition of singular spacetimes by incomplete curves is logically prior to the construction of missing points for singular spacetimes. All the missing point constructions I know of, and all the ways I can more or less easily imagine trying to concoct a new one, rely on probing the spacetime with curves of some sort or other to discover where points may be thought of as missing, just as in the Riemannian case one cannot complete a manifold until one knows which Cauchy sequences do not have a limit point, or equivalently which geodesics are incomplete. Even Scott and Szekeres, who make much of the fact that the construction of their a-boundary *per se* does not depend on the existence of any particular geometrical structure on a manifold, such as an affine connection and incomplete curves, cannot define singular points, which after all was the

54. See, *e.g.*, (Ellis and Schmidt 1977, p. 920) and (Hawking and Ellis 1973, §5.8).
55. (Scott and Szekeres 1994, p. 34).
56. In fact, Shepley and Ryan provide only the briefest and most tendentious of justifications for this position.

point of the whole affair, without probing their boundary with some specified class of curves. [57] One, however, does not need any conception of a missing point, much less a definition of such a thing, to define and investigate the existence of incomplete curves on a manifold. In sum, I disagree with the gist of much of the discussion of (Earman 1995, ch. 2), wherein he suggests that unclarity plagues the semi-official definition of a singular spacetime, in terms of b-incompleteness, in so far as, on the face of it, one does not know how it relates to the idea of missing points. Incomplete curves seem to me a fine definition of singular structure on their own. I will try to make these considerations more precise in the following section.

3.5 Global *vs.* Local Properties of a Manifold

There is at least one *prima facie* good reason why it would be useful to have a precise characterization of points missing from singular spacetimes: one would then be able to "paste the points to the boundary of the spacetime manifold" and so analyze the structure of the spacetime locally at the singularity, instead of taking troublesome, perhaps ill-defined limits along incomplete curves. The power and elegance of Penrose's construction of conformal infinity for asymptotically flat spacetimes lies precisely in the ability one gains to perform such analysis locally at infinity, without relying on limits. [58] The example of (Geroch, Can-bin, and Wald 1982) already discussed makes the prospects for a reasonable boundary construction for singular spacetimes grim. I believe this should not have been surprising.

In desiring a boundary so as to have a place to analyze structure locally, one ought to be clear on what one means by 'locally'. One sometimes hears talk of a global, as opposed to a local, feature of a spacetime, but I know of no precise characterization of the difference. I believe this distinction plays a crucial role in a proper understanding of the standardly proposed definitions of a singular spacetime in terms of incomplete curves. I therefore offer the following precise definition of this distinction. I formulate it initially for topological properties both for the sake of generality and because I think it easier to get a feel for the definition in the sparser arena of topological structure than in the more cluttered arena of differentiable manifolds with an affine structure.

Consider the class $\mathfrak{T}$ of all topological spaces. A *topological property* $\mathfrak{P}$ is a subclass of this class. A topological space $\mathcal{S}$ has the property $\mathfrak{P}$ if $\mathcal{S} \in \mathfrak{P}$.

Definition 3.1 *A topological property* $\mathfrak{P}$ *is* local *if it has the following feature: a given topological space* $\mathcal{S}$ *has the property* $\mathfrak{P}$ *if and only if* $\mathcal{S}$ *is such that every neighborhood of every point has a subneighborhood that, considered as a topological space in its own right, with the restriction topology, has the property* $\mathfrak{P}$*.* [59]

Roughly speaking, a local property must hold in arbitrarily small neighborhoods of every point of a topological space, but not necessarily in every neighborhood of every point of the space; and conversely, if the property holds in arbitrarily small neighborhoods of every point of a space, it must hold for the entire space for it to be local.

Definition 3.2 *A topological property is* global *if and only if it is not local.* [60]

57. I thus think that (Earman 1995, p. 42) was off-base when he suggested that the a-boundary might be used to "do justice" to the idea of missing points for singular spacetimes.

58. See (Wald 1984, §11.1) for an account of Penrose's construction.

59. This sense of 'local' has nothing to do with that often bandied about in discussions about the foundations of quantum mechanics.

60. By this 'not', I do not mean the logical negation of the definition of 'local' but rather the class complement of the class of local properties in the class of all topological properties—interestingly enough, these do not come to the same thing. Were the logical negation of the definition of 'local' used to define 'global', this would entail that a space with the global property $\mathfrak{P}$ would have a point and a neighborhood of that point such that *every* subneighborhood of

One could be sure of ascertaining for a given topological space whether the local property $\mathfrak{P}$ held or not by checking for $\mathfrak{P}$ at individual points of the space (quite a few points, to be sure), whereas a global property cannot be checked by examining the structure of the space at any collection of points. As one should expect, local compactness, local connectedness and local simple connectedness for example all come out to be local on this definition, whereas compactness, paracompactness, connectnedness and simple connectedness come out to be global.[61]

In an analogous manner, one can now straightforwardly characterize properties of differentiable manifolds and of differentiable manifolds with an affine connection as either local or global. Nontrivial examples of local properties for a manifold include any structure residing entirely on the tangent planes over every point. For our purposes, the most important fact about a manifold with an affine connection arising from a pseudo-Riemannian metric is that both the property of geodesic completeness and of geodesic incompleteness come out to be global properties, again as one should expect. One might initially have thought that geodesic incompleteness, at least, ought to have been a local property—if a geodesic came to an end abruptly, as it were, surely one ought to be able to pinpoint where this happens. If one could do this, however, then it also would seem that one could continue the geodesic. If there were a point on the manifold where the incomplete geodesic terminated, one could, around that point, take a chart diffeomorphic to some open set of $\mathbb{R}^n$ (assuming the manifold does not already have a boundary), push the geodesic and the connection down to $\mathbb{R}^n$, where the geodesic obviously would be extendible, and pull the extended version back to the manifold, contradicting the hypothesis that the geodesic could not be continued. This cannotbe done, however, for incomplete geodesics of a pseudo-Riemannian metric. All attempts to construct "missing points" founder on this rock.

A point of spacetime, in the usual way of thinking of these matters, represents an *event*, a highly localized occurrence in spacetime such as a snapping of fingers or the collision of two billiard balls. It represents an instant of some ponderable object, the specious 'now' of some sentient being. When thinking on cosmic scales, the sun, at a certain instant, can profitably be thought of as occupying a single point of spacetime. In short, spacetime points pertain to discrete objects, very broadly construed, that can be localized in an intuitive sense. There is no *a priori* reason to suspect that the existence of an incomplete curve, a global phenomenon, could be tied in any natural or reasonable way to the existence of a particular point in an extended manifold. Incomplete curves are not discrete, localizable objects in the appropriate sense.

A detractor will likely balk at this line of thought, pointing to the case of Riemannian manifolds, wherein incomplete curves can be naturally associated with points of an extended manifold. I would reply that it is merely a happy accident in the Riemannian case that one can arrange this. One has no grounds for suspecting that one will be able to do this in the general case, and in fact, as I endeavored to show, one has reasons to suspect that in general one will not be able to do this, since curve incompleteness is global and a missing point is, well, a point, and so *prima facie* "local". Of course, even for Lorentzian manifolds, in certain cases, one will be able to associate to an incomplete curve a missing point in a natural way—*e.g.*, in Minkowski spacetime (in some global coordinate system) with the origin removed, to continue all the geodesics aimed at the missing origin one pastes the origin back into the space and continues the geodesics through that point—in general, though, one ought not expect the two to have anything to do with each other.

The demand that singular structure be localized at a *place* bespeaks an old Aristotelian substantivalism that invokes the maxim, "To exist is to exist in space and time."[62] When I speak of 'Aristotelian substantivalism' here, I refer to the fact that Aristotle thought that everything that exists is a substance and that all substances can be qualified by the Aristotelian categories, two of

that neighborhood did *not* have $\mathfrak{P}$. Compactness is clearly not a local property, and yet does not satisfy the logical negation of the definition of 'local'.

61. *Cf.* (Hocking and Young 1988) for definitions of these topological properties.

62. This formulation of the maxim is due to (Earman 1995, p. 28).

which are location in time and location in space. In particular, not only substantivalists but also relationalists in debates about the nature of spacetime points could (and often do, I think) consistently fall prey to this particular brand of substantivalism. By focusing attention on the way that spacetimes can have actual features that do not rely on the existence or absence of any particular point, and are not instantiated at any particuar point, I suspect that this distinction between global and local properties of spacetime could have a salutary effect on the moribund debate between substantivialists and relationalists. To lay my cards on the table, I suspect one could parlay these considerations into a persuasive argument for the most salutary (to my mind) of effects on that debate, its dismissal as a *Scheineproblem*.[63]

(Geroch, Can-bin, and Wald 1982, p. 435) deserve the last word on this subject: "Perhaps the localization of singular behavior will go the way of 'simultaneity' and 'gravitational force.'"

3.6 The Finitude of Existence

In this paper I have examined the standard characterizations of singular spacetimes and rejected attempts to link singular structure to the existence of missing points, arguing that the characterization of singular structure in terms of incomplete curves is adequate for the purposes of all known sorts of physical investigations touching on the subject. In the end, this is the only criterion I know of that ought to matter when the issue is the cogency and cognitive content of a proposed physical notion and concomitant methods of physical investigation and argumentation framed in the terms of that notion. Before concluding, I turn to examine whether singular structure as thus characterized is objectionable on physical or interpretive grounds, and whether one is forced to or ought to take them as indicating the 'breakdown' of classical general relativity, as some would have it. In the process, I will examine whether b-completeness is wholly consistent with some of the explicit sentiments behind using curve incompleteness as a criterion for singular structure.

Two types of worries, one psychological, the other physical, give rise to the dissatisfaction with the existence of incomplete curves in relativistic spacetimes. Trying to imagine the experience of an observer traversing one of the incomplete curves provokes the psychological anxiety, for that observer would, of necessity, be able to experience only a finite amount of proper time's worth of observation, even were he, in Earman's evocative conceit, to have drunk from the fountain of youth. The physical worry arises from the idea that particles could pop in and out of existence right in the middle of a singular spacetime, and spacetime itself could simply come to an end, as it were, though no fundamental physical mechanism or process is known that could produce such effects. These two types of worries are not always clearly distinguished from each other in discussions of singular structure, but I think it important to keep in mind that in fact there are two distinct types of problems envisaged for incomplete curves, requiring to some degree two separate sorts of response.

The existence of incomplete spacelike curves is often felt not to be so objectionable as that of incomplete timelike or null curves, on the grounds that it represents structure beyond the direct experience of any observer.[64] I submit that, on this criterion, neither ought one be so bothered by the existence of incomplete timelike or null curves, *for an observer travelling along such a curve will never directly experience the fact that he has only a finite amount of proper time to exist*—there is no spacetime point, no event in spacetime, that corresponds to the observer's ceasing to exist. This is not to say that the person traversing this worldline cannot surmise the fact, perhaps based on observation of the curvature in his immediate neighborhood, that he has only a finite amount of time to exist; the claim, rather, is that there will never be an instant when the observer experiences himself as dissipating, popping out of existence as it were. To disarm possible misunderstanding, I

63. I would base such considerations and argument on those delivered by (Stein 1989), against the cogency of the traditional debate between realists and anti-realists, so called.

64. See, *e.g.*, (Hawking and Ellis 1973, §8.1).

emphasize that I am referring to, not the "popping out of existence" due to the observer's possibly being torn apart by unbounded tidal acceleration or being shot in the midst of his experiments by a Luddite lunatic, but the "popping out of existence" that would come about because the observer actually reached the "end" of his worldline, so to speak—for there is no end of the worldline to reach! We may be unable to conceive of experiencing such a state of affairs, but this reflects limitations in our psychological constitution, not an inherent flaw in general relativity. [65]

These considerations suggest as well a tension between the definition of singular structure by b-incompleteness on the one hand and the intuitions that drove some to look to incomplete curves as marks of singular structure in the first place on the other. Only the finitude of proper time matters so far as the experience of a possible observer goes—a generalized affine parameter has no clear physical significance—but, while a curve's being b-incomplete implies that the curve is of finite total proper time, the converse is not true: timelike curves of unbounded total acceleration in Minkowski space can be of finite total proper time and yet be b-complete. I would even say that such a curve should be more disturbing on reflection to those with such intuitions than an incomplete null geodesic, for the concept of 'proper time' does not apply to null curves at all, even though they are the possible paths of massless particles. The few people who even remark on the tension usually mouth a few vague generalities about particles' "reaching infinity", the implication seeming to be that, in so far as particles tracing out such worldlines are able to accomplish this *recherché* feat, one should have no qualms about the discomfiture they may feel in having only a finite amount of proper time in which to exist.

I speculate, with no hard evidence, that people have not wanted to count such curves as constituting singular structure in large part because of vague worries about energy conservation—an observer would require an "infinite amount of energy" to traverse a curve of unbounded total acceleration. In general relativity, however, there is no rigorous, generic notion of energy conservation, not globally or locally—there is not even a rigorous, generic, invariant definition of 'energy'. [66] Indeed, the structure of general relativity offers up no *a priori* reason to suspect that it in any way excludes a particle's getting shot out asymptotically "to infinity" in finite total proper time, having started from perfectly regular (in whatever sense of that term one likes) initial data. After all, it is not even difficult to construct solutions to Maxwell's equation on Minkowski spacetime in which a charged test particle gets shot off "to infinity" in finite total proper time. [67]

An example of a spacetime that was b-complete for all timelike curves of bounded total acceleration but not for timelike curves of unbounded total acceleration would clarify some of these issues, and I conjecture that examples of such spacetimes exist. Those who would not want to count such a spacetime as singular would be forced to give up b-incompleteness as the criterion for singular structure—which, given the lack of a clear physical interpretation of b-incompleteness in general,

---

65. I remark in passing that those disturbed by the prospect of an observer's having only a finite amount of proper time in which to exist into the future ought to be troubled by the Big Crunch, if there is to be one, but I have found no discussion of this point in the vast literature on singular structure, even by those relativists who display the germane intuitions. The thought seems to be that one ought to abhor singular structure in the 'interior' of spacetime, because one could imagine 'encountering' it on a walk through the park, so to speak. I do not think this a consistent stance, though.

Also, there seems to be a feeling among workers in the area that incomplete spacelike curves are not so bad in so far as they will have no observable effect on possible experiments one could perform in such a spacetime. I do not have space or time to go into it here, but I do not think this view is correct. For a sketch of the grounds for my reasons for saying so, see the account in (Synge 1960, ch. I, §14) of the physical significance of spacelike intervals.

66. See, *e.g.*, (Curiel 2009) and (Curiel 2000b).

67. To see how one might do this, consider solutions to Laplace's equation that exert along some particular, fixed direction a force constantly increasing as one moves along that direction. Perhaps one places positive charges at regular intervals along a geodesic starting from a given point, such that the magnitudes of the charges increase exponentially as one moves along the geodesic away from the point. If one then fires off a negatively charged particle essentially tangent to the given geodesic in the direction of increasing charge, the electric field of the charges will send it shooting off with an exponentially increasing acceleration in the direction orthogonal to the geodesic.

as opposed to incompleteness with respect to total proper time, I would not mind. [68] Really, so long as the idea of "reaching infinity" is given no precise content, and no argument is made to show why such a thing ought to alleviate anxiety about observers having only a finite amount of proper time in which to exist, there seems no reason to hold on to *b*-completeness. Of course, if incomplete timelike curves of unbounded total acceleration constituted singular structure, then every solution to Einstein's field equations would be singular. Many would reject this conclusion out of hand, but it does not seem intolerable to me. Singular structure would simply be one more type of global structure that all spacetimes necessarily had, along with, *e.g.*, paracompactness. Once so much was settled, then one could further classify spacetimes, according to the needs of the project at hand, by satisfaction of various more restrictive types of curve-completeness in order to produce more restricted, physically significant types of singular structure, as the compactness of a spacetime is a more restrictive type of paracompactness.

On physical grounds, curve incompleteness has been objected to because it seems to imply that particles could be "annihilated" or "created" right in the middle of spacetime, with no known physical force or mechanism capable of performing such a virtuosic feat of prestidigitation. [69] The demand that a spacetime be maximal, *i.e.*, have no proper extension, often rests on similar considerations: (Clarke 1975, pp. 65–6) and (Ellis and Schmidt 1977, p. 920) conjecture that maximality is required by the lack of a physical process that could cause spacetime to draw up short, as it were, and not continue on as it could have, were it to have an extension. This sort of argument, though, relies (implicitly) on a certain picture of physics that does not sit so comfortably with general relativity: that of the dynamical evolution of a system. From a certain quite natural point of view in general relativity, spacetime does not evolve at all. It just sits there, sufficient unto itself, very like the Parmenidean One. A solution to Einstein's equation, after all, is an entire spacetime *simpliciter*—and the topology may be naturally suggested by the form of the metric, as in the case of the Schwarzschild coordinates, but one can always put the 'same' metric on a space with an entirely different topology, and still have a solution to the field equations. From this point of view, the question of a physical mechanism capable of causing the spacetime manifold not to have all the points it could have had, as it were—which is essentially a topological question in the first place—becomes less poignant, perhaps even misleading. [70] Of course, an opponent of this point of view could argue that such a move could foreclose the possibility of deterministic physics, to which I would whole-heartedly agree, for we already know that general relativity does not guarantee deterministic physics: there may be no Cauchy surface in our spacetime, or there may even be so called naked

68. It is one of the few major shortcomings of (Earman 1995) that he does not analyze the physical significance of the various sorts of curve incompleteness.

69. *Cf.*, *e.g.*, (Hawking 1967, p. 189).

70. The invocation of problems arising from the principle of sufficient reason (if one thinks of these as problems at all!) in postulating maximality makes the same assumption: that a "creative force", in Earman's words, would create spacetime piece by piece, and not simply have it be there all at once, so to speak, in whatever form was desired. Under such a conception, one might wonder why the creative force would stop at any particular point and not continue on to 'complete' the spacetime. Such problems do not arise in the viewpoint I propound. Demanding maximality may lead to Buridan's Ass problems anyhow, for, as mentioned earlier, it can happen that global extensions exist in which one of a given set of incomplete curves is extendible, but no global extension exists in which every curve in the set is simultaneously extendible. Also, there may exist several physically quite different global extensions: the spacetime covered by the usual Schwarzschild coordinates for $r > 2M$, for instance, can be extended analytically to Kruskal-Schwarzschild spacetime, or it can be extended to a solution representing the interior of a massive spherical body. The three criteria usually invoked in choosing an extension are: analyticity (as in the Kruskal extension); preservation of a symmetry group (as in the interior Schwarzschild solution); limiting the Bondi news-function, *e.g.*, incoming radiation, in the extended spacetime. None seems very compelling. For example, why limit incoming radiation when relativity treats radiation as every bit as real as ponderable matter, in the sense of contributing to the stress-energy tensor and so to metric structure, irrespective of the presence of other fields or matter that may be considered sources of the radiation? Indeed, I know of no convincing, rigorous way to distinguish "radiation" from "ponderable matter" in general relativity.

singularities.[71]

Perhaps a more serious worry is that such a viewpoint would seem to deny that certain types of potentially observable physical phenomena require explanation, when on their face they would look puzzling, to say the least. Were we to witness particles popping in and out of existence, the mettle of physics surely would demand an explanation. I would contend in such a case, however, that a perfectly adequate explanation was at hand: we would be observing singular structure. If there were no curvature pathology around, such a response might appear to be ducking the real issue, *viz.*, why is there this anomalous singular structure when all our strongest intuitions and most dearly held metaphysical principles tell us it should be impossible?[72] Far from ducking the issue, the viewpoint I advocate is the only one I know of that gives us a toehold in looking for precise answers to such questions—or, more precisely, in making such questions precise in the first place. Note that those who balk at this viewpoint ought to be equally as troubled by the singular structure associated with the Big Bang as they are by the example under discussion, for it just as surely lacks an explanation. From the viewpoint I advocate, questions about what happened "before" the Big Bang, or why the universe "came into being", can come from their former nebulosity into sharper definition, for they become questions about the presence of certain global structure in the spacetime manifold, in principle no different from paracompactness, connectedness or the existence of an affine connection, and one can at least envisage possible forms of an answer to the (precise) question, "Are there any factors that necessitate spacetime's having such and such global structure?" And were we actually to observe particles popping in and out of existence, we could formulate and begin trying to answer the analogous questions.

The most serious problem I can imagine for the viewpoint I advocate is that of representing our subjective experience, experience that seems inextricably tied up with ideas of evolution and change. As I suggested earlier, this problem is not an idiosyncracy of the viewpoint I advocate, but in fact arises from the character of general relativity itself: 'dynamical evolution' and 'time' are subtle and problematic concepts in the theory no matter what viewpoint one takes, as attested by the most notorious and seemingly intractable problem in the drive to 'quantize' gravity, the so-called problem of time.[73] My viewpoint has the virtue of calling attention to this very fact, that, to judge by the preponderant mass of literature in both physics and philosophy, is often overlooked: general relativity, in its own way, requires us to refashion the conceptual apparatus we use to comprehend the physical world, to a rethink in a profound way several dearly held, deeply related concepts and the relations among them, just as quantum mechanics has.

It has become fashionable of late to say that such problems point to the need to find an "interpretation" of general relativity in the same sense in which the measurement problem in quantum mechanics is taken to require that that theory be interpreted. (Belot 1996), for instance, reaches this conclusion on the basis of an investigation into the problems encountered in trying to develop a quantum theory of gravity. I think this is a serious misunderstanding. Quantum mechanics demands an interpretation because it is not clear how to model physical phenomena, how to model the outcomes of experiments *simpliciter*: the predictions of standard quantum theory are in some sense in contradiction with the outcomes of experiments, but not in such a way as to invalidate the theory—an extraordinary state of affairs. There is no analogous problem in general relativity. We know how to model in the terms of the theory experiments that manifest and probe every phenomena suggested or predicted by the theory, with no inconsistency of any kind, for we know with no ambiguity what are the fundamental, physical terms and principles of the theory in which one articulates these models and draws conclusions on their basis. In a similar vein, the comprehension of special relativity's dismissal of the idea of absolute simultaneity did not require an interpretation

71. See (Earman 1995, ch. 3) for a discussion of these phenomena.

72. Speaking of which, I would love to have someone explain to me in a way that will not make me cringe what the difference is between a philosophical intuition and a metaphysical principle.

73. See (Ashtekar 1991, §12.3) for a brief discussion of this problem.

of the theory, in any sense of the term; it required only that investigators come to terms with the fact that the fundamental terms of the theory does not allow for the rigorous, physically relevant articulation of the fundamental terms of Newtonian physics. In quantum mechanics, we do not even know what the fundamental terms and principles—'measurement'? 'observable'?—ought to be.

In a paper on the foundations of quantum mechanics, discussing the lack in general relativity of an explicit representation of our experience of a privileged instant in our history, the "now", (Stein 1984, p. 645) makes a remark most *à propos* to the present case: "…although relativity does not give us a *representation* of that experience [, the psychologically privileged status of the "now"], there is no *incompatibility* between the experience and the theory: a gap is not a contradiction." [74] There is a gap between the raw materials the theory provides us and the rich content of our experience to be explained—but it is no flaw of or lacuna in general relativity—it is not in virtue of the lack of an "interpretation"—that the theory does not illuminate the psychological experience we imagine will accrue to an observer in any particular circumstance the theory predicts, no more than Newtonian mechanics fell short in so far as it did not show why I understand by certain irritations of my eardrum from perturbations in the ambient air pressure the import of the spoken word 'gap'. It cannot be an argument against general relativity that it predicts phenomena we find it difficult to envisage, when we also know perfectly well how to model experiments that manifest and probe the phenomena. On those grounds, I submit, every revolutionary physical theory ever proposed would have been DOA, in light of the historical evidence concerning the reception by contemporaneous scientists of every one of them. [75]

74. The italics are Stein's. The point Stein makes with this remark is somewhat different than the point I wish to extract from it, but they are akin enough for my profitable use of it.

75. I thank R. Geroch and D. Malament for stimulating conversations on all these topics. I'm also grateful to M. Dorato for writing a review of (Earman 1995) that made me realize the need to reread it and think more about singular structure, and to the History and Philosophy of Science Department at Pittsburgh, where I presented an earlier, briefer, version of this paper in a colloquium, for stimulating questions.

Chapter 4

On the Formal Consistency of Experiment and Theory in Physics

Abstract

The dispute over the viability of various theories of relativistic, dissipative fluids is analyzed. The focus of the dispute is identified as the question of determining what it means for a theory to be applicable to a given type of physical system under given conditions. The idea of a physical theory's *regime of applicability* is introduced, in an attempt to clarify the issue, along with a way of trying to make the idea precise. This construction involves a novel generalization of the idea of a field on spacetime, as well as a novel method of approximating the solutions to partial-differential equations on relativistic spacetimes in a way that tries to account for the peculiar needs of mathematical physics. It is argued, on the basis of these constructions, that the idea of a regime of applicability plays a central role in attempts to understand the relations between theoretical and experimental knowledge of the physical world in general, and in particular in attempts to explain what it may mean to claim that a physical theory models or represents a kind of physical system. This discussion necessitates an examination of the initial-value formulation of the partial-differential equations of mathematical physics, which suggests a natural set of conditions—by no means meant to be canonical or exhaustive—one may require a theory of mathematical physics satisfy in order to count as a physical theory. Based on the novel approximating methods developed for solving partial-differential equations on a relativistic spacetime by finite-difference methods, a technical result concerning a peculiar form of theoretical under-determination is proved, along with a technical result purporting to demonstrate a necessary condition for the self-consistency of a physical theory.

[[[*** 'colored' structures should rather be nominated 'piebald', 'pied', or 'mottled', 'motley', or 'stippled' ***]]

4.1 Introduction

In this paper, I intend to investigate a series of questions on the complex interplay between the theoretician and the experimentalist required for a mathematical theory to find application in modeling actual experiments and, in turn, for the results of those experiments to have bearing on the shaping and substantiation of a theory. On the one hand, we have the rigorous, exact and often beautiful mathematical structures of theoretical physics for the schematic representation of the possible states and courses of dynamical evolution of physical systems. [1] On the other hand, we have the intuitive, inexact and often profoundly insightful design and manipulation of experimental apparatus in the gathering of empirical data, in conjunction with the initial imposition of a classificatory structure on the mass of otherwise disaggregated and undifferentiated raw data gathered. Somewhere in between these extremes lie the mutual application to and qualification of each by the other.

It is one of the games of the experimentalist to decide what theory to play with, indeed, what parts of what theory to play with, in [*** planning experiments and designing instruments for them, designing the arrangments ***], modeling any particular experimental or observational arrangement, in light of, *inter alia*, the conditions under which the experiment will be performed or the observation made, the degree of accuracy expected or desired of the measurements, *etc.*, and then to infer in some way or other from the exact, rigorous structure of that theory, as provided by the theoretician, models of actual experiments so that he may explicate the properties of types of physical systems, produce predictions about the behavior of those types of systems in particular cirumstances, and judge whether or not these predictions, based on the schematic models contructed in the framework of the theory, conform to the inaccurately determined data he gathers from those experiments. It is one of the games of the theoretician to abduct exact, rigorous theories from the inaccurately determined, loosely organized mass of data provided by the experimentalist, and then to articulate the rules of play for those theories, by, *inter alia*, articulating the expected kinds and strengths of couplings the quantities of the theory manifest and the conditions under which they are manifested, leaving it to the experimentalist to design in light of this information probes of a sort appropriate to these couplings as manifested under the particular conditions of experiments. Jointly, the two try to find, in the physical world, common ground on which their games may be played. No matter what one thinks of the status of these sorts of decisions and articulations in science—whether one thinks they can ultimately be explained and justified in the terms of a rational scientific methodology or whether one thinks they are, in the end, immune to rational analysis and form the incorrigibly asystematic bed-rock of science, as it were—it behooves us, at the least, to get clearer on what is being decided and articulated, and on how those decisions and articulations bear on each other, if, indeed, they do at all .

I will not examine the actual play of the theoretician and the experimentalist in their attempts to find common, mutually fruitful ground on which to engage each other. I leave those issues, fascinating as they are, to other, more competent hands. Neither will I examine all the different sorts of games in which they engage in their respective practices, rather treating only those played in one small part of their common playground, that having to do with the comparison of predicted and observed

1. I follow the discussion of Stein (1994) here in my intended use of the term *schematic* to describe the way experiments are modeled in physics. That paper served as much of the inspiration for the questions I address in this paper, as well as for many of the ways I attempt to address the questions. Besides to that paper, I owe explicit debts of gratitude for inspiration to Geroch (2001), Stein (npub), Stein (1972) and Stein (2004), with all of which, I hope, this paper has affinities, in both method and conclusions.

values of a system as it dynamically evolves for the purposes of testing and substantiating a theory on the one hand, and refining experimental methods and design on the other [*** for this latter, *cf.* the suggestion by Lee and Yang of the experiments that showed violation of party; differentiate these more explicitly from using the predictions—though even in the example of the moon-shot, they surely used the actual observations of refine their methods of predicting approximating solutions, *etc.* ***]. I do not deal explicitly with others, such as predictions that have nothing to do with comparison to observations (for instance, the use of Newtonian gravity in calculating trajectories during the Apollo project's flights to the Moon), or the calculation of fundamental properties of physical systems based on theoretical models (for instance, the use of the quantum theory of solids to calculate the specific heat of a substance). The extension of the methods and arguments of this paper to those and other practices strikes me as straightforward, but the proof is in the pudding, which I do not serve here, and have, indeed, not thought much about preparing.

[*** Move this paragraph elsewhere ***]. Indeed, in my use of the terms 'experimentalist' and 'theoretician' throughout this paper, I am guilty of perpetuating the crudest of caricatures—as though the two lived in separate worlds, and had to travel some distance and overcome great obstacles even to meet each other. Physicists such as Newton and Fermi, masters of both theory and experiment, give my caricature the lie. I think Bob Geroch, unwittingly to be sure, exemplified the distinction I am trying to gesture at when, in lectures or conversation, he would declare that he was currently wearing his "physics-hat" or that he was now changing into his "math-hat", to emphasize the spirit with which his listeners should understand what he was saying. [2] Still, there is a grain of truth in the caricature—which is to say, it does strike home somewhat. When I had a desk in the Relativity Group, in the Fermi Institute at the University of Chicago, the other graduate students and I used to say, only half jokingly, that other groups of theoreticians—those studying quantum field theory, or solid-state physics, for example—spoke a different language than the one we spoke, and one had to work hard at translation if we were to avoid a complete breakdown of communication. One may extrapolate from there our feelings about experimentalists: "if experimentalists could speak, we would not understand them." [3]

I will examine in this paper only what one may think of as the logical structure of the relations between the practice of the theoretician and that of the experimentalist, and, *a fortiori*, of those between theory and experiment. I do not mean to claim that there is or ought to be a single such structure *sub specie æternitatis*, or indeed that there is any such structure common to different branches of physics, or indeed even one common to a single branch that remains stable and viable over arbitrary periods of time, in different stages of the scientific enterprise. I intend to investigate only whether one can construct such a structure to represent some idealized form of these relations. I am not, in this paper, interested in how exactly the experimentalist and the theoretician may make in practice the transitions to and fro between, on the one hand, inaccurate and finitely determined measurements, and, on the other, the mathematically rigorous initial-value formulation of a system of partial-differential equations, whether their exact methods of doing so may be justified, *etc.* I am rather concerned with the brute fact of its happening, whether there is indeed any way at all of constructing with some rigor and clarity a model of generic methods for doing so. Having such a model in hand would show that there need be no gross logical or methodological inconsistency in their joint practice (even if there is an inconsistency in the way physicists currently work, which I would not pretend to hazard a guess at). Indeed, it is difficult to see, on the face of it, how one may comprehend these two to be engaged in the same enterprise in the first place, difficult, indeed, to see even whether these two practices are in *any* sense consistent with each other, since it is not even clear what such consistency may or may not consist of. While I seriously doubt that any formal

2. Probably my favorite example was when he wrote scare-quotes on the wall in chalk on both sides of a chalkboard covered in mathematical formulæ, to emphasize the fact that we should understand it to be "physics-math", as it were, and not "mathematics-math".

3. To which the experimentalist replied, "When you speak, I cannot understand you!"

analysis of the relations between theory and practice I or anyone else may propose could answer this question definitively with regard to a real physical theory and its experimental applications, the sort of analysis I attempt to outline here, if successful, would perhaps have the virtue of underlining the sorts of considerations one must take account of in judging the consistency of a real theory. This may seem a Quixotic project, at best, on the face of it, but I think I can say a few words in defense of its interest in the remainder of the introduction. In defense of its feasibility, I offer the paper itself.

Without a doubt, one can learn an extraordinary amount about a physical theory (and about the world) by examining only its structure in isolation from the conditions required for its use in modeling phenomena, as is most often done in philosophical discussion of a technical nature about physical theories in particular, and about the character of our understanding of the physical world in general. I will argue, however, [*** I don't in fact give such an argument ***] that comprehensive understanding of a physical theory will elude us unless we examine as well the procedures whereby it is employed in the laboratory, and, moreover, that comprehension of the nature of such knowledge as we may have of the physical world will similarly elude us without a serious attempt to understand both the theoretical and the practical characters of that knowledge. In particular, the question I plan to address is not how one gets to a system of exact partial-differential equations from inaccurate data; nor is it how one gets from exact solutions of partial-differential equations to predictions that may or may not accord with actually observed, inaccurate data (though this latter will be touched upon *en passant* to some degree). It is rather a question of the consistency of, perhaps the continuity between, the two—a question, if you like, of whether the theoretician and the experimentalist can be understood as being engaged in the *same enterprise*,[4] that of modeling and comprehending the physical world, in complementary, indeed mutually supportive, ways. The answer I propose is constructive—a proposal for a more or less formal, explicit method of representing the connection between the stocks in trade of the two that remains true to the character of these two stocks. Another way of putting the point [*** no, it's a different point ***]: philosophers, when having tried to understand the relation between theory and experiment, tend to have been vexed by the problem of how a theory gets into (and out of!) the laboratory, often framed in terms of the putatively inevitable "theory-ladenness" of observations; I am concerned here with what one may call the converse problem, that of getting the laboratory into the theory, and the joint problem, as it were, whether the theory and the laboratory admit at least in part a consistent, common model. Along the way, I will present an argument, in large part constituted by the body of the construction itself, that the initial-value formulation of the partial-differential equations of a theory provides the most natural theater in which this sort of investigation can play itself out. Later in the paper, after the construction has been sketched, I will have more to say explicitly on the privilege, as I see it, accruing to the role of the initial-value formulation in the comprehension of physical theory.

I will focus the discussion around the idea of the *regime of applicability*[5] of a physical theory (or *physical regime* or just *regime*, for short). From a purely extensive point of view, a regime of a physical theory, roughly speaking, consists of the class of all physical systems *cum* environments that the theory is adequate and appropriate for the modeling of,[6] along with a mathematical structure used to construct models of these systems, and a set of experimental techniques used for probing the systems in a way amenable to modeling in the terms of that structure. It can be represented by, at a minimum [*** the difference between kinematics and dynamics: there is a set of (algebraic and differential) conditions, stated only in terms of variables representing environmental quantities, characterizing part of the idea of a regime, none of which appear in the dynamical partial-differential equations of the theory; whereas the variables representing the physical quantities treated by the theory of the

4. I choose the word *enterprise* with some care, as an explicit reference to Stein (2004).

5. I owe this term, as well as the germ of the idea for the concept, to Geroch (2001).

6. It is immaterial to my arguments whether one considers the class to comprise only existant systems or to comprise as well possible systems, in whatever sense one wishes to give the modal term.

physical system at issue appear only in equations and conditions where the environmental variables may also appear, coupled to the system's dynamical variables ***]

1. a set of variables representing physical quantities not directly treated by the theory but whose values in a given neighborhood are relevant to the issue of the theory's applicability to a particular physical system in that neighborhood, along with a set of algebraic and differential expressions formulated in terms of these variables, representing the constraints these ambient, environmental quantities must satisfy in order for physical systems of the given type to be susceptible to treatment by the theory when they appear in such environments
2. a set of algebraic and differential expressions formulated in terms of the variables and constants appearing in the theory's system of partial-differential equations, representing the constraints the values of the quantities represented by those constants and variables must satisfy in order for the system bearing those quantities to be amenable to treatment by the theory; these expressions may include as well terms from the set of variables representing relevant environmental quantities
3. a set of algebraic expressions formulated in terms of variables representing the measure of spatiotemporal intervals, constraining the character of the spatiotemporal regions requisite for well-defined observations of the system's quantities to be performed in; these expressions may include terms from the set of variables representing relevant environment quantities, as well as from the set of variables and constants appearing in the theory's system of partial-differential equations
4. a set of methods for calculating the ranges of inaccuracy inevitably accruing to measurements of the values of the system's quantities treated by the theory, depending on the sorts of experimental techniques used for probing the system, the environmental conditions under which the probing is performed, and the state of the system itself (including the stage of dynamical evolution it manifests) at the time of the probing—these methods may include, *e.g.*, a set of algebraic and differential expressions formulated in terms of the variables and constants appearing in the theory's system of partial-differential equations, the variables representing the relevant environmental factors, and the variables representing the measure of spatiotemporal intervals
5. a set of methods for calculating the ranges of admissible deviance of the predictions of the theory on the one hand from actual measurements made of particular systems modeled by the theory on the other, depending on the sorts of experimental techniques used for probing the system, the environmental conditions under which the probing is performed, and the state of the system itself (including the stage of dynamical evolution it manifests) at the time of the probing—these methods may include, *e.g.*, a set of algebraic and differential conditions formulated in terms of the variables and constants appearing in the theory's system of partial-differential equations, the variables representing the relevant environmental factors, and the variables representing the measure of spatiotemporal intervals

The idea of a regime is perhaps best illustrated by way of an example. For the theory comprising the classical Navier-Stokes equations to model adequately a particular body of fluid, for instance, elements of its regime may include these conditions and posits:

1. the ambient electromagnetic field cannot be so strong as to ionize the fluid completely
2. the gradient of the fluid's temperature cannot be too steep near equilibrium
3. only thermometric systems one centimeter in length or longer are to be used to measure the fluid's temperature, and the reading will be taken only after having waited a few seconds for the systems to have settled down to equilibrium

4. the chosen observational techniques to be applied, under the given environmental conditions and in light of the current state of the fluid, yield data with a range of inaccuracy of $\pm 1\%$, with a degree of confidence of 95%
5. a deviance of less than 3% of the predicted from the observed dynamic evolution of the system's temperature, taking into account the range of inaccuracy in measurement, is within the admissible range of experimental error for the chosen experimental techniques under the given environmental conditions, in light of the current state of the fluid

I neither promise nor threaten to offer in this paper a definitive analysis of the concept of a regime or indeed of any of its constituents. I will rather sketch one possible way one may construct a (moderately) precise and rigorous model of the concept, with the aim of illuminating the sorts of questions one would have to answer in order to provide a more definitive analysis. The hope is that such a model and correlative demonstration may serve as a contructive proof of the formal consistency of the practice of the experimentalist and the practice of the theoretician in physics, at least in so far as one accepts the viability of the sort of formal model I will construct, indeed, as a construction of the common playground, as it were, of the two, playing with the toys and rides and games of which we may pose precise questions of a technical nature about the interplay between theory and experiment, and attempt to answer such questions.

The structure of the paper is as follows. In order to illuminate the sort of considerations that motivate and found my proposed definition and analysis of the idea of a regime, I begin, in §4.2, by briefly analyzing the dispute over hyperbolic reformulations of the theory of relativistic Navier-Stokes fluids, as the dispute illuminates the issues nicely. The points drawn from this analysis lead naturally into the introduction in §4.3 of the notion of a regime, and the sketch of a construction of a formal model purporting to represent the notion. In §4.4, I offer a mildly technical analysis of the mathematical representation appropriate for the modeling of physical fields by theories with regimes, necessary for the culmination of the paper in §4.5, in which I analyze the initial-value formulation of the partial-differential equations of theoretical physics (as opposed to that of those in pure mathematics) based on my analysis of the idea a regime, and draw several consequences from the analysis, and in §4.6, in which I discuss the criteria one may want to demand a theory satisfy in order for it to be thought empirically adequate. One of the most interesting of the results of this discussion describes a peculiar form of theoretical under-determination necessarily attendant on a physical theory, in so far as the theory possesses a regime in the idealized sense proposed in this paper.

For the most part, I will deal only with the case of the interaction of theory and experiment when both the theoretical structures and the experimental practices are well worked out and well understood; the investigation of these relations when one is dealing with novel theory, novel experiments, or both, presents far too many difficult and unavoidable questions for me to treat with any adequacy or depth here.

The entire paper, if you will, may be considered an exercise in approximation and idealization in the philosophy of physics.

4.2 Relativistic Formulations of the Navier-Stokes Equations

4.2.1 Parabolic Theories and Their Problems

It is sometimes held that parabolic systems of partial-differential equations, such as the Navier-Stokes system or Fourier's equation of thermal diffusion, do not have well set initial-value formulations.[7] This, of course, depends on one's formulation of the idea of an initial-value formulation.

7. See, *e.g.*, Geroch (2001).

The following is known, for example, about the Navier-Stokes system in non-relativistic physics: [8]

1. for appropriate initial data on a 3-space of absolute simultaneity, say $t = 0$, there exists a $0 < \tau < \infty$ such that there is a unique, regular solution in the interval $[0, \tau)$
2. for appropriate initial data, a distributional solution exists for all future time and, in the two-dimensional case at least, this solution is unique

Whether global distributional solutions for the three-dimensional case are unique is apparently not known. Leray (1934) conjectured that global uniqueness does not hold (though, of course, he did not phrase this conjecture in the language of distributions), arguing that the break-down in uniqueness is associated with the onset of turbulence, which, he held, is not representable by the Navier-Stokes equations. Recently, Ruelle (1981) has attempted to argue that global uniqueness does hold, and that the onset of turbulence should rather be associated with the existence of strange attractors in the phase space of the Navier-Stokes system. So far as I know, no firm conclusions either way are known, and, in any event, this issue has not been treated in the context of relativity to any comparable depth.

What is indisputable is that, in parabolic systems, roughly speaking, although the solutions to boundary-value problems vary continuously with the specified boundary-values, perturbations in initial conditions can propagate with unbounded velocities in initial-value problems. In other words, there is no guarantee that the solutions to parabolic systems will not violate the causal strictures of relativity theory, no matter how exactly one poses those strictures. This observation underlies the sense theoretical physicists have of the inadequacy of parabolic partial-differential equations. Another, related problem with them involves the stability of their solutions. In the particular case of the relativistic Navier-Stokes system as formulated by Landau and Lifschitz (1975), for example, Hiscock and Lindblom (1985) found a solution that grows exponentially over microscopic time-scales in any coordinate system in which the representation of the net momentum-flux of the fluid is not zero. [9] Kostädt and Liu (2000) have disputed the admissibility of this solution, claiming that it arises from an ill-set initial-value formulation. They conclude that this parabolic formulation is in fact viable as a mathematical representation of a physical theory, at least in so far as such objections go.

These discussions and arguments are exemplary of the problems faced by theoreticians when attempting to model novel systems, or systems that can be investigated only with great difficulty. Of particular relevance for our study is the focus on whether or not the initial-value formulation of the partial-differential equations of a theory is well set or not. This notion will play an indispensable role in the characterization we offer, in §4.5, of a regime and of a theory possessing one.

4.2.2 Hyperbolic Theories

While the problems mentioned in §4.2.1 have served as stimulus for finding a hyperbolic extension of the relativistic Navier-Stokes system, in the attempt it was realized that there are two other perhaps even stronger reasons to find a viable such extension. In particular, it was suggested that, contrary to early assumptions, the hyperbolic theories might produce predictions differing from those of the parabolic system in certain tightly constrained circumstances in which both were applicable (the so-called *hydrodynamic regime*—more on this below), offering the possibility of experimental tests of the hyperbolic systems. [10] Even more enticingly, it was suggested that the hyperbolic

8. See Temam (1983, *passim*).

9. Their solution has its origin in the fact that Landau and Lifschitz (1975) define the mean fluid velocity by the net momentum-flux—the so-called *kinematic velocity*—rather than by the flux of the particle-number density—the *dynamic velocity*. Whereas in classical physics these two quantities are identical, this is not generically the case in relativity, though it may be in particular cases, such as for a system in complete thermodynamic equilibrium.

10. Of course, given the profound observational entrenchment of the parabolic Navier-Stokes system, one of the conditions demanded of such hyperbolic extensions will be that they (more or less exactly) recapitulate the predictions

theories could be applied in circumstances in which the parabolic system becomes in one way or another inapplicable. I will briefly discuss how it is hoped that the hyperbolic systems may resolve the problems mentioned in §4.2.1, but my primary focus for the majority of the section will be on the two novel suggestions just mentioned. [11]

In order to discuss these issues further, it will be convenient to be more precise than there has yet been call for. Fix a relativistic spacetime [12] $(\mathcal{M}, g_{ab})$. Then a *relativistic Navier-Stokes fluid* (or just *Navier-Stokes fluid*, when there is no ambiguity) is a physical system such that:

1. its local state is completely characterized by the set of dynamical variables representing the mass density ρ, particle-number density ν, mean fluid velocity [13] ξ^a, heat flow q^a, and shear-stress tensor σ_{ab}, jointly satisfying the four kinematic constraints

$$\sigma_{ab} = \sigma_{(ab)} \tag{4.2.2.1}$$

$$\xi^m \sigma_{am} = \xi^n q_n = 0 \tag{4.2.2.2}$$

$$\nabla_m(\nu \xi^m) = 0 \tag{4.2.2.3}$$

$$\nabla_m((\rho + p)\xi^a \xi^m + p g^{am} + 2q^{(a}\xi^{m)} + \sigma^{am}) = 0 \tag{4.2.2.4}$$

2. in the same physical regime, there are equations of state (specified once and for all), expressed in terms of the dynamical variables characterizing the state, defining the pressure p, temperature τ, thermal conductivity θ, shear-viscosity α and bulk-viscosity β

3. in the same physical regime, these quantities jointly satisfy the two equations of dynamic evolution

$$q_a + \theta\rho[(\delta^m{}_a + \xi_a \xi^m)(\nabla_m \log \tau) + \xi^n \nabla_n \xi_a] = 0 \tag{4.2.2.5}$$

$$\sigma_{ab} + \alpha\rho(\delta^m{}_{(a} + \xi_{(a}\xi^m)\nabla_{|m|}\xi_{b)}) + \rho(\beta - \alpha/3)(g_{ab} + \xi_a \xi_b)\nabla_n \xi^n = 0 \tag{4.2.2.6}$$

Equation (4.2.2.3) represents the conservation of particle number (all possible quantum effects are being ignored), (4.2.2.4) the conservation of mass-energy, (4.2.2.5) the flow of heat, and (4.2.2.6) the effects of viscosity and stress.

4.2.3 The Breakdown of Partial-Differential Equations as Models in Physics

Classically, every Navier-Stokes fluid has a characteristic length (or equivalently, characteristic interval of time), the *hydrodynamic scale*, below which the description provided by the terms of the theory breaks down. Typically there is only one such length, of the order of magnitude of the mean free-path of the fluid's particles; at this length scale, the thermodynamic quantities appearing in the equations are no longer unambiguously defined. Different sorts of thermometers, *e.g.*, with sensitivities below the hydrodynamic scale, will record markedly different "temperatures" depending on characteristics of the joint system that one can safely ignore at larger scales—the transparency

of the original system under appropriate conditions.

11. For arguments in support of both suggestions, see Müller and Ruggeri (1993a), Herrera and Martínez (1997), Anile, Pavón, and Romano (1998), Herrera, Prisco, and Martínez (1998), Herrera and Pavón (2001a), Herrera and Pavón (2001b) and Jou, Casas-Vázquez, and Lebon (2001), *et al.* For attempts actually to conduct such studies, see, *e.g.*, Müller and Ruggeri (1993b), Zimdahl, Pavón, and Maartens (1996), Herrera and Martínez (1998), Jou, Casas-Vázquez, and Lebon (2001), Eu (2002) and Herrera, Prisco, Martín, Ospino, Santos, and Troconis (2004).

12. For the purposes of this paper, a spacetime is a paracompact, Hausdorff, connected, orientable, smooth differential manifold endowed with a smooth Lorentz metric under which the manifold is also time-orientable. The imposition of temporal orientability a simplified presentation of the material dealing with the dynamic evolution of systems. It could be foregone by restricting all analysis to appropriate subsets of spacetime.

13. I leave it purposely ambiguous as to which definition of fluid velocity appears, the so-called kinematic or dynamic, as nothing hinges on it here.

of each thermometric system to the fluid's particles, for instance. The other quantities fail in similar ways.[14]

It is worth remarking that there is no *a priori* reason why the definitions of all the different quantities that appear in the Navier-Stokes system should fail at the same characteristic length, even though, in fact, those of all known examples do. Clearly, if there were different characteristic lengths at which the definitions of different quantities in the system broke down, the system itself would fail at the greatest such length-scale. Any phenomena that are observed at scales greater than the largest length at which one of the thermodynamic quantities becomes ill-defined are said to belong to the *hydrodynamic regime*; any observed below that scale are part of the *sub-hydrodynamic regime* (or *the regime of molecular effects*).[15]

As Geroch (2001) points out, the Navier-Stokes system can fail in another way, at a length-scale logically unrelated to the hydrodynamic length-scale, one at which the equations themselves may fail to hold even though all the system's associated quantities remain well-defined. In other words, there may be a characteristic length-scale at which the expressions on the left-hand sides of the equations (especially the last two) may differ from zero by an amount, *e.g.*, of the same order as that of the terms appearing in the equations, while the equations remain valid at scales greater than that length-scale. I will refer to such a length-scale as the *transient scale*, gesturing at the fact that it is reasonable to expect that any such failures would have their origins in the dissipative fluxes of the fluid's quantities transiently settling down as the quantities themselves approach their equilibrated, hydrodynamic values. This idea, in fact, inspires the preferred interpretations for the novel terms introduced in the hyperbolic theories. I will refer to the greatest length-scale at which the system for any reason is no longer valid—whether because the quantities lose definition or because the equations no longer hold—as the *break-down scale*, and I will refer correlatively to the regime below this scale as the *break-down regime*.

Geroch (2001) points out a possible complication in the notion of a characteristic length-scale at which the system of equations breaks down (for whatever reason). The system may fail in a way more complicated than can be described by a single, simple spatial or temporal length, or spatiotemporal interval. As an example, he points out that relativity itself imposes constraints on the experimental applications of the theoretical model: the model must fail at every combined temporal and spatial scale, t and s respectively, jointly satisfying

$$s^2 \lesssim \chi t \tag{4.2.3.1}$$

and

$$s > ct \tag{4.2.3.2}$$

where χ is the value of a typical dissipation coefficient for the fluid and c is the speed of light. Instead of a characteristic break-down scale, this requirement defines a characteristic break-down *area* in the t, s-plane. Note that the complement of this region of the plane, that is, the region in which the system remains valid (at least so far as these conditions are concerned), includes arbitrarily small s-values and arbitrarily small t-values (though not both at the same time!).

In this terminology, proponents of hyperbolic theories contend that the examples they exhibit are of relativistic, dissipative fluids for which the parabolic system adequately models the equilibrium behavior, yet which have transient scales measurably greater than their hydrodynamic scales, manifesting them in disequilibrated states—in other words, in certain kinds of disequilibrium, the

14. It may seem that this sort of constraint on the definition of physical quantities manifests itself only as one shrinks the germane spatial and temporal scales, but this is not so. Imagine the difficulties involved in attempting to define what one means by *the* temperature of a cloud of gas three billion light years across. How will one, for instance, calibrate the various parts of the thermometric apparatus, each with the others?

15. I have also seen this in several places referred to as the *Knudsen regime*, but no one ever gives a citation, explains why or says who this Knudsen is.

quantities in the equations are well-defined, but the equations themselves fail to hold to a degree that, for one reason or another, whether theoretical, experimental or pragmatic, is unacceptable. Geroch (2001), in turn, contends that there are no such fluids not even, as he puts it, any known *gedanken* fluids.[16] This is why such fluids represent an intriguing possibility: they would provide unambiguous examples of systems (presumably) amenable to theoretical treatment by the hyperbolic theories and (perhaps) accessible to experimental investigation.

Geroch (2001) offers an illuminating example of a particular way a system of equations may fail while the quantities in terms of which the equations are formulated remain well defined. I call it the *problem of truncation*, and, again, the hyperbolizations of the relativistic Navier-Stokes equations provide excellent illustrations. The hyperbolizations work, as we have said, by introducing terms of second-order or higher,[17] purportedly representing transient fluxes of the ordinary quantities treated by the parabolic Navier-Stokes system. There is, however, no natural, *a priori* way to truncate the order of terms one would have to include in the new equations to model the systems accurately enough, once one began including any higher-order terms, for the scales at which second-order effects become important, for instance, seem likely to be the same at which third-order, fourth-order and 839$^{\text{th}}$-order terms also may show themselves. It is, so far as I can see, a miracle and nothing more that there are physical systems capable of being accurately modeled by the first-order Navier-Stokes equations, ignoring all higher-order effects. As Geroch says,

> The Navier-Stokes system, in other words, has a "regime of applicability"—a limiting circumstance in which the effects included within that system remain prominent while the effects not included become vanishingly small.
> Geroch (2001, pp. 6–7)

The quantities modeled by the parabolic Navier-Stokes equations have a regime in which they are simultaneously well-defined, satisfy the equations and have values stable with respect to higher-order fluctuations. One cannot assume this for any amended equations one writes down, with novel terms purportedly representing higher-order effects. One must demonstrate it. On the face of it, this would be a fool's quest to attempt by a strictly theoretical analysis; in practice, I suspect, it could be accomplished only through experimentation, and this itself would be a difficult task, at best.

4.3 The Regime of a Physical Theory

Philosophical analysis of particular physical theories, such as non-relativistic quantum mechanics, often focuses on the more or less rigorous mathematical consequences of the structure of the theory itself, in abstraction from the necessary laboratory conditions required for application of the theory in modeling the dynamic evolution of particular, actual systems. To clarify what I mean, consider the usual schema of a Bell-type experiment considered by philosophers: an undifferentiated source of pairs of electrons in the singlet state, and an inarticulate, featureless Stern-Gerlach device to measure the spin of the electrons. This indeed constitutes a model of a physical system, but only in an abstract, even *recherchè*, sense. No consideration is given to the structure of the source of the electrons, the exact form of the coupling between the system under investigation (the pairs of electrons) and the instrument used to measure the relevant quantities of the system (the Stern-Gerlach device), or to the regime of applicability of the model they are using for this kind of system's coupling with that sort of measuring apparatus—it is a schematic representation of the experiment, in the most rarefied sense of the term.

It is taken for granted, for instance, that

16. He does exhibit an instructive but ultimately unsuccessful attempt to construct an example of such a Navier-Stokes fluid.

17. The *order* of a term here refers roughly to the moment of the distribution function one must calculate to express it.

1. the ambient temperature is not so high or so low as to disrupt the source's output of the electrons
2. the electrons are not traveling so quickly (some appreciable fraction of the speed of light), nor are the primary frequencies of the photons composing the magnetic field so high (having, *e.g.*, wave-lengths of the order of the Compton length), as to require the use of quantum field theory rather than standard non-relativistic quantum mechanics in order to model the observation appropriately
3. the spins of the electron are measured using a Stern-Gerlach type of mechanism whose physical dimensions are such as to allow its being treated as a classical device (as opposed to one whose dimensions are of such an order—a quantum-dot device, *e.g.*—for which the "measurement of the electron's spin" would become ambiguous, as one would have to account for the quantum properties of the measuring device as well in modeling the interaction)
4. the metric curvature of the region in which the experiment is being performed ought not be so great as to introduce ambiguity in the assignment of correlations among the spin-components of different directions at the different spacetime points where the spin of each electron is, respectively, measured [18]

In the literature in general, no effort is put into determining how such restrictions may, if at all, affect the expected outcome of the experiments. [19] While we are perhaps safe in blithely ignoring these sorts of issues in the case of Bell-type experiments (and I am not even convinced of that), the study of theories of relativistic, dissipative fluids provides a clear example of a case in which we may not safely ignore them.

The analysis of the debate over theories of a relativistic Navier-Stokes fluid shows that, at a minimum, the applicability of a theory to a set of phenomena is constrained by conditions on the values of environmental quantities, the values of the quantities appearing in the theory's equations, and the measure of spatial and temporal intervals: a theory can be used to treat a type of physical system it putatively represents only when the system's environment permits the determination, within the fineness and ranges allowed by their nature, of the system's quantities over the spatial and temporal scales appropriate for the representation of the envisioned phenomena. In this section, I will propose a possible model for dealing with these considerations precisely, *the kinematical regime*, requiring with regard to the observation and measurement, and hence to the well-definedness, of the quantities treated by a theory [*** explicate what sense *kinematical* has in this context—definition of quantities, and what else? ***]:

1. a set of constraints on the measure of spatial and temporal intervals, and perhaps as well on the behavior of the metric in general (*e.g.*, that some scalar curvature remain bounded by a given amount in the region)
2. a set of constraints on the values of the theory's quantities in conjunction with correlative constraints on environmental conditions
3. a set of methods for calculating the ranges of inevitable inaccuracy in the preparation or measurement of those quantities using particular sorts of experimental techniques under particular environmental conditions

18. If the curvature were so great that parallel transport of a tangent vector along different paths from the point of measurement of the spin of one of the electrons to the point of that of the other would yield markedly different resultant tangent vectors, then the question of the correlation of the spins along "opposite" directions at the two points becomes incoherent.

19. The analysis of Fine (1982) perhaps comes the closest in spirit in trying to take account of these sorts of issues with regard to Bell-type experiments. The discussion of Stein (1972) presents a much richer, somewhat complementary account to the one I sketch here.

I will not attempt to formulate any of the notions I discuss with rigor or to treat them to any depth. For those interested in the development of a rigorous technical apparatus for treating all these issues, as well as for treating the issues raised in §§4.4 and 4.5, see Curiel (2011).

Before starting the analysis proper, we fix some definitions.[20] Given a type of physical system, a *quantity* of it is a (possibly variable) magnitude that can be thought of as belonging to the system, in so far as it can be measured (at least in principle) by an experimental apparatus designed to interact with that type of system, in a fashion conforming to a particular coupling of the system with determinable features of its environment, which coupling may (at least, again, in principle) be modeled theoretically.[21] Fix, then, a type of physical system, along with a system of partial-differential equations. An *interpretation* of the equations in terms of the quantities of that type of physical system (or, more briefly: in terms of the physical system itself) is a complete, one-to-one correspondence between the set of variables and constants appearing in the equations on the one hand and some sub-set of the known quantities of that type of physical system on the other, in conjunction with a set of statements describing the coupling of these quantities to known and determinable features of the environment precise and detailed enough to direct the experimentalist in constructing probes and intruments tailored to the character of each quantity, as associated with that kind of system, for its observation and measurement. The system of partial-differential equations *models* the type of physical system if, given an interpretation of the equations in terms of that type of physical system, and given any appropriate set of initial data for the equations representing a possible state of a physical system of that type, the mathematically evolving solution of the equation continues to represent a possible state of that system if it were to have dynamically evolved from a state represented by the initial data of the equations. In other words, the equations model the system if the equations' solutions do not violate any of the system's inherent kinematic constraints. If, for example, a set of partial-differential equations as interpreted by the terms of a given type of physical system predicted that systems of that type, starting from otherwise acceptable initial data, would evolve to have negative mass, or would evolve in such a way that the system's worldline would change from being a timelike to being a spacelike curve, then we would likely conclude that those partial-differential equations do not model that type of system. Note, in particular, that modeling is a strictly kinematical notion. The accuracy of predictions produced by the partial-differential equations—whether or not its solutions, under the given interpretation match to an

20. None of these prefatory definitions ought to be considered attempts at even the slovenly rigor, as it were, I aim for in this paper, or, indeed, anything near it. These are rather in the way of marking off the field of play, much as children determine a bit of a meadow as a soccer-field with episodic markers of the boundary (jackets, frisbees, ...), which is to be interpolated between those markers as the niceness of the occasion demands.

21. This characterization of quantity involves (at least) one serious over-simplification. Not all quantities' values can be determined by direct preparation or measurement, even in principle, as this statement may suggest. Some, such as that of entropy, can only be calculated from those of others that are themselves directly preparable or measurable. Other quantities defy direct measurement for all intents and purposes, though perhaps not strictly in principle. Consider, for example, the attempt to measure directly distances of the order of 10^{-50} cm—the precision required of any measuring device that would attempt it would demand that the probes it uses have de Broglie wave-lengths of comparable scale, and so, correlatively, would demand the release of catastrophic amounts of energy in its interaction with another system—think of the energy of a photon whose wavelength was of that scale.

I am not sure whether the analysis I offer in this paper would or would not suffice for the treatment of these sorts of quantity, though, offhand, I see no reason why it should make a difference. Temporal and spatial constraints do not allow me to consider the question here, however.

[*** Note also that this is not an *instrumentalist* requirement. These measurements do not define the quantities, at least not in all cases. ***]

Perhaps this is not an inappropriate place to mention, *en passant*, that, were one to allow oneself the momentary luxury of Saturnalic speculation and wild extrapolation, it would be fun to imagine that the lack of such a thing, even in principle, as an "entropometer", and correlatively the lack of a unit of measure or scale for entropy, as the Joule is to energy, somehow has to do with the fact that the Wigner time-reversal operator in quantum mechanics is not a Hermitian operator but rather is anti-Hermitian. (The issue of the possible existence of an entropometer is, thank goodness, not directly related to the possibility of a Maxwell demon, for the demon does not putatively measure the entropy of a system, but rather only reduces it piece-meal.)

admissible degree of accuracy the actual, dynamic evolution of such systems—has no bearing on the question of modeling. Let us say, then, that a physical theory *comprises* a system of partial-differential equations if those equations model the types of systems treated by the theory, under the interpretation the theory provides. For example, the theory of relativistic Navier-Stokes fluids comprises equations (4.2.2.3)–(4.2.2.6), under their standard interpretation. Finally, by *physical theory*, I intend, very roughly speaking, an ordered set consisting of, at least,

1. a mathematical structure representing the states and the dynamical evolution of the physical systems treated by the theory
2. a set of experimental techniques for probing these systems
3. a mapping between the terms of the mathematical structure and the quantities associated with these systems as observed and probed by the experimental techniques
4. the set of data germane to knowledge of these quantities, collected from these systems by the given experimental techniques and analyzed and informed by application of the given mathematical structure

The last element I would include in the ordered set is a regime of applicability for the theory, to the articulation of which I now turn.

I feel I need to make one last remark before proceeding, however. One may be tempted to think that a "fundamental" physical theory, such as quantum field theory, ought not require specification of a regime for its applicability. This is not the case. Quantum field theory can not solve in closed form the dynamical equations representing the evolution of arguably even the simplest micro-system, the isolated Hydrogen atom. It rather relies on perturbative expansions, and thus requires the system to be not too far from equilibrium of one sort or another. Quantum field theory in general, moreover, can not handle phenomena occurring in regions of spacetime in which the curvature is too large. The Standard Model breaks down in regimes far above the Planck scale. Not even quantum field theory formulated on curved-spacetime backgrounds can deal rigorously with phenomena under such conditions.[22] Indeed, it appears that possession of a fairly well articulated regime of applicability, as we will characterize it, or something nearly like it, is necessary for a theory's being viable as a theory of *physics*, as opposed to merely a chapter of pure mathematics.[23] Bondi, in a paper on gravitational energy, puts his finger on the heart of the issue: "*Good physics is potential engineering.*"[24]

4.3.1 Constraints on the Measure of Spatiotemporal Intervals

The idea of a regime, at bottom, rests on twin pillars: the idea that certain types of operations associated with the theory make sense (in some fashion or other) only when carried out over spatiotemporal regions whose dimensions satisfy certain constraints and in which some appropriate measure of the intensity of the metric curvature does not become too great; and that certain types of operations associated with the theory make sense (in some fashion or other) only when the values

22. I know of no theory of quantum gravity mature enough for it even to attempt the claim that it could do so. Even if one could, and even were we able to observe and measure the peculiar quantities modeled by the theory, we presumably would measure them using technological apparatus of some stripe, which, again presumably, would be limited in its precision and its accuracy.

As an aside, I remark that I may appear to be leaving myself open to the charge of conflating two different ways in which inaccuracy can accrue to measurments and predictions, one based on the nature of the quantities (as with the statistical character of temperature, *e.g.*, or as in the constraints imposed by the Heisenberg principle on quantum phenomena), and the other based on *de facto* limitations due to the current stage of development of our technological prowess. I should rather say that, part of the point of this paper, is that this distinction may not be so sharp and clean as it appears at first glance. [*** discuss Newton's third rule ***].

23. The mathematician will balk at this *merely*, but she is not my primarily intended audience. Still, it would please me were she able to read this paper with some profit, so I hope such rhetorical flourishes do not put her off too much.

24. Bondi (1962, p. 132). Italics are Bondi's.

of some set of quantities relevant to systems treated by the theory satisfy certain constraints. We begin with a few considerations on how one may constrain the measure of spatiotemporal intervals, which will culminate in a few quasi-technical definitions and results needed for a quasi-technical analysis of the idea of a regime.

Real initial data for real physical problems are not specified with arbitrary accuracy over an arbitrary region of a spacelike hypersurface. It is less of an idealization to model initial-data as occupying a compact, connected region of spacetime, of non-zero metrical volume, determined by the spatial extent of the system in conjunction with the temporal interval during which the measurement or preparation of the initial-data takes place. As we have seen in the discussion of the dispute over the hyperbolic extensions to the Navier-Stokes system, moreover, the determination of the values of real physical quantities appropriate for use in initial-data for a given system will always be coarse-grained in the sense that the they are more properly modeled as being attached to compact, connected sets of non-zero metrical volume, satisfying certain collateral metrical conditions, contained in the region occupied by the physical system, rather than as being attached to individual spacetime points contained in the region occupied by the physical system. These sets, moreover, should be as small as possible, in order to maximize the accuracy of the modeling of the experimental apparatus, while still being large enough to satisfy the constraints the theory places on the definition and measurement of its quantities and on the satisfaction of its equations, under the specified environmental conditions, using the chosen methods of observation.

To study some physical phenomena modeled by a particular theory, then, we first need a compact, connected region of spacetime of non-zero metrical volume, which for the purposes of this discussion we may without a great loss of generality assume to have properties as nice as we choose (we may demand, *e.g.*, that it be the closure of an open, convex, normal set), as the stage on which the phenomena will unfold and the experiment be played out. The theory may impose further requirements on the region; it may demand, *e.g.*, that its spatial and temporal dimensions (as determined in a specified manner) satisfy a set of algebraic constraints, or that the curvature in the region satisfy a set of differential and algebraic conditions. Once so much is settled, the difficulty lies in partitioning our region into components appropriate for the determination of the values of the quantities modeled by the theory. Again, these components need to satisfy whatever constraints the nature of the quantities demand. It makes no sense in general to attempt to determine the temperature of a system, *e.g.*, on scales smaller than the mean free path and the mean free time of flight of the system's dynamically relevant constituents, as determined in a kinematically relevant way. For a sample of nitrogen gas under "normal" conditions on the surface of the Earth, for example, this would include the relevant measurements of the nitrogen molecules, not of their electrons and nucleons, as calculated in a frame co-moving with the surface of the Earth and not in one spinning wildly and moving at half the speed of light with respect to it. We therefore require that the individual regions to which values of temperature are to be ascribed be larger than, in an appropriate sense, those characterized by the theory's break-down scales. Similarly, if we are to try to model a sample of nitrogen gas using the Navier-Stokes equations, for instance, then we must ensure that the dynamical evolution of the system is such that the gradient of its temperature on those scales not be too great just off points of equilibrium (as it settles down to equilibrium, *e.g.*, during preparation).

A Maxwell-Boltzmann sort of partitioning of phase-space, and *eo ipso* of the spatiotemporal region occupied by the system itself, into scraps of roughly equivalent volume and shape offers the most obvious way forward at this point.[25] I do not find this solution satisfactory, however—or, rather, I find it satisfactory for the particular treatment of the statistical mechanics of a more or less ideal gas, but I do not find it satisfactory for the generic treatment of the modeling of constraints on the determination of the values of quantities for many other kinds of physical theory. Although

25. Synge (1957) has worked out such a device in great detail for the statistical-mechanical treatment of ideal, relativistic gases.

thermodynamics *cum* statistical mechanics provides the easiest and most straightforward examples of the kinds of constraints that interest us, I would argue that such constraints form an integral part of the nitty-gritty of *every* physical theory, no matter how seemingly "fundamental", as I gestured at above.

Let us try to sketch a construction of a different sort of partitioning of a spatiotemporal region.[26] Fix a compact, connected subset C of spacetime, of non-zero metrical volume, representing the spatiotemporal region in which the physical phenomena we would model play out. We demand that such regions satisfy a few basic, generic, topological and metrical conditions, mostly along the lines of guaranteeing that the region is not "too small along either its spacelike or its timelike dimensions", that its boundary is well-behaved, and so on. We will call such a region a *canvas*. More precisely, a canvas is a convex, normal, compact, connected, 4-dimensional, embedded submanifold of $\mathcal{M}$.[27] We will use canvases to model the spatiotemporal regions physical systems occupy in which a specified family of observations and measurements occurs, as well as to model the elements into which such regions will be carved for the purpose of serving as "points" of the system to which values of its associated quantities may be meaningfully ascribed (as opposed to ascribing the values of the quantities to points of spacetime itself).

To give a flavor of the sorts of algebraic conditions one may demand of the elements of such a partition, we first require terms in which to express the conditions. There is an endless supply of theoretical terms one could employ to do so. I offer here only a sampling, by way of example. I do not think that these have a preferred status over others one could propose. I offer them because they seem to me to be reasonably clear, to be easy to visualize and to have straightforward, meaningful physical content. Other sets of terms could well serve better the purposes of a particular investigation. Such choices are, I think, fundamentally of a pragmatic and æsthetical character. Choose, then, an element O of the proposed partition of C and a point q on the boundary of O, and consider the family of all spacelike geodesics whose intersection with $\breve{O}$ (the interior of O) consists of a connected arc one of whose points of intersection with the boundary of O is q. Calculate the supremum of the absolute values of the proper affine length of all these arcs. Finally take the infimum of all these suprema for every point q on the boundary. This is the *infimal spacelike diameter* of O. The *infimal timelike diameter* is calculated in the analogous way, using timelike rather than spacelike arcs. We take the infimum of the suprema, as the simple infimum of the lengths, for a Lorentzian metric, would in general be zero, as the arcs may approximate as closely as one wishes to a null arc. Note that the spacelike or the timelike infimal diameter of a connected set with non-zero metrical volume will always be greater than zero (so long as the metric is "well behaved", which we henceforth assume). It thus follows directly from the definition of a canvas that its infimal spacelike and timelike diameters are always both greater than zero. Also, any 4-dimensional set with non-zero infimal spacelike and timelike diameters has non-zero volume with respect to the spacetime's volume element (we deal only with measurable sets in this paper), as one can always fit a non-trivial open set inside it (*e.g.*, a small tubular neighborhood of a geodetic, spacelike arc whose length is within some $\epsilon > 0$ of the infimal spacelike diameter). Thus it also follows that a canvas has non-zero measure with respect to the volume element of spacetime. We will use these sorts of properties of canvases, especially those relating to their infimal diameters, to articulate the first kind of constraints a regime imposes on a theory, those directly addressing characteristic spatial and temporal measures of spatiotemporal regions appropriate for the application of the theory. [*** briefly sketch one possible way to think of the physical content of these diameters—that the 'longest' way across the region for a particle or rod crossing near the center of the region will never be smaller

26. Again, for the rigorous details, see Curiel (2011).

27. The full definition (see Curiel (2011)) includes the proviso that ∂C not be a null 3-space with respect to the spacetime metric. The exclusion of null hypersurfaces ensures that certain integrations and operations on the boundary are always well defined. Since light never travels, so far as we know, in a true vacuum in any real physical situation, this is a negligible exclusion for the goal of modeling real, inaccurate data over finite spatiotemporal regions.

than this amount ***]

It is not so easy to articulate terms in which the second half of the possible constraints on the character of spatiotemporal regions appropriate for the definition of physical quantities, those pertaining to the general behavior of the metric in the region, may be formulated. For instance, one can impose constraints on the intensity of the curvature in a region in any of a number of ways, from, say, fixing an upper bound on the total integral of any scalar curvature-invariant over the region to fixing an upper bound on the average of such a scalar or an upper bound on the value of that scalar at any given point in the region; one may as well, for example, fix an upper bound on the integrated components of the Riemann tensor as measured with respect to a parallel-propagated frame-field along any timelike geodetic arc contained in the region; and so on. There are more general sorts of considerations one may bring into play as well, including the imposition of some kind of causality conditions (*e.g.*, that the region contain no almost closed, timelike curves), an exclusion of certain kinds of singular structure (*e.g.*, that the region contain no incomplete timelike geodesics), a restriction on the topology of the spacetime manifold (*e.g.*, that its second Stiefel-Whitney class vanish, the necessary and sufficient condition for a spacetime manifold to admit a globally defined, unambiguous spinor-structure—see Geroch (1969) and Geroch (1970b)), and other general, metrical considerations (*e.g.*, that the spacetime be asymptotically flat). I will not attempt to characterize with any formality these sorts of constraints, restricting myself mostly to speaking only of constraints on spatial and temporal measures, primarily because I see no way of doing so for the former in light of their amorphous nature, not because I think they are unimportant or not worth considering. On the contrary, I think it would be of enormous interest to construct a formalism for studying these sorts of constraints. In any event, the reader should bear in mind that, from hereon, when I speak of constraints on spatial and temporal measure I do not mean to exclude the other sort from consideration. [*** remark—for reasons like those adumbrated in Curiel (1999)—that it is far more difficult to lend such constraints clear physical content ***]

A set of algebraic constraints on the measure of temporal and spatial intervals, then, is a formal system of equations and inequalities with some number (greater than zero!) of unknown terms, each term representing a characteristic temporal or spatial scale associated with the quantities modeled by the theory. For the sake of simplicity, we will assume that, for any set of algebraic constraints on temporal and spatial measures associated with the regime of a theory, there are only two unknowns used in all the expressions in the set, which we will interpret respectively as the spacelike and timelike infimal diameter of any region that is a candidate for having the values of the theory's quantities legitimately determined on it. A canvas *satisfies* such a set if its two diameters jointly satisfy the elements of the set. In the case of a relativistic Navier-Stokes fluid, for instance, we know that, for any element of a partition of the region it occupies, the infimal spacelike diameter ought to be strictly greater than c times the infimal timelike diameter (see equation (4.2.3.2)). We also know that the two infimal diameters ought to be, respectively, at least of the order of the mean length of the free path and the mean time of free flight of the fluid's molecules, as determined in a "reasonable" frame. [28]

4.3.2 Infimal Decoupages

We will now sketch the proposed manner of generically partitioning a canvas into elements to which we may apply our algebraic constraints. Fix a set of algebraic constraints on spatial and temporal measures and a canvas C satisfying the chosen constraints in such a way that the canvas

28. With a little more effort, one can state this last condition in a relativistically invariant way, by stating it in terms of the measure of intervals along and separations between timelike geodesics contained in the canvas such that two otherwise free particles instantiating two given timelike geodetic arcs contained in the canvas will, with a given probability, collide with a certain number of other particles traversing timelike geodetic arcs contained in the canvas closer than the given distance—get it?

contains as proper subsets other canvases also satisfying the constraints. A *scrap* S of the canvas is itself a canvas such that

1. it is a proper subset of C
2. it satisfies the constraints
3. its interior is topologically $\mathbb{R}^4$

The *decoupage* of a canvas is the family of all its scraps. A rich family of mathematical structures accrues to the decoupage in a natural way. It has, for instance, a natural topology under which it is Hausdorff, connected, and compact, if C itself is so. This topology can be extended to a σ-ring on which a Lebesgue measure, and thus integration of scalar fields, can be defined. See Curiel (2011) for details. The decoupage as characterized so far has the structure of an infinite-dimensional space. I found it convenient in Curiel (2011) to construct by the use of equivalence classes and to use in place of the decoupage a finite-dimensional space capturing in approximate form all the essential structure of the decoupage. The construction of this space raises several interesting questions about the nature of the sorts of approximations one deals with in physics, which we will not be able to address in this paper. In any event, from hereon, the term 'decoupage' will refer to this finite approximation rather than to the full, infinite-dimensional space.

We will attempt to capture the idea of "spatiotemporal regions whose dimensions satisfy certain constraints", the ones appropriate for taking as the elements of the partition of the region in which the phenomena occur, by using decoupages. There are, again, several ways one may go about it. I will sketch only one. The following consideration will be our primary guide. On the one hand, the details of the physical state of the system on regions smaller than the break-down scale are, if not irrelevant, then at least *ex hypothesi* not sensibly representable in the theory at issue or do not yield results consonant with the solutions of the equations, whereas, on the other, those regions significantly larger than the break-down scale are not so fine-grained as one can in principle make them for the purpose of maximizing the accuracy of observation and measurement. Given a theory with its attendant set of algebraic constraints on spatial and temporal measure, we require a way of specifying a family of subsets of a region that are in some sense or other as small as possible while still conforming to the theory's contraints. In general, neither the set of spacetime points constituting the region itself nor the whole decoupage itself of the region will serve the purpose.

Fix, then, a canvas $C \subset \mathcal{M}$ and a set $\mathfrak{m}$ of algebraic constraints on spatial and temporal intervals. The *infimal decoupage* of C, $\mathcal{C}^{\text{inf}}$, consists of all the scraps of C whose volumes are, in a certain precise sense,[29] as small as possible while still being consistent with $\mathfrak{m}$. An *infimal scrap* is a member of an infimal decoupage. Alternative definitions of an infimal decoupage could minimize, for instance, the volumes of the boundaries of the scraps, or a weighted average of the lengths of all the spacelike and timelike arcs contained in each scrap, or the average scalar curvature of each scrap, or some combination of these, and so on. I choose the definition based on volume not because I think it is *a priori* superior to the alternatives, but rather because it is simple, intuitively clear, and suggestive of the usual Maxwell-Boltzmann partition of phase space in statistical mechanics. One of the alternatives could well fit the purposes of some particular analysis or investigation more closely.

It makes no sense to talk about the temperature, *e.g.*, in regions on a scale finer than that characteristic of the break-down of the modeling of the given system, nor indeed to speak of possible solutions to the partial-differential equations of a theory on a finer scale, for the equations are no longer satisfied to the desired degree of accuracy in that regime, if they are well posed at all. This is why, in a substantive sense, a real-valued function whose domain is an infimal decoupage more appropriately models the details of the physical state associated with the fluid's temperature, *e.g.*, than does a scalar field on a subset of spacetime: it forces attention to be focused on those details and

29. See Curiel (2011).

only those details both relevant to the experimental problem at hand and sensical with respect to the theory being applied. For a Navier-Stokes fluid contained in a spacetemporal region, for example, the break-down scale, as discussed in §4.2.3, defines part of the set of constraints on the spatial and temporal intervals over which the fluid's quantities are well defined and over which the solutions to the equations themselves model the fluid's actual dynamical evolution to the desired degree of accuracy, and so fixes the infimal decoupage over the scraps of which the quantities associated with the fluid should be considered fields.

Still, this all may sound more than superficially similar to the standard Maxwell-Boltzmann sort of partitioning. It differs from that device in important ways, however. Primary among the differences are two. First, in this scheme the scraps overlap in a densely promiscuous fashion. Thus, even though one can speak of, *e.g.*, the temperature only on finite scraps rather than as associated with individual spacetime points, one can still speak of the temperature on such scraps arbitrarily "close" to each other in a topological sense. This may seem a slight advantage at best, but, as is shown in Curiel (2011), exactly this aspect of the machinery developed here allows one, in complete contradistinction to the ordinary Maxwell-Boltzmann partition of phase space, to bring to bear with complete rigor the full battery of mathematical structures one most often employs in attacking both theoretical and practical problems in physics, including topology, measure theory, differential topology, differential geometry, the theory of distributions and the theory of partial-differential equations on finite-dimensional manifolds. One can then use these structure to articulate and prove results of some interest (*e.g.*, theorem 4.6.4.1 below) illuminating the relations among ordinary scalar fields as employed in theoretical physics and fields defined on these decoupages that, I argue, more appropriately model the data gathered during the course of and used for the modeling of actual experiments. Second, and at least as important, I do not see any other way of attempting to define such a partition in a relativistically meaningful and useful way. The standard Maxwell-Boltzmann device fixes the partition of the observatory once and for all into a finite lattice of scraps. This partition may provide excellent service for one observer but be next to useless, or worse, for another. The idea of the infimal decoupage allows one to take account of all such partitions all at once, as it were, in an invariant manner.

4.3.3 The Kinematical Regime

Recall that the first type of failure of a theory's applicability to a given system stemmed from the ceasing to be well defined for one reason or another of the quantities of the theory for the given system. The International Practical Temperature Scale of 1927, as revised in 1948, 1955 and 1960, provides an excellent, concrete example of this phenomenon. [30] For example, the thermodynamical temperature scale between the primary fixed point 0.01° Celsius (the triple point of water at one standard atmosphere) and the secondary fixed point 630.5° Celsius (the equilibrium point between liquid and solid antimony at one standard atmosphere) is defined [31] by interpolation, using the variation in resistance of a standard platinum wire according to the equation of Callendar (1887):

$$t = 100\left(\frac{R_t - R_0}{R_{100} - R_0}\right) + \delta\left(\frac{t}{100} - 1\right)\left(\frac{t}{100}\right)$$

where

— R_0 is the resistance of platinum as measured with the thermometer immersed in an air-saturated ice-water mixture at equilibrium, at which point the ice-point temperature is unaffected (to an accuracy of ±0.001° Celsius) by barometric pressure variations from 28.50

30. *Cf.*, respectively, Burgess (1928), Stimson (1949), Hall (1955) and Stimson (1961).

31. See, *e.g.*, Benedict (1969, §§4.1–4.4, pp. 24–9). This reference is not the most up-to-date with regard to the international agreement on defining the standard, practical methods for the determination of temperature, but I have found no better reference for the nuts and bolts of thermometry.

inches to 31.00 inches of mercury, and the resistance of the wire is independent of the static water pressure up to an immersion-depth of 6 inches at sea-level

— R_{100} is the resistance of platinum as measured with the thermometer immersed in saturated steam at equilibrium under atmospheric pressure (as determined using a hypsometer), though corrections must be carefully made in this determination, the steam-point temperature being greatly affected by variations in barometric pressure (for which, standard tables may be consulted)

— R_t is the resistance of platinum at temperature t (the temperature being measured), *i.e.*, R_t is itself the quantity being measured that allows the calculation therefrom of the environment's temperature

— δ is a characteristic constant of the particular type of thermometer employed, defined at the primary fixed point 444.6º Celsius (the equilibrium point between liquid and solid sulphur at one standard atmosphere)

Below 0.01º Celsius and above 630.5º Celsius, the Callendar equation quickly diverges from the thermodynamic scale. From 0.01º Celsius down to the primary fixed point -182.97º Celsius (the equilibrium point between liquid oxygen and its vapor at one standard atmosphere), the temperature is also based on the resistance of a standard platinum wire, the interpolation being defined by an emendation of Callendar's equation (transforming it from one quadratic to one cubic in the unknown temperature), known as van Dusen's equation; above 630.5º Celsius up to the primary fixed point 1063.0º Celsius (the equilibrium point between liquid and solid gold at one standard atmosphere), the temperature is based on the electromotive force generated by a 90%-platinum/10%-rhodium versus 100% platinum thermocouple, the interpolation being defined by the so-called parabolic equation of thermocoupling; above 1063.0º Celsius, the temperature is based on the measurement of the spectrum of radiation by an optical, narrow-band pyrometer, the interpolation being defined by Planck's radiation formula.[32] In all these cases, moreover, it is clear that one cannot speak of the temperature's being measured on a spatial scale more finely grained than that corresponding to the physical dimensions of the thermometric device employed, or on a temporal one more finely grained than that of the time it takes the state of the device to equilibrate when placed in proper thermal contact with the system under study, under the influence of fluctuations in the temperature of the system itself and its environment, under the given conditions.

As this example illustrates, the constraints on the definability and measurability of a quantity in a given theory must be variously given with regard to the parameters of particular types of systems under certain kinds of conditions, not generically, once and for all, in an attempt to constrain the definability and measurability of that quantity *simpliciter*. It is in part this very variability in the specification of a quantity's definition—that it is possible to make in a variety of ways—that leads us to think that we have cottoned on to a "real" quantity, and not one artifactual of this particular experimental arrangement under these particular conditions.[33] This example makes clear, moreover, that in modeling different ranges of values of a given quantity different theories must be used. If one treats phenomena in which temperatures rise above 1063° Celsius, for instance, one's theory must include, or have the capacity to call upon the resources of, at least that part of quantum field theory required for a Planckian treatment of electromagnetic radiation.[34] We will therefore assume, at a minimum, that the range of admissible values for any quantity modeled by a theory is bounded both from below and from above. In technical terms, this means that the family of scalar fields admissible for representing the distribution of the values of a quantity for any spatiotemporally extended system

32. See, *e.g.*, Benedict (1969, pp. 27).

33. See Newton's Third Rule of Natural Philosophy, at the beginning of Book III of the *Principia*, for a remarkably concise and incisive discussion of a few facets of this issue, with particular emphasis on the nomination of certain properties of a system as being *simple* or *fundamental* with regard to that system.

34. This remark suggests that my use of 'theory' does not entirely harmonize with standard usages in physics and in philosophy. I will discuss this point below in §4.5.5.

treated by the theory is itself uniformly bounded from below and from above. In a similar vein, we assume, roughly speaking, that the first several derivatives (the exact number being idiosyncratic to each theory) of all the scalar fields are uniformly bounded from above and below—there is no sense, for example, in employing scalar fields that oscillate wildly in regions smaller than the breakdown scale.[35]

To make these ideas precise, fix a physical theory comprising a system of partial-differential equations.

Definition 4.3.3.1 *A* kinematical regime of applicability *of a theory (or a* kinematical regime*, for short) is an ordered quintuplet* $\mathfrak{K} \equiv (\mathfrak{e}, \mathcal{E}, \mathfrak{k}, \mathfrak{m}_k, \mathcal{K})$ *such that*

1. $\mathfrak{e}$ *is the set of variables and constants the partial-differential equations of the theory are formulated in terms of*
2. $\mathcal{E}$ *is a finite set of variables and constants none of which appear in* $\mathfrak{e}$*, and thus not in any of the theory's equations*
3. $\mathfrak{k}$ *is a set of differential and algebraic conditions on the values of the elements of* $\mathfrak{e} \cup \mathcal{E}$*, including an upper and a lower uniform bound on values of the family of fields admissible for modeling values of the elements of* $\mathfrak{e}$*, as well as uniform upper and lower bounds on some fixed number of the derivatives (appropriately defined) of all fields admissible for modeling values of the rates of change of fields modeling the values of the elements of* $\mathfrak{e}$[36]
4. $\mathfrak{m}_k$ *is a set of algebraic conditions, possibly involving elements of* $\mathfrak{e} \cup \mathcal{E}$*, on the measure of spatial and temporal intervals*
5. $\mathcal{K}$ *is a set of particular types of interactions with the environment using particular kinds of experimental apparatus, in conjunction with methods for calculating the intervals of possible inaccuracy in preparing or measuring the quantities of the theory by dint, respectively, of those interactions, within given levels of confidence, under any particular set of circumstances conforming to the requirements imposed jointly by* $\mathfrak{k}$ *and* $\mathfrak{m}_k$

$\mathfrak{e}$ represents the quantities directly modeled by the theory through its partial-differential equations. $\mathcal{E}$ is to represent a set of environmental quantities the values of which play a role in the determination of the applicability of the theory but which are not themselves explicitly treated by the theory. The elements of $\mathfrak{k}$ constrain the values of those environmental variables, in addition to constraining the values of the quantities directly treated by the theory. $\mathfrak{k}$ attempts to capture the fact that the theory's quantities will remain well defined only under certain environmental conditions, and only while the quantities the theory treats do not exhibit behavior pathological with regard to other quantities treated by the theory. $\mathfrak{k}$ contains constraints on both collateral environmental quantities and the theory's own quantities because the two often are not extricable from each other. In the case of the relativistic Navier-Stokes fluid, for example, the ambient Maxwell field ought not be so intense as to ionize the fluid, but the value at which the Maxwell field ionizes the fluid will itself in general depend on the temperature of the fluid; the temperature of the fluid, likewise, should not be so high as to denature the molecules constituting the fluid. In either of those two cases, for example, the definition of the fluid's shear-stress would become ambiguous, dependent on how one accounted for the contributions to it of the various particles as they ionize, denature and recombine.[37] The

35. We will characterize more precisely this property, of a field's partials being uniformly bounded in a kinematically relevant way, below in §4.5.2. We will there, moreover, be able to extend these ideas to tensorial quantities by imposing a more or less natural, kinematical norm on such quantities, the one developed in purely formal terms in §4.4.5, and then demanding that this norm, and the norm of some appropriately derived set of tensorial quantities, be uniformly bounded.

36. We will work out in some detail what this may mean in a specific example, §4.4.8 below

37. I have particular qualms about the idealizations involved in positing a fixed set of kinematic constraints, to wit, a fixed, probably infinite conjunction of (at least) second-order quantified statements. It seems likely to me that, in

conditions contained in $\mathfrak{m}_k$ delimit the spatiotemporal ranges over which the quantities represented by terms in $\mathfrak{e}$ are well defined. As we have seen in the case of the relativistic Navier-Stokes system, these constraints on spatial and temporal measures may employ terms in $\mathfrak{e} \cup \mathcal{E}$. A strong Maxwell field, *e.g.*, could affect the hydrodynamic scale of a gas by affecting the value of the mean free-path of the gas's molecules.[38] Finally, the interactions and the associated measurement techniques and methods of calculation contained in $\mathcal{K}$ allow one, at least in principle, to calculate the range of possible, inevitable inaccuracy in a given experimental determination of the value of a quantity under particular conditions.[39] Confidence in these techniques and methods will itself, presumably, depend in large measure on the results of *other* theories, those treating the measuring instruments and the relevant environmental factors. How (or whether!) the theoretical dependencies sort themselves out in the end in a more or less logical fashion is a fascinating question, but one well beyond the scope of this paper. I take it for granted, for the sake of my argument, that the details of this sorting out are irrelevant here.

The kinematical regime of a theory allows one to characterize those spacetime regions that may serve as appropriate arenæ of observation and measurement of the quantities of a theory, irrespective of whether or not the dynamical evolution of those quantities in that region match the predictions of the theory to any desired degree of accuracy.

Definition 4.3.3.2 *Given a theory and its kinematical regime $\mathfrak{K}$, a $\mathfrak{K}$-*appropriate observatory *(or* observatory *for short) is a canvas O such that*

1. *the values of $\mathfrak{e} \cup \mathcal{E}$ in O satisfy $\mathfrak{k}$*
2. *O satisfies $\mathfrak{m}_k$*
3. *one can consistently define the infimal decoupage of O*

Observatories are where good experiments relating to the theory may be performed. It is worth remarking that, in certain spacetimes and for certain sets of conditions $\mathfrak{k}$ and $\mathfrak{m}_k$, a theory may have no observatories at all, or may have no observatories in large swaths of the spacetime.

I believe it is acceptable to restrict observatories to compact subsets, even though this prevents us from specifying initial data on an entire Cauchy surface of a globally hyperbolic spacetime with a non-compact Cauchy surface, such as Schwarzschild spacetime. Only the relations between solutions to the partial-differential equations of mathematical physics on the one hand and actual data specified and collected in actual experiments concerns us here. No matter how much we may wish to (or

---

practice, nothing remotely approximating such a fixed set of conditions exists, even *could* exist, covering all possible experimental circumstances as modeled by a particular theory. The case rather seems to me to be more along the lines of the conclusion of the analysis of Anscombe (1971) of the conditions under which ordinarily caused events, *e.g.*, the lighting of a match, are taken to occur and not to occur. She argues that rather than stipulating a fixed list once and for all—in this case, perhaps to include the statements that the match not be wet, that the temperature not be too low, *etc.*—it is more appropriate to stipulate that, when the expected event does not occur, it behooves one to look for a contravening cause as best one may, without recourse to such a list. Though she fails to remark on this, one ought to note that this analysis, spot-on in many ways as it may be, raises the question—when ought one *expect* a given outcome, without something very like a tentative list of necessary conditions? Needless to say, this issue is too vexed to address here—or, really, I fear, anywhere.

38. Note that the constraints represented by the elements of this set do not depend on the types of measurements and interactions considered—they are, as it were, absolute constraints—the character of a particular experiment, depending, *e.g.*, on the sorts of apparatus used for measuring a quantity, may place coarser constraints on the measure of spatiotemporal intervals than those imposed by $\mathfrak{m}$, but we will not consider this complication at the moment.

39. This set is the cause of yet another in a growing list of qualms I have about the idealizations and simplifications with regard to my characterization of a regime. I doubt seriously that anything even vaguely approximating a complete set of this sort concerning the actual practice of physicists could be compiled, even for a fixed, single moment of time with its fixed state of technical competency for the field as a whole. This state of competency changes, advances and regresses coninually, in all sorts of ways. Again, I seek solace in the fact that the task I have set myself consists only in this, a demonstration that such a thing as a regime can in principle be characterized so as plausibly to represent *in specie* the way the practices of the theoretician and of the experimentalist dovetail, if they indeed do.

be glad we think we cannot) have the capacity to perform experiments unbounded in spatial and temporal extent, we in fact cannot, given the current state of our theoretical knowledge and technical prowess.

It is worth remarking that, even at this early stage of the game, the idea of a regime makes itself useful: it shortcuts the problem of truncation, discussed at the end of §4.2.3. To ensure the applicability of one of the Navier-Stokes hyperbolizations, for example, one need only demand that the only higher-order terms of a size to manifest effects at the considered scales be those involved in the explicitly introduced novel terms in the equations (assuming that one has laid down an interpretation of these terms by reference to physical quantities amenable to physical probing). It is also important to remark, however, that this is a purely formal solution to that particular ill of the Navier-Stokes hyperbolizations. This may not be a satisfactory physical solution, for it is not obvious at all that it is physically plausible to demand of a Navier-Stokes fluid, when considered at scales comparable to any of its breakdown-scales, that only some small subset of the envisioned transient fluxes be large relative to the others.

4.4 Physical Fields

Crudely speaking, a *physical theory* is one possessing a fixed regime. In §4.5 we will be more precise and propose a somewhat formal definition of a representation of the type of theory in appropriate possession of all the features we have been discussing. In order to get there, we must first complete the work begun in §4.3.2 above, by making precise the sorts of mathematical objects to be used in the modeling of physical quantities in conformance with a theoretical regime. Ordinary scalar fields on spacetime will not serve the purpose, for their range, ordinary scalars, does not account for the inevitable inaccuracy in the determination of the values of physical quantities, as articulated in the kinematical regime of a theory; and such fields do not have the proper domain of definition, which should be the infimal decoupage of a canvas rather than (some subset of) spacetime. The first order of business, then, is to define a space to serve as the appropriate range for our fields to have and to characterize the structure of this space, before using it to define fields over infimal decoupages, which will constitute the desired representation of fields of physical quantities as modeled by a theory with its regime.

Before we begin constructing a representation of such quantities proper for use in physical theories, we must delineate the roles these quantities will be expected to play, which is to say, the sorts of properties they ought (and ought not) to have, and the sorts of operations on them we require they make available to us. We begin by taking up this issue in §4.4.1, before moving on, in §4.4.2, to propose a way to define a space of scalar objects suitable to play the delineated role, and in §4.4.3 to endow this space with algebraic operations in conformance with the results of the reflections in §4.4.1. In preparation for treating fields of such objects on infimal decoupages, in §4.4.8, we first consider, in §4.4.4–§4.4.7, these scalar fields on ordinary manifolds, extensions of these fields to the analogues of tensorial fields, and the analogue of linear operators on them, such as derivations and integrals of these.

4.4.1 Algebraic Operations on the Values of Quantities Treated by a Physical Theory

Since our proximate goal is the defining of operations akin to integration and derivation on the fields we will construct to represent physical quantities in a way conformable to the requirements of a regime, it would be pleasant to have something akin to a linear, normed structure on a space comprising them, to mimic as closely as possible the behavior of the space of scalars $\mathbb{R}$ and the space of fields Σ composed of these scalars as used in theoretical physics. Before attempting to define and impose such structures, however, we must pause to consider the intended physical meaning of such operations and mappings, what it may mean in the context of physical theory to add together several values of a quantity associated with a physical system, or to multiply such a value by a scalar, and so on. This discussion, while interesting in itself, may not seem required here, but I think it is, as suggested by a simple example. Say we are considering the subtraction, one from another, of two values of a physical quantity, along with their respective, associated inaccuracies. Say that the modeling of a physical interaction requires that we subtract one of the magnitudes of the determined values of the quantities from the other. This seems straightforward enough—one applies the standard, additive group-operation on $\mathbb{R}$. How ought we combine the inaccuracies, though? One cannot apply the same additive operation, as this may yield a negative value, which makes no sense for the inaccuracy in measurement of a physical quantity (assuming, as we will, that the possible inaccuracy measures the absolute length of the interval within which the determined magnitude of the value may fall). How one does it in practice would seem to depend on circumstances such as the nature of the physical quantity, the nature of the experimental apparatus and techniques employed, *etc.* On the other hand, if one is trying to strike an average over time of the inaccuracies or measure their deviance from some fixed value over time, or some operation of this sort, it may make perfect sense to have a negative value for the inaccuracy. It seems, then, that how one handles the inaccuracy depends at least in part on the sort of operation one wants to apply to the values of the physical quantities. Indeed, I will argue that the signification of standard algebraic operations as applied in physics is not unambiguous in and of itself. We must, therefore, get clear on the different senses they may have, so that, when defining operations on our constructed space, including those on the inaccuracies, we may fix the intended sense our operations are meant to schematize, and thus have a partial guide in constructing the operations.

In physics, the application of the same algebraic operation in form and appearance can have one of at least three distinct kinds of signification. [40] Consider addition.

1. One can add, at the same point in spacetime, the Maxwell tensors representing the Maxwell fields associated with two separate charge-distributions, in computing the total Maxwell field at that point, in virtue of the linearity of Maxwell's equation.
2. One can add the vectors representing the velocities of two different bodies with respect to a third, in computing the velocity of one of the two in a new frame of reference, in virtue of the linearity of the Galileian transformations.
3. One can add the values of the gradient of the temperature of a body at two of its separate points, or at two different times, in striking an average, in virtue of the linear, additive group-operation available on the vector-space $\mathbb{R}^3$.

The first exemplifies an operation that represents an aspect of the true, physical character, as it were, of a state or process associated with a given type of physical system, in this case the superposability of Maxwell fields. The third exemplifies an operation with no true physical signification whatsoever (this bald statement will be explained in a moment), the computation of the spatially or temporally distributed average of the temperature of a body, but rather one whose employment we find handy for a variety of practical reasons, some of them tending to the furtherance of physical investigation and others to the furtherance of more pedestrian concerns. The second occupies a funny no-man's land: on the one hand, it embodies nothing more than the preferences we often have for the particular form in which we represent to ourselves the states and processes of physical systems and signifies nothing about the true, physical character of the system under study; on the other hand, the nature of mathematical representations of physical theory often, if not always, demands that we muster

40. I do not mean to suggest either that the classification I propose includes all possible significations of algebraic operations as applied in physical theory, or that it exemplifies the only method of classification of such operations. This is merely the one I require for the task at hand.

such a preference even when we would rather not, demands that we choose one from among a fixed class of such superficially different yet physically equivalent forms on the basis of nothing more than our preferences, if theory is to find application in the quantitative modeling of physical phenomena.

To keep these three straight, I will nominate them as follows. I will call operations as used in the first context, those applicable to the representations of the true, physical quantities of the system, *physical*; those as used in the second, reflecting our preferences in choosing the representation of these quantities, *psychological*; and those in the third, bearing on the use we make of our representations, *pragmatic*. I will extend this usage promiscuously, for the qualification of the names of scalars, structures, *etc.* We are not used to distinguishing among these three, I think, because the peculiarly simple properties of the mathematical structures standardly used in theoretical physics allow the use of formally identical algebraic operations to represent all three, and so mask the differences in signification.

Let me try to clarify what I mean with an example illustrating the difference between operations of the first and the second kind. One is reading in a text-book on Newtonian mechanics a description of the modeling of a bicyclist who has been trundling along at 8 mph; the book proceeds to claim that, at a certain time, one ought to multiply 8 by 2 to represent the bicyclist's current speed. Though we do not often think so, there is a possible ambiguity in what the writer is claiming (though, I must emphasize, the ambiguity is almost never a problem for the reader's grasping of the sense of the writer, as context tends to disambiguate it—indeed, context tends to disambiguate it with such an immediacy, clarity and finality as to make us almost never aware of the possible ambiguity in the first place). She may be saying that the bicyclist is now traveling twice as fast as before. She may rather be saying, however, that, for whatever reason, we are changing our units of measurement from miles per hour to half-miles per hour. Likewise, if she says that one ought to add 2 to 8 to represent the current speed of the bicyclist, she may be saying either that the bicyclist is now traveling with a speed of 10 mph, or else that, for whatever reason, we are now changing to a system of units the zero-point of which is what we would have referred to as '2 mph' in the original one (say, the "laboratory frame", of which text-book writers are so fond).[41] Such operations do not manifest themselves in physics only in the choosing of scales and zero-points for units of measure. The inevitable arbitrariness inherent in formulating a Lagrangian representation of a system provides another example. In the case of Lagrangian mechanics, for instance, the presentation of the space of states is, up to trivial isomorphisms, fixed once and for all; the Lagrangian function itself, however, is wildly indeterminate, in the sense that one can, without changing the solution to the equation, add to the Lagrangian any function that will not contribute to its total variation over any path.[42] The adding of such a function represents only a preference we may have for the representation of the system at hand, and nothing of true physical substance *vis-à-vis* the character of the system.

To illustrate the differences between the first and the third types of operation, the physical and the pragmatic, consider, again, the operation of addition. Naively, that the values of all physical quantities are represented by real numbers suggests that these values may always be added, and thus that any quantity represented by such a structure satisfies a principle of linear superposition, such as a Maxwell field does. Otherwise, what sense can there be in adding and subtracting the values

41. I make this example intentionally out of the ordinary, with respect to the modification of units, to emphasize the point that the manipulation of the values of quantities treated by every branch of physics displays this ambiguity. The reader may find the point easier to swallow by reframing the example with the use of temperature, in terms of the Fahrenheit and Celsius scales.

42. To be more precise, the first statement means that, if one is given a manifold on which one can formulate and solve the Euler-Lagrange equation, then it follows that the manifold is the tangent bundle of some manifold, and, moreover, the structure of the space of solutions of the Euler-Lagrange equation suffices for the complete determination of the structure of the original manifold as a tangent bundle over a determinate configuration space. If one then fixes a Lagrangian on this tangent bundle and adds to it a scalar field defined by a 1-form on configuration space, the modified Lagrangian will determine the same dynamical vector-field on the tangent bundle as the original one. This is the content of the second statement. See, *e.g.*, Curiel (2014) for details.

representing the quantity, as seems to be done when, say, striking averages, as may be done with the values of any physical quantity? In fact, however, this need not be the case. This addition does not, in general, represent the physical superposition of two manifestations of the quantity; rather, it represents a purely formal operation we perform to compute the value of a factitious quantity, such as the average or a certain approximation of the gradient. It makes no sense, *e.g.*, to add the values of the mass-density or the temperature of two perfect fluids mixed together, but it does make sense to ask for the average of those densities and temperatures (a pragmatic operation), just as it makes sense to calculate the resultant density and temperature of the mixture by the addition of weighted terms and to calculate the spatial variation of the values of these quantities (physical operations), even when no physical significance attaches to the adding to or subtracting from the value at one point of that at another.

The real differences among the three can, I believe, be summed up in the following observations. Assume we are treating a physical system with 6 degrees of freedom.[43] Then the physical operations apply to those quantities (the *physical quantities*) of which 6 taken together are necessary and sufficient for the complete determination of the state of the system at a given moment; these operations, furthermore, are such that their employment either signifies some actual modification or qualification to the physical state or dynamical evolution of the system (*e.g.*, by modeling an interaction of the once isolated system with its environment, such as the addition of a non-constant scalar field to the Hamiltonian), or else signifies the calculation of a physical quantity from some already known (or in principle knowable) other physical quantity (*e.g.*, the calculation of the gradient of the temperature of a body from knowledge of its temperature). More precisely (but not rigorously by any means), an operation is physical just in case, given a representation of the space of states of a system accomodating the operation, the operation acts either: 1) as a non-trivial mapping, to itself, of the class of vector-fields representing solutions to the partial-differential equations comprised by the theory treating the system; or 2) as a non-trivial mapping taking (in our example) a set of 6 physical quantities as represented by scalar fields, to a different set of 6 physical quantities as represented by 6 scalar fields, the values of which at a point represent the same state as those of the first set at that point.[44]

The pragmatic operations apply to quantities (the *pragmatic quantities*) that are such that, though calculable from physical quantities (and indeed calculable *only* from physical ones, perhaps mediately by the use of other pragmatic quantities that are themselves calculated from physical ones), no number of them taken together determine the state of the system at any moment. This statement will perhaps clear up a misunderstanding that may have been engendered by my use of examples. So far I have spoken blithely of averages as essentially unphysical. This is certainly not true in every theory. In statistical mechanics, for instance, the temperature of a body is, roughly speaking, defined as the average of the kinetic energies of the fundamental constituents of a body (fundamental, that is, with respect to the theory employed), which surely is a physical quantity in my sense of the term. It indeed is, in a theory essentially expressive of statistical mechanics, and it can in fact serve as one component in a determination of the state of a system as represented by the theory. In thermodynamics, on the contrary, temperature is not an average of anything; it is, if you will, a brute quantity. The point is that a quantity's counting as "physical" or as "pragmatic" depends on the nature of the theory at issue purporting to represent it—it makes no sense to declare a quantity or operation to be physical or pragmatic, in my senses of the terms, absent the context of any theory representing it or within which it finds application.

The *psychological operations* are such as to apply to the same quantities as the physical operations, but only in a Pickwickian sense: their use does not signify any modification or qualification

43. My use of of the term *degrees of freedom* is perhaps somewhat unusual—I mean by it the dimension of the reduced phase space, not of the reduced configuration space (assuming that any constraints on the system are integrable).

44. Note that this characterization holds for quantum systems as well. It can, as well, with some care, be extended to cover systems with an infinite number of degrees of freedom, such as Maxwell fields.

of the physical state or dynamical evolution of the system. More precisely, but again not rigorously, an operation is psychological just in case, given a representation of the space of states of a system accomodating the operation, the operation, up to appropriate isomorphism, commutes with the action of the operator representing the partial-differential equations comprised by the theory. In other words, speaking loosely, solving the equations for a given set of initial data and then executing the psychological operation on the resultant dynamical vector-field yields the same vector-field as does first executing the psychological operation and then solving the equations. It follows that, among many other things, a theory ought to specify what counts as an "appropriate isomorphism" (*e.g.*, a symplectomorphism in Hamiltonian mechanics). In fact, as we have seen, there are (at least) two distinct sub-types of psychological operations, those having to do with the defining of units of measure for physical quantities and those commonly thought of as gauge-transformations.[45]

It follows from these observations that, whereas the pragmatic operations available to us in the manipulation of the values of physical quantities are fixed once and for all, irrespective of the theory at issue—in physics as commonly practiced, comprising all the richness accruing to the space of real numbers in all its many *personæ* (as an additive group, multiplicative group, field, affine space, vector-space, Hilbert space, topological space, smooth manifold, Lie group, measure space, *et al.*)—, the physical and the psychological operations available to us are dictated by the character of the theory at issue. The spaces representing physical scalars in general are real, one-dimensional, differential manifolds (the minimum structure we demand), as, for example, those representing temperature and mass-density.[46] The structure of a differential manifold neatly and precisely captures all the fundamental properties required of such scalars—that, *e.g.*, integrals and derivatives make sense and have true physical significance even when addition and subtraction do not—in so far as they represent the values of such quantities. They do not have in general any further structure.

In special cases, such as with the space representing the values of the electric potential in electrostatics, one can impose further, richer structures on the space, such as that of a real affine space. The space of objects representing electric charge (non-quantized) has the further structure of a full vector-space. None of this can be assumed, however; it depends entirely on the nature of the physical quantity under study. In the case of the electrostatic potential, for example, the affine structure represents the fact that, while such potentials satisfy a principle of linear superposition, they have no natural zero-point; in the case of electric charge, we have both a linear operation and a natural zero-point, so we use a vector-space.[47] The mass-density of a system composed of two fluids that may be mixed in different proportions provides perhaps a more interesting example. In this case, for two masses, we know how to add and to multiply them, we know how to take their ratio, and we know, up to a point, how to subtract them from each other. In the absence of negative mass, however, we do not have a fully linear structure. This space has two natural structures accruing to it, which are isomorphic in a certain sense, though not naturally so. The first is a modular structure,

45. Since I do not think that diffeomorphisms in general relativity are properly thought of as gauge-transformations, for reasons too involved to enter into here, I must bracket their status *vis-à-vis* this classification. I will say only that their character seems to me closer to that of a change in the definition of units of measure than that of a gauge-transformation.

It perhaps points up a weakness in the paper as a whole that general relativity time and again offers up a structure not easily amenable to treatment by my definitions and arguments. Why it is always general relativity that seems to cause trouble in this discussion, as opposed to classical mechanics or quantum mechanics, is, I suspect, a question worth thinking about. Perhaps it has something to do with the inextricability, in the theory, of kinematics from dynamics.

46. In the case of quantum mechanics, such spaces are sometimes discrete, composed of the eigenvalues of a bounded, self-adjoint operator on the Hilbert space representing the states of the system at issue. I believe that all my definitions and arguments carry over essentially intact to this case, with only cosmetic alteration, though I will not go into any details on the matter.

47. We will not concern ourselves, for the sake of this example, with such high-falutin' stuff as magnetic charge and dual rotations of the Maxwell field.

over the associative, commutative ring whose fundamental group is the non-negative, real numbers under multiplication (and so the non-group operation in this case is addition, which implies that the ring has zero-divisors and so is not an integral domain). This represents the fact that, in general, the mass-density of the body consisting of the mixture of the two fluids will be a linear, strictly non-negative combination of the mass-densities of the component fluids. This structure suffices for the defining of operators whose actions correspond to those of integrals and derivatives respectively. The second structure accruing to it is that of a real measure space with a natural Lebesgue measure, which makes available exactly the same set of operations, so long as the restriction on the subtraction of one mass from another is adhered to.

Finally, in virtue of the fact that the space of any physical quantity has, at a minimum, the structure of a differential manifold, we are now in a position to see the proper interpretation of psychological scalars: those associated with a change in the definition of units are the components of particular coordinate presentations of (subsets of) such manifolds; those associated with gauge-transformations, on the other hand, live in a fiber bundle associated with the space of states of the system bearing the quantity, in the sense that the elements of the associated bundle have a natural action on the kinematically and dynamically relevant geometrical structure of that space. [48] Pragmatic quantities do not, so far as I can see, "live" anywhere. They are simply abstract, mathematical structures, such as the vector space of real numbers.

4.4.2 Inaccurate Scalars

With these considerations in mind, we turn now to the definition of our proposed space of scalars (we need not specify whether we are dealing with physical, psychological or pragmatic scalars until we attempt to introduce operations on the space). It will be convenient to define the following abbreviations. '$\mathbb{R}^+$' denotes the set $(0, \infty)$ of all strictly positive, real numbers, '$\mathbb{R}^\uparrow$' the set $[0, \infty)$ of all strictly non-negative, real numbers. For $\gamma > 0$, '$\mathbb{R}^+_\gamma$' denotes the set (γ, ∞) of all real numbers greater than γ, and '$\mathbb{R}^\uparrow_\gamma$' the set $[\gamma, \infty)$ of all real numbers greater than or equal to γ. For $\omega > 0$, '$\mathbb{R}_{<\omega}$' denotes the set $(-\omega, \omega)$ of all real numbers with absolute value less than ω. For any two real numbers γ^- and γ^+, $\mathbb{R}_{\gamma^-, \gamma^+}$' denotes (γ^-, γ^+), the space of all real numbers greater than γ^- and less than γ^+, and $\mathbb{R}_{[\gamma^-, \gamma^+]}$' denotes $[\gamma^-, \gamma^+]$, the space of all real numbers greater than or equal to γ^- and less than or equal to γ^+.

Now, let $\Re$ be the space of compact, connected, real intervals of non-zero length. For example, $[0, 1]$ is an element of this space, but $[0, 1)$ is not, nor is $[0, 1] \cup [2, \pi]$, nor $[\pi, \infty)$. Call it the *space of real intervals*. Because we are dealing only with compact, connected, real intervals of non-zero length, the standard Hausdorff metric on a space of sets is in this case a true distance function (*i.e.*, two intervals are at a separation of zero from each other if and only if they are identical). [49] $\Re$ is a two-dimensional Hausdorff topological space under the topology induced by this metric. I will refer to the greater value of an interval as its *top* and the lesser as its *bottom*. In this parametrization of the space, we will denote the element representing, *e.g.*, the interval $[0, 1]$ by '$(0, 1)$'; note that this denotes the ordered pair whose first element is the real number 0 and whose second is 1, and not the open, real interval from 0 to 1. Context should make clear which is meant.

In the event, however, the Hausdorff metric is not the most useful for our purposes, as it has little physical relevance under the interpretation we will impose below on $\Re$. We will rather use the

48. Specifying what counts as "kinematically" or "dynamically" relevant geometrical structure, and correlatively specifying conditions such an action must satisfy to be considered gauge, are issues beyond the scope of this discussion. I know of no single work in which all these questions are discussed taken as a whole. For discussions of various combinations of subsets of them, and related matters, see, *e.g.*, Geroch (1996), Trautman (1962), Trautman (1970a), Trautman (1970b) and Trautman (1980).

49. See, *e.g.*, Kolmogorov and Fomin (1970) for the definition of the Hausdorff metric on a set of sets of real quantities.

following variant of the Hausdorff metric, $\Delta : \Re \times \Re \to \mathbb{R}^{\uparrow}$,

$$\Delta((a,b),\,(y,z)) \equiv |y-a| + |z-b|$$

This is easily shown to be a Euclidean metric. In particular, for all $(a,b), (y,z), (m,n) \in \Re$,

1. $\Delta((a,b),\,(y,z)) \geq 0$
2. $\Delta((a,b),\,(y,z)) = \Delta((y,z),\,(a,b))$
3. $\Delta((a,b),\,(y,z)) + \Delta((y,z),\,(m,n)) \geq \Delta((a,b),\,(m,n))$

as easily verified. It is also easily shown that the topology induced by this metric is the same as that induced by the Hausdorff metric.

Now, one may think of $\Re$ as follows. Let the x-coordinate of the Cartesian plane represent the bottom, and the y-coordinate the top. Because we deal only with intervals of non-zero length, points on the line $y = x$ do not represent elements of the space, nor, by the nature of our chosen representation, do points below this line, at which the value of x is greater than that of y. This mapping of $\Re$ into the open half of the plane above the line $y = x$ is one-to-one and onto as well, and so is a point-wise isomorphism. Because, moreover, a continuous curve in the top half of the plane represents the shrinking and expanding in a continuous fashion of an interval on the real line (*i.e.*, the top and the bottom each trace out a continuous curve on the real line), it is natural to endow $\Re$ with the topology induced by this isomorphism, so that it is homeomorphic to $\mathbb{R}^2$. This topology, therefore, has all the nice properties one could wish for it, and so we will employ it in what follows. This contruction does not essentially depend on the fact that we consider intervals of length greater than zero. For $\gamma^-, \gamma^+ > 0$, let $\Re_{\gamma^-,\gamma^+}$ be the space of intervals of length strictly greater than γ^- and strictly less than γ^+. By the same argument, using this time the open strip [50] between the lines $y = x + \gamma^-$ and $y = x + \gamma^+$ rather than the half-plane above $y = x$, it follows that $\Re_\gamma$ is naturally homeomorphic to $\mathbb{R}^2$ as well.

The parametrization of $\Re$ by top and bottom is not, in the event, the most useful for our purposes. Because we are interpreting the elements of $\Re$ as ranges of possible inaccuracy, it seems not unreasonable to treat them as though the idealized, determined values about which they are ranges is the mid-point of the interval, which we will take as the first component of a representation of an element of $\Re$ in our new system of coordinates; we take the length of the interval as the second. In this scheme, the interval, say, [1, 2], would represent a determined value of 1.5 with a range of inaccuracy of ± 0.5, and so would be represented in our new system of coordinates by $(1.5, 1)$. From hereon, unless specifically stated otherwise, $\Re$ will be assumed to be parametrized with respect to these coordinates. For ease of expression, we will sometimes refer to the first component in this parametrization as the *magnitude*, and to the second as the *inaccuracy*. The idea is to have $\Re$, or some modification of it, serve as the appropriate range of values of fields modeling physical quantities in so far as they conform to the regime of a theory: the interval represents all the values a physical quantity may take, within the range of its possible inaccuracy in measurement and preparation. I will refer to $\Re$ in this guise as the *space of inaccurate scalars*.

In so far as the second component represents the possible or allowed inaccuracy of the magnitude of a quantity according to the regime of a given theory, it would seem that we ought to work exclusively with the space $\Re_{\gamma^-,\gamma^+}$, for some $\gamma^-, \gamma^+ > 0$, or some modification of it, depending on the particular theory at issue. The thought is this. In any experimental arrangement, a non-zero inaccuracy inevitably accrues to the measurement or preparation of initial data. The nature of physical quantities, moreover, as characterized in this paper, strongly suggests that this inaccuracy is in principle strictly bounded from below, away from zero, and strictly bounded from above by some finite value, for every physical quantity treated by a theory with a non-trivial regime. It makes no sense, for instance, to conclude that the inaccuracy in a determination of the time-of-arrival of

50. I am tempted to call this a *Las Vegas strip*—always open.

a particle at a sensor is greater than the known age of the universe, nor does it make any sense to conclude that the inaccuracy is less than 1 over 10-raised-to-the-power-of-itself-10-times seconds.

In practice, applying these structures to the modeling of a particular physical theory with its associated regime, one would sometimes want to work with only a single such pair of infimal and supremal inaccuracies for all quantities by the theory. In this case, one may take γ^- to be the supremum of the set of infimal inaccuracies accruing respectively to each of the quantities treated by the theory, and γ^+ to be the infimum of the set of supremal inaccuracies accruing respectively to each of the quantities treated by the theory,[51] so we will sometimes refer to γ^- in what follows as the *sup-inf inaccuracy*, and to γ^+ as the *inf-sup inaccuracy*. Recall from the discussion just before definition 4.3.3.1, moreover, that we demand as well that the absolute value of the magnitude of a given quantity, in so far as it is amenable to modeling by the theory, have a supremum, say $\omega > 0$ (for tensorial quantities, the magnitude may be some more or less natural norm imposed on the values of the quantities—see §4.4.5 below). The magnitudes of our scalars, then, in so far as they are to model only systems amenable to treatment by our theory, will take their values in $\mathbb{R}_{<\omega}$, for some $\omega > 0$, the *space of scalars of kinematically bounded absolute value* (or *kinematically bounded scalars*, for short). We thus really want our scalars to take their values in $\Re_{\omega,\gamma^\pm}$, the space of real, connected, compact intervals of length at least $\gamma > 0$, the supremum of the absolute values of the tops of which is strictly less than ω. Our chosen coordinates, then, take their values in $\Re_{\omega,\gamma^\pm} = \mathbb{R}_{<\omega} \times \mathbb{R}_{\gamma^-,\gamma^+}$. These considerations notwithstanding, we will not bother to keep explicit track of the value of the sup-inf and inf-sup inaccuracies in play. Neither will we bother to keep track of all the suprema of the kinematically bounded values of the magnitudes of all the quantities. Keeping track of either of these two numbers for the (more or less) strictly formal purposes of this section would complicate the exposition without a real gain in perspicacity. Except in a few places where it will be convenient or of interest to re-introduce γ^-, γ^+ or ω explicitly, we will use $\Re$.[52] All arguments and results in §4.4 can be modified so as to be stated in the terms of and hold for $\Re_{\omega,\gamma^\pm}$.

In fact, one can go farther than treating $\Re$ as merely a topological space. One can show that $\Re$ naturally has the structure of something akin to a 2-dimensional smooth manifold "almost" diffeomorphic to $\mathbb{R}^2$ (in its guise as a two-dimensional manifold). The 'something akin' and the 'almost' come from the peculiar nature of the intended interpretation of $\Re$, which requires a few modifications in how we treat its differential structure. When raising issues bearing on or relying on the differential structure, we will treat as admissible only charts that respect the difference, as it were, between the components of $\Re$. We demand that a chart mapping a subset of $\Re$ to $\mathbb{R}^2$ never "mix" the two components and, moreover, that the part of the chart mapping the second component restrict its range to $\mathbb{R}^+$, in order to comply with the kinematical constraints that led to our construction of $\Re$ in the first place.[53] In effect, we are treating $\Re$ as a two-dimensional space that locally has the structure of $\mathbb{R} \times \mathbb{R}^+$ rather than that of $\mathbb{R}^2$. In order to state this a little more precisely, define the projection operators $\pi_1 : \Re \to \mathbb{R}$ and $\pi_2 : \Re \to \mathbb{R}^+$ to be, respectively, projection on the first and second components of elements of $\Re$: for $(a,\, \chi) \in \Re$, $\pi_1(a,\, \chi) = a$ and $\pi_2(a,\, \chi) = \chi$. An admissible chart $\phi : \Re \to \mathbb{R} \times \mathbb{R}^+$, then, is one that can be expressed as a pair of diffeomorphisms $\phi_1 : \mathbb{R} \to \mathbb{R}$ and $\phi_2 : \mathbb{R}^+ \to \mathbb{R}^+$, in the sense that

$$\pi_1(\phi(a,\, \chi)) = \phi_1(\pi_1(a,\, \chi))$$

and

$$\pi_2(\phi(a,\, \chi)) = \phi_2(\pi_2(a,\, \chi))$$

51. This assumes that every theory treats only a finite number of quantities, or, at least, that the set of infimal inaccuracies of all the quantities is bounded from above, but this does not seem to me an onerous assumption.

52. "Pay no attention to the man behind the curtain!"

53. If we were keeping explicit track of ω and γ, we would demand that the part of the chart mapping the first component restrict its range to $\mathbb{R}_{<\omega}$, and that mapping the second restrict the range to $\mathbb{R}^{\uparrow}_{\gamma}$.

The meaning of fixing such a chart is a strictly psychological one, having to do with how one ought to change a given system of units for geometric quantities into another in such a way that respects the relation between the expression of the magnitude in the units of each and the expression of the inaccuracy in each of them. We will not consider it in any detail here, satisfying ourselves with the following observations. One can, in two different ways, decompose $\Re$ into a family of equivalence classes with a group-operation by $\mathbb{R}^+$ imposed on it, though we will not use these presentations in what follows after this discussion. For the first, consider, for some fixed $\gamma > 0$, the equivalence class of all elements of $\Re$ under the relation "being of the same length". For example, the intervals $[0,\, \gamma]$ and $[29,\, 29+\gamma]$ are in the same equivalence class, denoted (suggestively) by '$\bar{\Re}_\gamma$'. The group-action of $r \in \mathbb{R}^+$ is a multiplicative one, mapping, for example, the equivalence class $\bar{\Re}_\gamma$ to $\bar{\Re}_{r\gamma}$. The space of all such equivalence classes inherits from $\Re$ the structure of a 1-dimensional manifold. For the second, consider the space of equivalence classes of all elements of $\Re$ under the relation "equal up to a multiplicative constant r". Thus, for example, the intervals $[1,\, \pi]$ and $[2,\, 2\pi]$ are in the same equivalence class. Denote the equivalence class by the mid-point of the unique interval in the equivalence class of length 1. In our example, the equivalence class would be written '$\ll\!\frac{\pi+1}{2\pi-2}\!\gg$'. This is a real, 1-dimensional, affine space, where the affinity is given by the additive group-operation: the addition of a real number $r \in \mathbb{R}^+$ to an element of the space maps, for example, $\ll\!\frac{\pi+1}{2\pi-2}\!\gg$ to $\ll\!\frac{\pi+1}{2\pi-2}+r\!\gg$. These two spaces of equivalence classes and their group-actions have clear significance: they are both psychological. Multiplication of an element of the first by a strictly positive real number represents a re-scaling of one's units of measurement in $\Re$ by that factor. Addition of a scalar to an element of the second represents the choice of a new zero-point, a distance away from the old equal to the magnitude of the added scalar, for one's unit of measurement in $\Re$.[54]

Finally, before moving on, I must pause long enough to escort one issue high up into the nose-bleed seats of the bleachers, the very back of the gymnasium. Our chosen parametrization of $\Re_{\gamma,\omega}$, and our nomination of its elements as 'magnitude' and 'inaccuracy', suggests that there may be such a thing as "the actual magnitude of the quantity at issue, which our measurement approximates to, with, we hope, ever smaller error", in the sense of some variant or other of philosophical realism—*e.g.*, that there may be a single number that represents the actual, determinate value of the pressure in Torricellis inside this yet corked bottle of Taittinger, Brut, 1975, considered as a *Ding an sich*—which we could determine, if only we could make our probes sensitive and accurate enough. Well, there may be and there may not be. I take no stand on the issue in this paper. I do not feel I need to. Nothing in this paper hinges on it. I will sometimes use such words and speak in a manner that may suggest I have firm positions on these matters, but all such talky-talk should be taken with a large pebble of salt. I engage in it for the sake of brevity and ease of expression. The reader may

54. There seems to be something about the choosing of the scale of units that is richer in physical content than is the choosing of a zero-point for one's units. For instance, in general relativity, how a geometric object scales when the metric is multiplied by a constant, strictly positive number encodes a lot of information about that object, in particular about the so-called "dimension" of the object, whether it has, *e.g.*, the dimensions of stress-energy or of some other type of physical quantity. It is not difficult to show, for example, that the Riemann tensor, and so the Einstein tensor as well, does not rescale when the metric is multiplied by a constant (which shows, incidentally, that the scalar curvature must rescale as the inverse of the constant). It follows that, since the gravitational constant is dimensionless, a proper stress-energy tensor ought not rescale either, if it is to be a viable candidate for constituting the righthand-side of the Einstein field equation (see Curiel (2009) for details). Indeed, the fact that multiplying the metric by a constant does not alter the Riemann or the Einstein tensor shows it is a physically well defined operation; were it to have led to a different Riemann tensor, it would have altered the fundamental physics of the spacetime—it would in effect have defined a different spacetime. No similar proposition holds, so far as I know, with respect to the choosing of a zero-point for one's units. In fact, one cannot even in general do this in general relativity, in so far as the concept of adding a "constant", symmetric, two-index, covariant tensor to the metric is not defined in general, and, even when it is (say, in the vector-space $\mathbb{R}^4$ considered as a manifold), would yield an entirely different Riemann tensor than the first. Why is this?

supply such sense as he or she will (or won't) for the words during such periods of play. Having disposed now of the unruly spectator, expect to hear no more from him. [55]

4.4.3 Algebraic Operations on Inaccurate Scalars

Let us now try to use the considerations of the previous two sections to guide the attempt to impose various sorts of operations on $\Re$. We begin by dealing with pragmatic operations, as they are easier to manage, being fixed once and for all for all theories. On the face of it, the definition of the pragmatic operations are trivial. We are, after all, defining operations on the magnitudes and the inaccuracies of (potential) measurements of physical quantities the results of which do not purport themselves to be such magnitudes or inaccuracies. The average of a set of temperatures of a body over time is not itself the temperature of anything, and is, indeed, not a physical scalar at all. It's just a number. There is a serious worry, though: in striking averages, normalizing data-sets, computing standard deviations, and so on, how one ought treat the inaccuracy? In particular, ought one treat the magnitude in isolation from the inaccuracy, so that, *e.g.*, in adding two elements of $\Re$, the sum of the first components of each would give the first component of the result, irrespective of how the inaccuracies are dealt with? We will hold the second question in abeyance for the moment, assuming its answer to be 'yes', though we will return to consider it in our treatment of the physical scalars. Indeed, assuming the answer 'yes' makes the definition of pragmatic operations trivial. We simply treat the separate components as the real numbers they are, and pay no heed to any possible relation they may have. It does not matter, moreover, that, *e.g.*, the difference of two inaccuracies may turn out to be negative, for we are not, as already stressed, computing an inaccuracy with these operations, rather only numbers that purport to give us useful information about the inaccuracies. Thus, we need not worry about whether or not these operations respect the restrictions placed on the values of the inaccuracy in $\Re$. The pragmatic, algebraic operations on $\Re$, then, are the ordinary, algebraic operations of $\mathbb{R}$ applied to elements of $\Re$ component by component.

Matters are far more difficult when we turn to physical operations, as we now do (we will not treat the psychological—they are beyond the scope of this paper). We will spend some time working through some ultimately unsuccessful attempts, before settling on one that seems to me acceptable, as the failures will be edifying. It will be convenient, for the moment, to re-introduce an explicit value γ for the sup-inf inaccuracy. Let us try first making the simplest choices in defining operations on these spaces, to see how far *naiveté* will carry us. We begin with an additive, a subtractive, a multiplicative and a divisive operation defined, respectively, as follows:

$$\begin{aligned} (a,\, \chi) +_\phi (b,\, \psi) &= (a+b,\, \chi+\psi) \\ (a,\, \chi) -_\phi (b,\, \psi) &= \Big(a-b,\, \gamma^-_\phi(\chi,\psi)\Big) \\ (a,\, \chi) *_\phi (b,\, \psi) &= \big(ab,\, \gamma^*_\phi(\chi,\psi)\big) \\ (a,\, \chi)/_\phi(b,\, \psi) &= \Big(a/b,\, \gamma^/_\phi(\chi,\psi)\Big) \end{aligned} \tag{4.4.3.1}$$

55. For what it's worth, it is difficult for me to imagine that the question depends on anything more than the kind of semantical analysis one prefers for such words as "real" and "actual" and "empirical". I suppose if I were pressed on the issue, I would claim to be an atheist 6 days a week, declaring all such questions to be *Scheineprobleme*, but backslide on the Sabbath and come over all religious, declaring myself a knight of faith, and make the leap of the absurd into the waiting embrace of some variant of Peirce's convergent pragmatacism. That, at least, is how I strike myself today.

where

$$\begin{aligned}
\gamma_{-_\phi} &= \begin{cases} |\chi - \psi| & \text{if} \quad |\chi - \psi| > \gamma \\ \gamma & \text{otherwise} \end{cases} \\
\gamma^*_\phi &= \begin{cases} \chi\psi & \text{if} \quad \chi\psi > \gamma \\ \gamma & \text{otherwise} \end{cases} \\
\gamma^/_\phi &= \begin{cases} \chi/\psi & \text{if} \quad \chi/\psi > \gamma \\ \gamma & \text{otherwise} \end{cases}
\end{aligned} \tag{4.4.3.2}$$

The 'ϕ' subscripted to '+', '−', '∗' and '/' signifies that these are physical operations. Note that the operations on the right-hand sides of the equals-signs in equations (4.4.3.1) and (4.4.3.2) represent physical operations on $\mathbb{R}$, which is to say, the familar algebraic operations on real numbers. These all define closed operations, albeit ones with no additive, subtractive, multiplicative or divisive identity in general. The divisive operation is not, in general, commutative, though the other three are. More problematic is the fact that the last three operations are not associative. Assume, for instance, that $0.1 < \gamma < 1$; then $\gamma_{-_\phi}(5,\, \gamma_{-_\phi}(1.1,\, 1)) = 5 - \gamma$, whereas $\gamma_{-_\phi}(\gamma_{-_\phi}(5,\, 1.1),\, 1) = 2.9$. It is difficult to know how to proceed in the definition of other structures such as derivations and integrals without associativity. *Naiveté* has been suggestive, but has not taken us far.

The difficulties involved appear to be twofold. First, while $\mathbb{R}^+$ has a multiplicative group structure, it lacks the vector-space structure $\mathbb{R}$ we ordinarily rely on in performing these operations. Even were $\mathbb{R}^+$ to have had this, however, it would have been by no means clear that the correct way to have dealt with the inaccuracies associated with two magnitudes of a quantity, in subtracting them, one from another, for example, would have been by subtracting the inaccuracies as well—this could yield a value of zero or even a negative value for the inaccuracy, which is strictly *verboten*, in so far as, in this case, the number is meant to represent the inaccuracy in our knowledge of the magnitude of a physical quantity. The straightforward, unsubtle attack on the problem, in the persons of equations (4.4.3.1) and (4.4.3.2), ran squarely into this problem and failed to get past it.

It does not seem far-fetched, moreover, to imagine that, contrary to our assumption in the pragmatic case, in adding two elements of $\Re$, *e.g.*, the sum of the first components, the magnitudes, will have a non-trivial dependence both on the possible inaccuracies themselves, and on the fine details of how those inaccuracies may combine. This brings us to the second difficulty. In so far as the goal of this paper is to construct a generic model of the joint practice of the theoretician and the experimentalist, we want to define generic operations, once and for all, so as to be applicable to the magnitude and inaccuracy of any quantity, in any physical theory, without any notice taken of any idiosyncratic character of the quantity and the theory, much as the operation of the striking of an average of the value of a quantity in theoretical physics is defined once and for all, and applied promiscuously to all comers, irrespective of the character of the quantity or of its associated theory. I see no way, however, of answering such questions once and for all, with any fineness of grain, in a way applicable to the interplay between real physical data garnered from experiment and the descriptions and predictions offered by theory. The answer, for any particular case, will surely depend on (at least) the nature of the quantity, the nature of the interactions of the system being modeled, the nature of the experimental arrangements employed for observing the evolution of the system during its interactions, and, *a fortiori*, on the nature of the theory and its regime as well. Indeed, I wager there is no way to take account of all these factors even were one to attempt to construct, with even a moderate fineness of grain, for only a particular theory, and for only a restricted class of systems and experimental arrangements treated by the theory, a model of the transformation of the magnitudes of quantities and their associated inaccuracies during the dynamical evolution of such systems by the use of algebraic operations. The way, in particular, that the inaccuracies may combine seems to me to have an irremediably *ad hoc* character, albeit one governed by over-arching, generic, if highly abstract, principles (*e.g.*, that, in the long run, we expect the inaccuracies in determinations of a quantity to decrease, as more and more measurements are taken), just as the rules of hide-and-seek will be freely adapted by children to suit the particular characters of the field of play, the age and condition of the players, temporal constraints on the length of the game, and so on, while still remaining true to the core tenets of the game (for instance, that most of the children will hide and one, or at most a few of the rest, will try to find them). It is lucky for us that we do not require our model to have a fine grain. The nature of the project of this paper demands only that we construct *some* plausible model of the common playground and game of the theoretician and the experimentalist, one that, as it were, "has some seeming to it", not that we construct one that is natural in some sense, or that is the most accurate (such a thing as which I doubt the existence of, in any sense of the term 'accurate', even for a single theory).

To address the issue, we need some at least heuristic considerations to guide us. Consider what is known around the physics department at the University of Chicago as a 'Fermi problem'. Two of my favorite examples are "How far can a duck fly?" and "How many piano-tuners are there in Chicago?". The idea is to take a seemingly unanswerable question (in the absence of empirical investigation) and break it down into as many simple components as possible, the measure of simplicity for the components in this case being susceptibility to somewhat accurate, back-of-the-envelope estimation. The hope, then, is that, when one combines all the estimated answers to the simple components to compute the answer for the original question, the errors will tend to cancel each other out and the final result will be reasonably accurate. The name of such problems comes, passed down by word of mouth, from Fermi's almost preternatural ability to pose and solve them. To give an example, probably the most famous: at the detonation of the first nuclear device, during the Manhattan project in the deserts of Nevada, just before the explosion occurred, Fermi licked his index finger and reached it out just beyond the protective, concrete shield the observers stood behind; at the moment the explosion occurred he reached his clenched fist out just beyond the shield and released a flurry of shredded paper; after the shock-wave passed (about 40 seconds after the explosion), Fermi walked over whither the shock-wave had pushed the shreds, studied them for a moment, turned around and, to what, I am sure, must have been the utter bewilderment of his colleagues (Oppenheimer, Von Neumann, Bethe, Feynman, *et al.*), declared that the explosion had released an amount of energy equivalent to the explosion of 10 kilotons of TNT. 8 weeks later, when the Los Alamos computers, churning away day and night, had finished calculating the energy released on the basis of data collected from the most sophisticated instruments of the day, the result, 18.6 kilotons, differed from Fermi's estimate by only about 80%, well within an order of magnitude. Fermi already knew (roughly) or had good guesses at data such as: the distance of the shelter from the epicenter of the detonation; the density of the ambient air; the viscosity of the ambient air; the atmospheric pressure; the velocity of the ambient air just before the shock-wave passed; and so on. Based on these data, and on estimates he made on the spot, such as for the volume of the body of air the shreds of paper encompassed, the distance the shock-wave had pushed the shreds back, and how long it had taken it to do so, he computed the amount of energy that would have needed to have been released to have moved a body of air at the given distance from the epicenter, under the given conditions, the distance the air traveled in the time it took to travel that distance, as measured by the flight of the shreds of paper. [56]

It's easy enough to say that the errors tend to cancel out, but what does this really mean? In the example I gave of Fermi's computation of the output of energy by the nuclear device, it means something like this. Let's say that he overestimated the distance from the shield to the epicenter by 5%, underestimated the density of the ambient air by 3%, and so on. With enough such estimates in hand, the distribution of errors should begin to approximate a Gaussian curve centered on 0

56. After I first heard this story, a friend of mine and I were inspired to try it ourselves—we calculated the number of piano-tuners in Chicago. I came up with 100 and my friend came up with 32. There were something like 40 in the phone-book. Since my friend is an experimental neurophysiologist and I am a philosopher, I didn't feel so bad. I urge the reader to try it. It's fun.

(counting underestimates as negative numbers and overestimates as positive). In the worst case, the errors will be concentrated on one side or the other, strongly skewing the total, resulting error; in the best case, one will get something like a perfect distribution and the total, resulting error will approach zero. In the long run, the total, resulting error will tend to oscillate around zero with an average, absolute value of a smaller order of magnitude than the smallest (absolute) error in the bunch, with a variance of an even smaller order of magnitude. No matter how one algebraically combines the magnitudes of the quantities, the same reasoning should apply, that the errors will, in the long run, tend to cancel each other out, whether one is "adding" or "multiplying" or "subtracting", or what have you, the inaccuracies. We will adopt, therefore, only one template for physical operations on inaccuracies. To err on the simple side, let's say, then, that, to represent the way the errors combine in such computations as we have just discussed, we require an operation taking two arguments that is associative, commutative, monotonically decreasing in each component separately, and that always yields a value somewhat smaller than the smallest of the two, but never zero. Denote the result of combining two inaccuracies χ_1 and χ_2 by $\alpha(\chi_1,\, \chi_2)$. Then something like the following suggests itself.[57]

$$\alpha(\chi_1,\, \chi_2) \equiv \begin{cases} \exp\left(-1/(\chi_1+\chi_2)\right) & \text{if } 0 < (\chi_1+\chi_2) \leq 1 \\ \frac{1}{e} + \ln\left(\chi_1+\chi_2\right) & \text{if } 1 < (\chi_1+\chi_2) < \infty \end{cases} \tag{4.4.3.3}$$

While this proposal has much going for it, it has one marked demerit: it does not meet our requirements, for, while satisfying three of the conditions, it is not associative.[58] For example, for the values of three inaccuracies χ_1, χ_2 and χ_3 for which $\chi_1+\chi_2+\chi_3 < 1$,

$$\exp\left(\frac{-1}{\chi_1+\exp(-1/(\chi_2+\chi_3))}\right) \neq \exp\left(\frac{-1}{\exp(-1/(\chi_1+\chi_2))+\chi_3}\right)$$

I believe that, as I have posed it, the problem has no solution. I have not found a proof of the following conjecture (albeit, I have not yet had much time to look for one), but I am reasonably confident it is true.

Conjecture 4.4.3.1 *There is no $\alpha : \mathbb{R}^+ \times \mathbb{R}^+ \to \mathbb{R}^+$ simultaneously satisfying these conditions:*

1. *α is commutative: for every $r, s \in \mathbb{R}^+$, $\alpha(r,\, s) = \alpha(s,\, r)$*
2. *α is associative, in the sense that, for every $r, s, t \in \mathbb{R}^+$, $\alpha(r,\, \alpha(s,\, t)) = \alpha(\alpha(r,\, s), t)$*
3. *for every $r, s, s' \in \mathbb{R}^+$, if $\alpha(r,\, s) = \alpha(r,\, s')$, then $s = s'$*
4. *for every $r, s \in \mathbb{R}^+$ such that $r < s$, there exists a unique $t \in \mathbb{R}^+$ for which $\alpha(s,\, t) = r$*
5. *for every $r, s \in \mathbb{R}^+$, $\alpha(r,\, s) < \min\{r,\, s\}$*
6. *for every $r, s, t, u \in \mathbb{R}^+$ such that $r < t$ and $s \leq u$, $\alpha(r,\, s) < (t,\, u)$*

I have a feeling a proof could run along these lines: show that it follows from the first three conditions that one can construct a homeomorphism $\phi : (1,\, \infty) \to (1,\, \infty)$ such that $\phi(r) < r$ for all $r \in (1,\, \infty)$ and, if one restricts the action of α to the open interval $(1,\, \infty)$, then $\alpha(\phi(r),\, \phi(s)) = rs$; it would follow that α could not satisfy the fourth condition (much less the fifth, which I include only because it seems to me a condition one wants to demand of such a function), since $\phi(r) < r < rs$ and $\phi(s) < s < rs$, and so $\alpha(\phi(r),\, \phi(s)) < rs$. This line of argument suggests itself by dint of the fact that, if one restricts the domain to the open interval $(0,\, 1)$, then ordinary multiplication satisfies

57. If we were keeping explicit track of the sup-inf inaccuracy γ, we would have to multiply the exponential by a normalizing factor to ensure it approaches γ rather than 0, and perhaps also add γ to the first and $\gamma - 1/e$ to the second of the constituent functions, if $\gamma > 1/e$.

58. I do not think the discontinuity in the derivative at $\chi_1+\chi_2 = 1$ matters. In any event we can always multiply each constituent function by a smoothing factor in a small region encompassing the values for which $\chi_1+\chi_2 = 1$, to make the entire function smooth.

all the conditions, as it does on the domain $(1, \infty)$ as well so long as one reverses all the less-than signs and changes 'min' to 'max'. In any event, these seem to me the minimum conditions a generic, physical operation combining inaccuracies as we require should satisfy, and I can find no consistent way of defining a function that satisfies all the conditions (though, I emphasize, I also have not found a proof that none exists). If this conjecture is true, one could, if one liked, take it as one way to encapsulate precisely my earlier ruminations on the inexorably *ad hoc* and inexact character of such an operation, which will depend on the vagaries of the particular theory, system and experimental arrangement at issue.

We appear to have reached an impasse. All of our attempts to define a generic, physical operation on inaccuracies have come to naught. Indeed, I see only one way forward, and it requires yet another in an ever-growing list of approximations, fudges and hand-waving. Since I do not readily see how to pose the problem differently, we will have to do without a generic, physical operation on inaccuracies that is associative. I think the lack of associativity will turn out to be less of a problem than it may initially seem. I remarked earlier, after our first, naive attempt failed, that it is difficult to know how to proceed without associativity in the definition of other structures such as derivations and integrals. Difficult, yes, but I do not think impossible, at least not in practice. We want only a rule fine enough to guide us without ambiguity in our computations, which at the same time captures adequately the ideas drawn out in our discussion of Fermi problems. In this spirit, I offer the following proposal.

Definition 4.4.3.2 *A* compounding family $\mathfrak{F}$ *is a family of mappings* $\{\alpha_i\}_{i \in \mathbb{I}_2^{\uparrow}}$*, where* $\mathbb{I}_2^{\uparrow} = \{2, 3, \ldots\}$*, such that, for each* $n \in \mathbb{I}_2^{\uparrow}$*,*

1. $\alpha_n : \underbrace{\mathbb{R}^+ \times \cdots \times \mathbb{R}^+}_{n} \to \mathbb{R}^+$ *is continuous, surjective and totally symmetric*
2. *for every collection* $r_1, \ldots r_{n-1}, s, s' \in \mathbb{R}^+$*, if* $\alpha_n(r_1, \ldots r_{n-1}, s) = \alpha_n(r_1, \ldots r_{n-1}, s')$*, then* $s = s'$
3. *for every collection* $r_1, \ldots r_{n-1}, s \in \mathbb{R}^+$ *such that* $s < r_1, \ldots, s < r_{n-1}$*, there exists a unique* $t \in \mathbb{R}^+$ *for which* $\alpha_n(r_1, \ldots r_{n-1}, t) = s$
4. *for every collection* $r_1, \ldots r_n \in \mathbb{R}^+$ *,* $\alpha_n(r_1, \ldots r_n) < \min\{r_1, \ldots r_n\}$
5. *for every collection* $r_1, \ldots r_n, s_1, \ldots s_n \in \mathbb{R}^+$ *such that* $r_1 < s_1$ *and* $r_i \leq s_i$ *for* $i \in \{2, \ldots n\}$*,* $\alpha_n(r_1, r_2, \ldots r_n) < \alpha_n(s_1, s_2, \ldots s_n)$

We will refer to a member of such a family as a *compounder*, and, in particular, to a compounder taking n arguments as an *n-compounder*. To see how we would apply a compounding family in practice, take our earlier proposal, equation (4.4.3.3). For $(a_1, \chi_1), (a_2, \chi_2) \in \Re$, say for $(\chi_1 + \chi_2) < 1$, the combined inaccuracy is $\alpha_2(\chi_1, \chi_2) = \exp(-1/(\chi_1 + \chi_2))$ and, in general, the combined inaccuracies of n values, for $\sum_{i=1}^{n} \chi_i < 1$ will be $\alpha_n(\chi_1, \ldots \chi_n) = \exp\left(-1/\sum_{i=1}^{n} \chi_i\right)$. Although none of these functions is associative, the entire family does allow us to compute without ambiguity the result of any particular physical, inaccurate, algebraic operation; moreover, it will allow us to define the inaccurate analogue of partial derivatives and Lebesgue integrals on the inaccurate fields we will discuss in §4.4.4 below, analogously to the methods usually employed, using convergent approximations. In so far as we require only these operations, our definition will suffice. We therefore decree that the algebraic operation of combining inaccuracies is represented by a compounding family, the details of which will depend on the nature of the theory and the physical quantity at issue. We will, from hereon, not specify the exact form of the compounding family in play, using only the notation introduced in the definition when we need to refer to compounders in formulæ. We will use the following abbreviation for an exponentiated inaccuracy writing, *e.g.*, the "inaccurate square" of $\chi \in \mathbb{R}^+$, $\alpha_2(\chi,\, \chi)$, as 'χ^{α_2}'. We will also allow ourselves the occasional abuse of notation by writing such things as '$\alpha_2(\eta,\, \zeta)$', for $\eta, \zeta \in \Re$.[59]

I emphasize again that these seem to me only the most minimal conditions one would want to demand of such functions. I can easily imagine more that the character of a particular theory or quantity or experimental sitatuation may require. For example, it seems to me almost certain that a theory would have not a single compounding family but rather a family of such families, one for each possible algebraic combination—physical coupling—of the different quantities among themselves. In the case, say, of the equation of state for an ideal gas, the compounder one would use in computing the inaccuracy accruing to the algebraic product of the magnitude of the pressure and that of the volume may well differ from the compounder used in computing the inaccuracy accruing to the algebraic division of the magnitude of the temperature by that of the volume. The choice of operation, in this case—multiplication or division—indicates the nature of the measurements taken, whether one jointly measures the pressure and the volume in order to calculate the temperature of an ideal gas at equilibrium, or whether one jointly measures the temperature and the volume in order to calculate the pressure. Speaking more generally, it is easy to imagine that the compounder one uses to calculate the resultant inaccuracy after a dynamic process mediated by two separate quantities of the same system will differ from the compounder one uses for a dynamic process mediated by the self-interaction of one of those quantities (a non-linear process). In any event, since none of our arguments and results depends on the use of only a single compounding family for a given theory, we lose nothing by not taking account of such issues in what follows.

This leaves us still with the problem of combining the magnitudes of two inaccurate scalars by physical operations. Let us play for a moment with a toy model, to make the problem slightly more concrete. Fix a theory, a system modeled by that theory, and one of the system's quantities treated by the theory, without regard to the idiosyncratic character of any of the three. We want operations for algebraically combining and comparing the determined values of the magnitudes of the quantity, during the course both of the system's isolated dynamical evolution and of its interactions with other systems; we want these operations to be general enough, moreover, to apply when these magnitudes are determined with any of a number of different methods, perhaps depending on the application of different experimental techniques in different environmental circumstances, while the system is in markedly different states. Say we are to add the values $(100,\, 10e)$ and $(0.8,\, 0.1e)$, both in $\Re$, representing a physical magnitude of the given type of system, for which an additive operation seems required for the representation of an aspect of its physical character (perhaps this quantity satisfies a principle of linear superposition). According to our exemplary compounding family (as per equation (4.4.3.3)), no matter what form we settle on in the end for the additive operation on the magnitudes, the value of the resultant inaccuracy will be 1, which is greater than the magnitude of the second value. Should the result of adding the two, then, be $(100 + .8,\, 1)$? $(100,\, 1)$? Or something entirely different? And is it, in the end, reasonable to demand that the answer to this question not depend in any way on the nature of the theory, the system, the quantity or the experimental circumstance modeled?

It is a remarkable fact that all known physical quantities, in so far as they are modeled by physical theories, can find their mathematical representation among a narrowly circumscribed set of mathematical entities and that, correlatively, all known physical interactions can find their mathematical representation among a narrowly circumscribed set of algebraic operations on these sorts of entities. At bottom, for instance, the physical mixing or combination of two physical systems, no matter how exotic are the systems and no matter how exotic are the forces mediating the process, can be represented by the operation of a group. I believe that we often lose sight of how remarkable this fact

59. According to Pavlov, mixed conditioning, mostly positive with some negative thrown in at random, is the fastest type of habituation and leads to the most deeply ingrained habits. Perhaps the abuse's being merely occasional in this case will serve us well in the end.

is, our vision obscured by the very familiarity and seeming "naturalness" of the group-operation.[60] I can see no reason, *a priori* or otherwise, why matters stand thus, much less why they *must* do so, if indeed they must, in some sense of the term. Every sort of physical operation could have found its natural representation in a structure not translatable into the terms of any other also modeling a physical operation. This would entail no logical inconsistency. In any event, the actual state of affairs suggests that our search for a single operation or set of operations to represent the requisite form of the combination of magnitudes during the course of physical operations may not present so formidable a face as it first seemed to. If we can find a reasonable archetype for the elements of such a set of operations on the elements of $\Re$, enough to recapitulate, *mutatis mutandis*, the basic algebraic structures of $\mathbb{R}$ used in physics as ordinarily performed, I believe we will have done enough.

Our analysis of errors arising from our discussion of Fermi problems provides once again a clue to a way forward. As we remarked there, we expect in the long run that the errors in the determination of a quantity will distribute themselves evenly around the mid-point of the interval of possible inaccuracy, which is to say, around the magnitude, approximating to a Gaussian. In particular, this means that we expect, again in the long run, for the actual value of the quantity to lie half the time above the magnitude (the mid-point of the interval of possible inaccuracy), and half the time below. If you squint your eyes just right, it sort of follows from these considerations, in conjunction with our definition of a compounding family, and helped along by our faithful crutch, the demands of the nature of this project, that it would not be unreasonable to compute the value of the resultant magnitude under a physical operation by use of the ordinary, corresponding algebraic operation that $\mathbb{R}$ makes available to us, applied directly to the magnitudes of the elements of $\Re$ at issue. At least, I believe this will suffice for a first-order approximation, as it were, and will not lead us too grossly astray. Our proposed set of basic, physical, algebraic operations on $\Re$, then, are as follows.

$$\begin{aligned} (a,\,\chi) +_\phi (b,\,\psi) &= (a+b,\,\alpha_2(\chi,\,\psi)) \\ (a,\,\chi) -_\phi (b,\,\psi) &= (a-b,\,\alpha_2(\chi,\,\psi)) \\ (a,\,\chi) *_\phi (b,\,\psi) &= (ab,\,\alpha_2(\chi,\,\psi)) \\ (a,\,\chi)/_\phi(b,\,\psi) &= (a/b,\,\alpha_2(\chi,\,\psi)) \end{aligned} \tag{4.4.3.4}$$

Note that the algebraic operations applied to the magnitudes on the righthand sides of equations (4.4.3.4) are the ordinary ones from $\mathbb{R}$. We will call such an algebra an *inaccurate (scalar) algebra*.

Whatever else may be the case about these operations, we demand at a minimum that they be "as linear as they can". In this case, that means we require a group action on $\Re$ that interacts with the operations in the appropriate way. There at least three ways one may impose such an action, the first two by the group $\mathbb{R}$ and the third by the group $\mathbb{R}^+$: for $r \in \mathbb{R}$

$$r *_\phi (a,\,\chi) = (ra,\,\chi) \tag{4.4.3.5}$$

and

$$r *_\pi (a,\,\chi) = (ra,\,r\chi) \tag{4.4.3.6}$$

60. Rilke puts it finely:

und Das und Den
die man schon nicht mehr sah
(so täglich waren sie und so bewöhnlich)
auf einmal anzuschauen: sanft, versöhnlich
und wie an einem Anfang und von nah...

– from "Der Auszug des Verlorenen Sohnes" and for $r > \mathbb{R}^+$

$$r *_\psi (a,\, \chi) = (ra,\, r\chi) \tag{4.4.3.7}$$

The second corresponds to the pragmatic operation of multiplying an element of $\Re$ by a real number r, which we will make some use of below. It is linear over the pragmatic operations (*i.e.*, the ordinary algebraic operation on $\mathbb{R}$, applied component by component to elements of $\Re$) in the ordinary sense. The third corresponds to the psychological operation of rescaling one's units, and we will not bother with it further. The first is the physical one, used, for instance, when, in calculating the kinetic energy of a particle from its mass and velocity, one multiplies the product of the mass and the square of the velocity by 1/2: this represents the calculation of a physical quantity from one already given, and so represents a physical operation, as the notation suggests, but not one that increases or decreases the possible inaccuracy in any way, since the operation represents no process by which there could have been an increment or decrement in physical knowledge; thus this group-action does not affect the inaccuracy. it is "linear" in the sense that, for all $(a,\, \chi), (b,\, \psi)$ and $r \in \mathbb{R}$,

$$\begin{aligned} r *_\phi [(a,\, \chi) +_\phi (b,\, \psi)] &= (ra + b,\, \alpha_2(\chi,\, \psi)) \\ &= (ra + rb,\, \alpha_2(\chi,\, \psi)) \\ &= (ra,\, \chi) +_\phi (rb,\, \psi) \\ &= r *_\phi (a,\, \chi) +_\phi r *_\phi (b,\, \psi) \end{aligned} \tag{4.4.3.8}$$

I will call this property *inaccurate (physical) linearity*; we will say that any space having the structure of inaccurate algebra with such a group action satisfying inaccurate linearity defined on it an *inaccurately linear space.* From hereon, I will drop the subscripted 'ϕ', *etc.*, from the signs denoting algebraic operations, as context should disambiguate the sort of operation denoted by a given sign.

Having a notion approximating to a linear group action by $\mathbb{R}$ suggests the possibility of having a norm on $\Re$ as well. Again, the discussion of inaccuracy in the light of Fermi problems points to a natural way of imposing one. We want all the values in the interval of possible inaccuracy to contribute in some way or other to the value of the norm, but not all equally, in so far as the values furthest from the mid-point are, we posit, the least likely to occur. In the long run we expect the errors, in the determination of the magnitude of a quantity, more or less to distribute themselves evenly around zero, approximating to a Gaussian. The notion that, in the long run, the actual magnitudes will tend correlatively to distribute themselves in a Gaussian around the mid-point of that interval suggests that, to compute the norm of an element of $\Re$, we integrate over the interval using a Gaussian-weighted measure. There are many ways of doing this. The one I propose seems to me simple, clear, and not devoid of physical content. Given $(a,\, \chi) \in \Re$, its norm will have a form something like

$$||(a,\, \chi)|| \equiv \frac{1}{\nu(a,\, \chi)} \int_{(a-\frac{\chi}{2}}^{a+\frac{\chi}{2})} \frac{-y^2(y-a)^2}{2\left(y - (a - \frac{\chi}{2})\right)\left(y - (a + \frac{\chi}{2})\right)}\, \mathrm{d}y \tag{4.4.3.9}$$

where $\nu(a,\, \chi)$ is a normalizing factor that guarantees the value of the integral shrinks smoothly to a as the interval itself shrinks to zero (note that $||(a,\, \chi)|| > |a|$), and the open parenthesis prepended to the lower-limit of the integral and that appended to the upper limit jointly indicate that the integral is to be taken over the open interval rather than the closed one. [61] If one likes, this represents the "expectation value" of the quantity. It is straightforward to show that this mapping satisfies the definition of a norm (using the physical group-action posited above), *i.e.*, for all $(a,\, \chi), (b,\, \psi) \in \Re$ and $r \in \mathbb{R}$,

1. $||(a,\, \chi)|| \geq 0$

61. More precisely, the integral should be taken over the closed interval $[a - \frac{\chi}{2} + \epsilon,\, a + \frac{\chi}{2} - \epsilon]$, where $\epsilon < \frac{\chi}{2}$, and then the limit of this integral taken as $\epsilon \to 0$.

2. $||r(a,\,\chi)|| = |r|\,||(a,\,\chi)||$
3. $||(a,\,\chi)|| + ||(b,\,\psi)|| \geq ||(a,\,\chi) + (b,\,\psi)||$

This norm induces the same topology as does the metric Δ. (To see this, note that there is a homeomorphism h of $\Re$ into itself such that $\Delta((a,\,\chi),\,(b,\,\psi)) \mapsto ||h((a,\,\chi) - (b,\,\psi))||$.)

This mapping, strictly speaking, satisfies the letter of the definition of a norm, but does not seem to exemplify its spirit. It fails only in so far as there is no element in $\Re$ whose norm is 0: for example, $||(a,\,\chi) - (a,\,\chi)|| = ||(0,\,\chi^{\alpha_2})|| > 0$. This may seem problematic, but I think it makes physical sense. Let us say that the subtraction in this case represents the difference in values of a particular quantity associated with two numerically distinct but otherwise identical physical systems. This difference will be zero only inaccurately, as it were, in so far as there is a non-zero inaccuracy accruing to the magnitude in the determination of each of the two values. This norm will indeed approach arbitrarily closely to zero in the limit as the inaccuracy shrinks to zero, but it will never make it there, as the inaccuracy will never itself be zero. We will work around this issue in the following way. Equation (4.4.3.9) will remain our "official" definition of the norm on $\Re$, the one we will refer to and exploit when we need explicit use of one; in secret, however, we will know that any quantity whose sup-inf inaccuracy is γ has as its "real" norm

$$||(a,\,\chi)||_\gamma \equiv ||(a,\,\chi)|| - ||(0,\,\gamma)|| \tag{4.4.3.10}$$

where the norm-signs on the righthand side of the formula refer to that defined by equation (4.4.3.9).

In the end, one ought to have no illusions about the adequacy of this treatment of the algebraic structure of inaccuracy and error as they appear in all their multifarious roles in physics; it is only a crude, and still very nuch naive, treatment, but one, I hope, that suffices for the aims of this paper.

4.4.4 Inaccurate Scalar Fields and Their Derivations

Let us call any mapping that has $\Re$ as its range an *inaccurate field*, and in particular one whose domain is (a subset of) spacetime an *inaccurate scalar field*. We will sometimes refer to ordinary scalar fields on spacetime as *exact scalar fields*, to emphasize not only the difference, but the fact as well that inaccurate scalar fields constitute a certain sort of approximation to ordinary scalar fields, which, in the limit as the inaccuracy goes to zero, converges to an ordinary scalar field. The asymmetry in the naming reflects the fact that I do not want to seem to have a bias in favor of the theoretical structures by bestowing on them the honorific 'accurate'. [62]

We want to define the analog of fields of compact support. Since the second component of $\Re$, $\mathbb{R}^+$, has no natural additive identity e, we cannot define the support of an inaccurate field to be the set of points at which the value of the field equals $(0,\,e)$, as we would otherwise naturally do. Let's consider how such a thing as the idea of the support of an inaccurate field would be used in practice, then. For all intents and purposes, inaccurate scalar fields used to model the values of physical quantities in a $\mathfrak{K}$-appropriate observatory will have well defined values only in (some subset of) the spatiotemporal region representing the observatory. Outside that region, the values of the fields are assumed to be negligible with regard to the dynamic evolution of the fields in that region. As we remarked earlier, in practice every physical quantity, as modeled by a specific theory with a regime, will have associated with it a sup-inf inaccuracy $\gamma > 0$. We stipulate, then, that outside the region of the observatory, that physical quantity be represented by the value $(0,\,\gamma) \in \Re$. It is therefore natural to define the support of an inaccurate field ζ, Supp$[\zeta]$ to be the closure of the set of points at which it takes the value $(0,\,\gamma)$. We will also say that ζ is *inaccurately zero* outside its support. Thus, we can restate the definition: the support of an inaccurate field is the closure of the

62. It is well to keep in mind that 'exact' may be a term of derogation as well as of approbation. I invite the reader to recall the last sciolistic, pedantic lecture he or she has heard, in all its brilliant and empty exactitude.

complement of the set of points at which it is inaccurately zero. More generally, we say that an inaccurate scalar field is *inaccurately constant* if its value at every point is (k, γ), for fixed $k \in \mathbf{R}$. It is worth keeping in mind that the notion of being inaccurately zero, and so that of the support of an inaccurate scalar field, makes sense only in so far as one fixes such a γ. We will not bother doing so explicitly in what follows, as it is easy enough to do so as the requirements of the case at issue warrant; indeed, as we have already remarked, doing so would complicate the exposition needlessly, with no corresponding gain in perspicuity.

Let, then, $\Sigma_{\Re}$ be the space of inaccurate scalar fields of compact support on $\mathcal{M}$. We need deal only with fields of compact support in virtue of definition 4.3.3.2, that of a $\mathfrak{K}$-appropriate observatory. In order to define the analogue of derivations on inaccurate fields, we need a class of operators on $\Sigma_{\Re}$, analogous to the linear ones on Σ, to consider. $\Sigma_{\Re}$ as a whole inherits the algebraic structure of an inaccurately linear space from $\Re$, just as Σ inherits a linear structure from $\mathbb{R}$. By dint of the topology and differential structure of $\Re$, moreover, inaccurate scalar fields have natural notions of continuity, n-times differentiability, and smoothness accruing to them. Denote the subspaces of $\Sigma_{\Re}$ comprising only inaccurate scalar fields having those properties by '$\Sigma^0_{\Re}$', '$\Sigma^n_{\Re}$' and '$\Sigma^\infty_{\Re}$' respectively. Each of these is clearly an inaccurately linear space as well. They have as well a natural notion of boundedness, in virtue of the norm on inaccurate scalars defined by equation (4.4.3.9). An inaccurate scalar field is bounded if and only if the supremum of the norms of the values it takes at all points of its domain is finite. This supremum in turn defines a norm on $\Sigma^b_{\Re}$, the space of bounded inaccurate scalar fields, the analogue of the sup-norm for exact scalar fields.

Let us say, then, that an operator $\mathbf{\Gamma} : \Sigma_{\Re} \to \Sigma_{\Re}$ is *inaccurately linear* if it respects the inaccurately linear structure on $\Sigma_{\Re}$ respectively. In more detail, an operator $\mathbf{\Gamma}$ is inaccurately linear if, for $\zeta, \eta \in \Sigma_{\Re}$ and $r \in \mathbb{R}$,

$$\mathbf{\Gamma}(r\zeta + \eta) = r\mathbf{\Gamma}\zeta + \mathbf{\Gamma}\eta$$

Note that we use the physical group-operation for multiplication by r, that defined by equation (4.4.3.5). We impose an inaccurately linear structure on the space of all inaccurately linear operators in the standard way, by, *e.g.*, defining the addition of two of them, $\mathbf{\Gamma} + \mathbf{\Psi}$ by its action on inaccurate scalar fields: $(\mathbf{\Gamma} + \mathbf{\Psi})\zeta = \mathbf{\Gamma}\zeta + \mathbf{\Psi}\zeta$.

Turning now to differential operators in particular, the analogy with ordinary differential structure on real manifolds suggests that we define a smooth, inaccurate vector-field ξ^A on $\mathcal{M}$ to be an inaccurately linear operator on $\Sigma^\infty_{\Re}$, satisfying a few collateral conditions. We continue to deal only with physical operations unless explicitly stated otherwise. In particular, the derivations we define are those appropriate for use in physical computations, not pragmatic. By parity of reasoning, the differential, pragmatic, inaccurate operations are as trivial to define as were the algebraic ones. If we rely on the analogy with ordinary differential structure in defining derivations, we want the analogy to go as deep as it can, as it were. We would like, *inter alia*, to be able to associate with an inaccurate vector at a point a curve (with a fixed parameterization) passing through that point: the curve to which the vector is tangent. Standard treatments deal with this by fixing a chart around the point, pushing the field down to $\mathbb{R}^n$ via the chart, finding the curve such that the derivative of the field with respect to its affine parameter equals the action on the field by the vector on the manifold, and pulling the curve back up to the manifold by the chart. The derivative of a field on $\mathbb{R}^n$ with respect to the affine parameter of a curve is defined, when using charts, by one's taking the limit of the difference in values of the field at the point of the curve in question and at a neighboring point on the curve, dividing the difference by the affine distance between the two points, and taking the limit as the distance goes to zero.

We want to define the analogous operation on inaccurate scalar fields by use of a similar procedure. Say we are to compute the directional derivative of the inaccurate scalar field ζ along the curve η at the point q, where the affine parameter of η is s_q. Fix a chart (U, ψ) such that $q \in U$. What is the physical content of such an operation? Naively, one may picture it something like this. How ever we end up characterizing an inaccurate derivative, we expect it will consist of an ordered pair, the first component of which is something like an exact tangent vector, and the second component a representation of the inaccuracy accruing to the determination of that exact tangent vector. Computing the derivative of a quantity in the laboratory generally involves making (at least) two measurements of that quantity very close to each other in spacetime, taking the difference and dividing the magnitude by the separation of the events of measurement. The representation of physically combining these two inaccuracies comes precisely to computing the value of the two by the operation of a compounder. This suggests that we employ an operation as straightforwardly analogous to the ordinary directional derivative as possible, something like

$$\begin{aligned}\left.\frac{\mathrm{d}(\zeta\circ\eta)\,(s)}{\mathrm{d}s}\right|_{s=s_q} &\equiv \lim_{h\to 0}\frac{\zeta\circ\psi^{-1}(\psi\circ\eta\,(s_q+h))-\zeta\circ\psi^{-1}(\psi\circ\eta\,(s_q))}{h}\\ &= \lim_{h\to 0}\frac{\zeta\circ\eta\,(s_q+h)-\zeta\circ\eta\,(s_q)}{h}\end{aligned} \tag{4.4.4.1}$$

In this case, since the physical, subtractive operation acts on the inaccuracy by application of a compounder, that part of the difference in the numerator will not tend to zero but rather to

$$\lim_{h\to 0}\alpha_2(\pi_2\circ\zeta\circ\eta(s_q+h),\,\pi_2\circ\zeta\circ\eta(s_q))=(\zeta\circ\eta(s_q))^{\alpha_2} \tag{4.4.4.2}$$

where, recall, '$(\zeta\circ\eta(s_q))^{\alpha_2}$' denotes the inaccurate square of $\pi_2\circ\zeta\circ\eta(s_q)$. It follows that the division by h must use the physical operation of the group $\mathbb{R}$ on $\Re$.

Several manifest difficulties attend on this way of doing it, in virtue of the fact that the inaccuracy at a point accruing to the inaccurate tangent vector, as defined by this method, is, in essence, a scalar element of $\Re$, for it is the same for each component of the directional derivative at that point, and, indeed, the same for the computation of any directional derivative of the inaccurate scalar field at that point. It follows that the space of inaccurate tangent vectors at a point, according to this method, is a 5-dimensional, inaccurately linear space, each element of which consists in effect of an ordered pair the first component of which is an exact tangent vector and the second an element of $\mathbf{R}^+$. This fact raises three puzzles. First, it seems as though the inaccuracy accruing to a determination of the magnitude of the directional derivative of a quantity may depend on the direction along which one makes the measurements. Second, and on a related note, what may it mean to represent the inaccuracy accruing to the determination of a vectorial quantity by a scalar? Of what exactly is it the inaccuracy in the measurement of?

Let us try a second proposal for the inaccurate directional derivative, to address these questions. The awkwardness in the first way arose almost entirely from the fact that the limit's definition in equation (4.4.4.1) ensured that the value of the limit at a point for the inaccuracy depended on nothing else but the value of the inaccuracy at that point, and so the computation yielded a scalar to represent the inaccuracy of a vectorial quantity. The spirit rather than the letter of the ordinary, exact operation of taking a directional derivative points to perhaps the simplest way to avoid this problem,

$$\pi_1\circ\left.\frac{\mathrm{d}(\zeta\circ\eta\,(s))}{\mathrm{d}s}\right|_{s=s_q} \quad\equiv\quad \lim_{h\to 0}\pi_1\circ\frac{\zeta\circ\eta\,(s_q+h)-\zeta\circ\eta\,(s_q)}{h} \tag{4.4.4.3}$$

$$\pi_2\circ\left.\frac{\mathrm{d}(\zeta\circ\eta\,(s))}{\mathrm{d}s}\right|_{s=s_q} \quad\equiv\quad \lim_{h\to 0}\pi_1\circ\frac{\zeta\circ\eta\,(s_q+h)-\zeta\circ\eta\,(s_q)}{(\zeta\circ\eta(s_q))^{\alpha_2}+h} \tag{4.4.4.4}$$

The inaccuracy yielded by this computation manifestly depends on the direction along which the derivative is taken, and does so in a natural way. We thus obtain a vectorial kind of quantity for the total inaccuracy as determined using an orthonormal quadruplet of tangent vectors—the *vectorial* *inaccuracy*—, making the space of inaccurate tangent vectors at a point an 8-dimensional, inaccurately linear space (four dimensions for the magnitude and four for the inaccuracy associated with this vectorial magnitude, one member of $\mathbb{R}^+$ for each component of the magnitude, the inaccuracy in the direction of the coordinate-axis the component of the magnitude was computed for). Thus, as with inaccurate scalar fields, inaccurate tangent vectors have no natural, additive identity.

This formulation of the directional derivative presents its own set of difficulties, primary among them the question whether or not it has physical content relevant to our project. Looking only at the math suggests that to measure a vector is to fix a structure encoding a general rule for determining the rate of change of any given scalar quantity, no matter what that quantity may be, as measured along a particular spatiotemporal direction. [63] From a physical point of view, such a goal is nonsense. When measuring vectorial quantities, as, for example, the electric 4-current in special relativity, one attempts to determine the particular rate of change of a particular scalar quantity—in this case, an electric charge-density—in a particular spatiotemporal direction. There is no sane way to derive from the numbers resulting from such a measurement, or from the methods employed in coming to them (the use of galvonometers, *etc.*), the sort of generic rule the math encodes. [64] The measurement of the 4-current, nonetheless, exemplifies the application of such a rule. Although the tools and techniques one may use to determine the spatiotemporal rate of change of any particular scalar quantity will depend on the nature of the quantity, the schema, as it were, of the determination remains the same: make as many measurements as one can, along the line that one wants to determine the spatiotemporal rate of change of the quantity, as close to the point of interest as possible, and grind through a computation of the standard form (equation (4.4.4.1)). Perhaps the most striking fact about the nature of the information one needs to give content to this schema is how basic it is—one need know only the differential structure of the spacetime manifold to compute the directional derivative of a scalar field at a particular point. Computing that value, however, is not the end-all, be-all of physics. One wants to compute the total rate of change of the quantity, itself a vectorial quantity, and one wants to be able to compare in a meaningful way the magnitude of this vectorial quantity with that of others of the same type. To perform these operations one needs, at a minimum, knowledge of the affine structure of spacetime, and, in general, that of the metric structure. We do not ordinarily need to invoke this knowledge explicitly, as we are almost never in a position requiring fine knowledge of the affine structure for the planning of a measurement—"flat" is almost always an excellent approximation—which, perhaps, is why we rarely realize the nature of the operation we are performing. It thus becomes clear that, in measuring a vectorial quantity such as the electric charge-current, we are not attempting to abduct from the results of the measurement the structure of a general rule; we are rather applying an already known general rule to a particular case. This constitutes the physical content in the application of equation (4.4.4.3).

The question about the physical content of equation (4.4.4.4) remains. To investigate it, let us try to refine somewhat our example of the measurement of a vectorial quantity in the laboratory. Let us say that, for whatever reason, we are attempting to measure the directional derivative of a scalar quantity along two different lines at the same point, on the first of which, as one moves away from the point of measurement, the inaccuracy sharply and monotonically increases, whereas

63. We do not consider the measurement of vectorial amplitude-fields representing fundamental particles (*e.g.*, photons) in quantum field theory. It is not clear to me whether or not this treatment can be extended to treat that case. I suspect not.

64. *Cf.* Born (1943, p. 39): "I cannot see what experimental 'operation' could be devised in order to define a mathematical operator." Compare also Eddington (1923, p. 120–1):

> If we are to surround ourselves with a perceptual world at all, we must recognize as substance that which has some element of permanence. We may not be able to explain how the mind recognizes as substantial the world-tensor $[R_{ab} - \frac{1}{2} g_{ab} R]$, but we can see that it could not well recognize anything simpler. There are no doubt minds which have not this predisposition to regard as substantial the things which are permanent; but we shut them up in lunatic asylums.

on the second of which it sharply and monotonically decreases. The question whether or not to use equation (4.4.4.4) in calculating the inaccuracy of each directional derivative reduces in this case to the question whether the inaccuracies accruing to the measurement of both ought to differ from each other in accordance with our general rule, as the magnitudes do. It seems in fact they ought to. To measure the quantity along the line of increase, one will perform a sequence of ever more inaccurate measurements as one moves farther from the point. *Prima facie*, a greater inaccuracy will accrue to the total, resulting determination of the quantity along this line than will accrue to that along the line of decreasing inaccuracy, and, indeed, the more quickly the inaccuracy increases or decreases along these lines, the greater the difference of the two should be. Considerations of these sort justify the use of equation (4.4.4.4) in determining the inaccuracy accruing to the measurement of the directional derivative of a scalar quantity.

There remains another problem concerning the physical cogency and possible significance of the idea of a vectorial inaccuracy itself. The inaccuracy of an inaccurate scalar serves to define a bounded region of the real line containing what we have referred to as the magnitude of the inaccurate scalar, that region in which we have reason to believe the "actual" value of the quantity being measured lies. In what way, if at all, may a vectorial inaccuracy define an analogous region around the exact vector constituting the magnitude of an inaccurate vector? The answer seems clear enough on the face of it. It seems that a vectorial inaccuracy may, in one sense, be considered nothing more than an exact vector with strictly positive components in all coordinate systems. This suggests that we take the convex hull in the exact tangent vector-space determined by the magnitude of an inaccurate vector and by its vectorial inaccuracy, considered, as suggested, as an ordinary exact vector. As inviting as this sounds, at this point the suggestion can not even be wrong. It does not make enough sense to be wrong. First of all, even if we could make sense of thinking of a vectorial inaccuracy as an exact vector, such a convex hull would be only two-dimensional. We would expect, however, that, in so far as the inaccuracy of a vectorial quantity includes uncertainty about its direction, the possible directions in which it may point subtend a non-trivial, three-dimensional solid angle in spacetime (as determined by the ambient spacetime conformal structure). If we also assume, as seems not unreasonable, that the inaccuracy in the determination of the direction of the vector cannot be so severe as to make it possible that the real value of the vector points in exactly the opposite direction from the determined magnitude, and, moreover, that it must be such as to permit the real value of the vector to lie only in that half-space of the whole tangent space bounded on one side by the three-dimensional hypersurface orthogonal to the determined magnitude and containing that magnitude, then it follows that the entire region in which we have reason to believe the actual value of the vector lies forms something like a four-dimensional cone, with its vertex consisting of the point lying in the orthogonal hypersurface at the foot of the vector representing the determined magnitude. I say "something like" a cone, because, if the possible inaccuracy does not permit the real value of the vector to be zero, then we will end up with a topological 4-sphere of exact vectors. In any event, we cannot make sense of thinking of a vectorial inaccuracy as an exact vector: it obeys utterly different transformation laws. We will not be able to address adequately the issue of the physical congency and significance of a vectorial inaccuracy until §4.4.5.

This discussion also points to a subtle, at this point strictly mathematical problem with the characterization of the inaccuracy of tangent vector-field thus far, considered as an inaccurately linear operator on $\Sigma_{\Re}^{\infty}$. It is clear how to deal with the compounding of these vectorial inaccuracies when considering the sum of two inaccurate tangent vector-fields: we side-step the issue, noting that the sum of two inaccurately linear operators on $\Sigma_{\Re}$ is just that one defined by applying each summand separately to the argument of the sum and summing the two resultant inaccurate scalar fields. Because we already know how to compound inaccuracies for inaccurate scalar fields, this presents no problem. We define the compounded inaccuracy of the summed inaccurate tangent vector-fields to be the second component of that inaccurate tangent vector-field that acts in the way defined by the sum. It is not clear, however, how one is to compound the inaccuracies when one multiplies an inaccurate tangent vector by an inaccurate scalar, as we surely will want to do, for example, when calculating the static Coulomb force on a charged particle in a central field by multiplying the value of the charge by the value of the Coulomb field at its position. A compounding family as we have characterized it will not serve the purpose: in so far as it is not clear what one may mean by comparing a scalar inaccuracy to a vectorial inaccuracy using the "less-than" relation, as required by items 3, 4 and 5 of definition 4.4.3.2, we have no way of defining a compounder to meet this need. We will postpone this discussion as well, and its resolution, until §4.4.5 below.

We need now demonstrate only that our proposal satisfies a Leibniz rule in order to declare it an appropriate representation of the directional derivative, modulo the difficulties we have postponed. This is easily done, in the same way as in the exact case.

Definition 4.4.4.1 *A* smooth, inaccurate, tangent vector-field *on* $\mathcal{M}$ *is an inaccurately linear operator* $\xi^A : \Sigma^\infty_\Re \to \Sigma^\infty_\Re$ *satisfying the Leibniz rule: for* $\phi, \chi \in \Sigma^\infty_\Re$

$$\xi^A(\phi\chi) = \phi\xi^A(\chi) + \chi\xi^A(\phi)$$

In the same way as in the exact case, one can as well characterize these vectors by a slightly more general characteristic, that of being *inaccurately affine*, in the sense that their action on $\Sigma^\infty_\Re$ is determined only up to the addition, to the operand, of a constant inaccurate scalar field. As in the definition, we will indicate the indexical structure of these objects using the abstract-index notation of Penrose and Geroch (see, *e.g.*, Wald (1984) for an account of the notation), with majuscule indices. Exact tangent vectors and tangent vector-fields will be denoted as well using the abstract-index notation, with miniscule indices, *e.g.*, 'ξ^a'. We extend to inaccurate tangent vector-fields (and, later, to higher-order tensorial and affine objects) the action of our projection operators π_1 and π_2 in the obvious way: for every $q \in \mathcal{M}$, $\pi_1 : T_{[q,\Re]}\mathcal{M} \to T\mathbb{R}^4$ and $\pi_2 : T_{[q,\Re]}\mathcal{M} \to T(\mathbb{R}^+)^4$ are, respectively, projection on the first and second components of elements of $T_{[q,\Re]}\mathcal{M}$. Note, finally, that there will be, in general, an endless family of inaccurate tangent vector-fields the actions of which on the same inaccurate scalar field all agree on the first component of their respective, resultant inaccurate scalar fields.

The smooth, inaccurate tangent bundle on $\mathcal{M}$, $T_\Re\mathcal{M}$, is constructed in the usual way from these fields. Though this bundle is analogous in many ways to that of the ordinary, exact tangent bundle over a real manifold, there are important disanalogies as well. First of all, $T_\Re\mathcal{M}$ is 12-dimensional rather than 8-dimensional, as its fibers themselves are 8-dimensional, diffeomorphic to $\mathbb{R}^4 \times (\mathbb{R}^+)^4$. We will write the fiber over the point $q \in C$ as '$T_{[q,\Re]}\mathcal{M}$'. As with $\Sigma^\infty_\Re$, $T_\Re\mathcal{M}$ has no distinguished zero cross-section. Still, there are important analogies. It is easy to see, for instance, that if the group-action on the fibers is topologically trivial then $T_\Re\mathcal{M}$ is the trivial bundle, consisting of the topological product of the base space by the fiber.[65] Thus, non-trivial, global cross-sections do exist for such inaccurate tangent bundles. One can always impose an orientation on $T_\Re\mathcal{M}$, moreover, in terms of the inaccurate structures, in the person of a non-zero, inaccurate 4-form, as in the ordinary, exact case (see §4.4.5 for a sketch of a characterization of inaccurate differential forms and tensors in general). Let $\mathcal{T}^{1,0}_\Re$ be the space of smooth, global sections of $T_\Re\mathcal{M}$ (*i.e.*, of smooth, inaccurate tangent vector-fields), and the space of sections of the exact tangent bundle '$\mathcal{T}^{1,0}$'.

4.4.5 Inaccurate Tensorial Fields and Their Derivations

We want to construct inaccurate tensorial spaces of all orders and indexical structures in, again, the usual way, by marching up the ranks of indices, as it were, starting with the definition of

65. More precisely, if the group-action on the fibers of the inaccurate tangent bundle arises from a group that is *solid* (for a precise characterization of which, see, *e.g.*, Steenrod (1951, §12.1)), then the bundle space is trivial.

cotangent vector-fields as inaccurately linear operators on inaccurate tangent vector-fields, and so on. We know the general form we want an inaccurate tensor of a given indexical structure to have: a first component consisting of an ordinary, exact tensor of the given indexical structure, and a second component consisting of the inaccuracy that, in some way or other, accrues to the measurement of this exact tensor. Recall the considerations that led us to take equations (4.4.3.4) as the definition of an inaccurate scalar algebra, in particular how we arrived at the form the operations should take when restricted to the first component, that in the long run the errors should more or less wash out and we should end up with the magnitude one would have gotten by applying the ordinary operations to the magnitudes in the first place. I believe these same considerations are as suitable (or not) here as there, and so we conclude that, when applying any sort of algebraic operation to an ordered set of inaccurate tensors (how ever we end up defining these things in full)—whether it be contravection on multiple indices of multiple inaccurate tensors, or multiplication of an inaccurate tensor by an inaccurate scalar, or contraction of indices on a single inaccurate tensor, or what have you—the calculation of the resultant first component, the magnitude of the resultant inaccurate tensor, will be independent of the calculation of the resultant second component, the inaccuracy, and will be, moreover, the result of applying the algebraic operation to the first components of the inaccurate tensors in the ordered set.

The only delicacy in the process lies in characterizing the way the inaccuracies combine under these algebraic operations, and characterizing, indeed, what form the inaccuracies should take in general for tensorial objects. We will take a cue from our treatment of inaccurate tangent vector-fields as inaccurate operators on $\Sigma^{\infty}_{\Re}$, for one can demand that, as in the exact case, inaccurate cotangent vector-fields ought to be inaccurately linear operators on inaccurate tangent vector-fields, and inaccurate tensorial fields of arbitrary indexical structure ought to be inaccurately linear operators on ordered sets of inaccurate tangent and cotangent vector-fields. We will not go deeply into the details here (I suspect you know by now where to find those), limiting ourselves rather to a brief sketch.

After the statement of definition 4.4.3.2 we remarked that different interactions mediated by different quantities, as treated by the same theory, likely would require the analogue of different compounding families for calculating the inaccuracies resulting from such different interactions. In so far as each index of a multi-index tensor potentially represents a physical interaction of a sort different from those represented by the other indices—or, if you like, represents "half" of such an interaction—the second component of an inaccurate tensor, that representing the inaccuracy accruing to the determination of the value of a quantity the tensor models, will in general have associated with it something like a family of compounding families, one family for each way of contravecting one of the tensor's indices with the indices of other objects representing quantities the theory models. In particular, this shows that, as with inaccurate tangent vectors, the second component of an inaccurate tensor will in general be a tensor-like quantity—the *tensorial inaccuracy*, as we will call it—with the same indexical structure as the first component of the inaccurate tensor, its magnitude. Note that, as an ordinary, exact tensor-space is isomorphic as a vector-space to $\mathbb{R}^n$ for some $n \in \mathbb{I}+\uparrow$ (though not naturally so), a space of tensorial inaccuracies, all of the same indexical structure, is diffeomorphic to $(\mathbb{R}^+)^n$ for some $n \in \mathbb{I}+\uparrow$ as well (though, again, not naturally so).

Now, in light of our considerations after equation (4.4.4.4) about the way to handle the mathematical issue of compounding the innacuracies when one adds two inaccurate tangent vector-fields, we know already how to compound the inaccuracies when adding two inaccurate tensor-fields of arbitrary indexical structure, in so far as we consider those fields to be inaccurate operators over ordered sets of inaccurate tangent and cotangent vector-fields (over every such set, to be more precise, the contravection of whose elements, term by term, with the indices of the given inaccurate tensor-field will saturate the indices of the tensor-field, yielding an inaccurate scalar field). In fact, we now know even a little more, for one can clearly use the same techniques to define the inaccuracy accruing to the result of a contraction of two indices on an inaccurate tensor. The problem we postponed in the same discussion, though, that of the proper way to compound the inaccuracies when multiplying an inaccurate tangent vector-field by an inaccurate scalar field, not only remains, but has become aggravated by the introduction of new algebraic operations we wish to impose on our inaccurate quantities, to wit, tensor-products and contravections. This observation suggests that the proper way to treat the compounding of inaccuracies for tensorial quantities is to extend definition 4.4.3.2. Items 3, 4 and 5, however, prevent us from performing such an extension in any more or less obvious, straightforward way, in so far as no obvious analogues of the "less-than" relation suggests itself for tensorial objects. We need a way to apply something like the "less-than" relation not only to ordered pairs consisting of a scalar and a vectorial inaccuracy, but as well to ordered pairs consisting of tensorial inaccuracies of arbitrary indexical structure. We will circumvent this problem by the introduction of a device that, at the moment, will likely appear purely formal and mostly *ad hoc*, but will prove itself, in §4.5.2 below, to have, under an appropriate interpretation, a use of real physical significance integral to the completion of this project. We will use this device here to extend the notion of a compounding family to cover the compounding of tensorial inaccuracies, which will complete our account of the fundamentals of inaccurate tensorial fields. The device consists of the imposition of a norm on tensorial inaccuracies, the values of which can be compared using the ordinary "less-than" relation.

We begin by imposing a norm on the space $\mathcal{T}^{1,0}_{\mathfrak{R}}$ of inaccurate tangent vector-fields. We cannot define this norm with respect to the length of vectors as determined by a distinguished inaccurate semi-Riemannian metric, as, at this point, we have no such distinguished metric (indeed, we do not at this point, strictly speaking, know what such a thing is). Fix some $0 < k < \infty$, for $k \in \mathbb{R}^+$. Let $\Sigma^{\infty}_{\mathfrak{R},k}$ be the subspace of the space $\Sigma^{\infty}_{\mathfrak{R}}$ consisting of all smooth, inaccurate scalar fields uniformly bounded by k, *i.e.*, all those fields contained in the ball in $\Sigma^{\infty}_{\mathfrak{R}}$ of radius k, as determined using the sup-norm. More precisely, let $\Sigma^{\infty}_{\mathfrak{R},k}$ be the interior of this ball. In other words, no field in $\Sigma^{\infty}_{\mathfrak{R},k}$ has a norm of k; rather all are strictly less than k. Let us say, then, that an inaccurate tangent vector-field ξ^A of compact support is *k-bounded* if

$$\sup_{\zeta \in \Sigma^{\infty}_{\mathfrak{R},k}} \left\{ ||\xi^A(\zeta)|| \right\} < \infty \tag{4.4.5.1}$$

(where $||\xi^A(\zeta)||$ is the sup-norm of the inaccurate scalar field resulting from the application of ξ^A to ζ). The value of this supremum for ξ^A is its k-norm, which we will write '$||\xi^A||_k$', to emphasize its dependence on k. Fix now $\mathcal{T}^{1,0}_{\mathfrak{R},k}$, family of smooth, inaccurate tangent vector-fields uniformly (and strictly) bounded by k with respect to the k-norm. Again, this family constitutes the interior of the ball in $\mathcal{T}^{1,0}_{\mathfrak{R}}$ of radius k with respect to the k-norm. By considering inaccurate cotangent vector-fields to be inaccurately linear operators on $\mathcal{T}^{1,0}_{\mathfrak{R}}$, we can now define the family of k-bounded inaccurate cotangent vector-fields by repeating essentially the same procedure, define its k-norm, and so define the subset $\mathcal{T}^{0,1}_{\mathfrak{R},k}$ of $\widehat{\mathcal{T}}^{0,1}_{\mathfrak{R}}$ consisting of the open ball of radius k as defined by its k-norm.[66] Marching up the ranks of indices in the usual way yields a k-norm on the space of inaccurate tensorial fields of any indexical structure (m, n), defined as inaccurately linear operators on ordered sets consisting of n inaccurate tangent and m inaccurate cotangent vector-fields. We will speak, therefore, of k-bounded inaccurate tensor-fields promiscuously, irrespective of their exact indexical structures.

Recall, moreover, that, as part of the definition of a kinematical regime, we demanded that the values of the fields and of some number of their "partial-derivatives" be uniformly bounded. k-norms provide the means for making this requirement precise. An inaccurate field ζ is *uniformly k-bounded*

66. Note that this account so far makes sense in the terms of our previous definitions and arguments, for we know how to characterize the compounding of inaccuracies for inaccurately linear operators whose range is $\Sigma_{\mathfrak{R}}$, even though we do not yet know how to do so for inaccurately linear operators with ranges other than $\Sigma_{\mathfrak{R}}$.

to first-order if $\zeta \in \Sigma^{\infty}_{\Re,k}$ and

$$\sup\left\{||\xi^N(\zeta)|| : \xi^A \in \mathfrak{T}^{1,0}_{\Re,k}\right\} \; < \; \infty \tag{4.4.5.2}$$

Similarly, ζ is *uniformly k-bounded to second-order* if

$$\sup\left\{||\eta^M(\xi^N(\zeta))|| : \eta^A, \xi^A \in \mathfrak{T}^{1,0}_{\Re,k}\right\} \; < \; \infty \tag{4.4.5.3}$$

and so on.[67] These suprema are the values to be used in determining whether or not the fields satisfy the regime's constraints on the boundedness of "partial-derivatives" of admissible fields. Note that this is a different question from whether or not the field itself satisfies some particular differential constraint. This is rather an algebraic constraint on the derivatives of the field. We will see at the end of this section how to extend the notion of -$^{\text{th}}$order k-boundedness to inaccurate tensorial quantities.

This gets us closer to what we want, but we have not yet arrived. We require, at the moment, a norm on the tensorial inaccuracy, not on the inaccurate tensor-field as a whole. In fact, we can extract the appropriate norm on the inaccuracies by extending our projection operators π_1 and π_2, in the obvious way, to inaccurate tensorial objects. To begin, we will use $\pi_2[\Sigma^{\infty}_{\Re,k}]$ for this rather than $\Sigma^{\infty}_{\Re,k}$, *i.e.*, we will use the space of exact scalar fields composed of the fields of scalar inaccuracies of those fields in $\Sigma^{\infty}_{\Re,k}$. Given $\zeta \in \Sigma^{\infty}_{\Re,k}$, for example, the corresponding field $\zeta_2 \equiv \pi_2(\zeta)$ in $\pi_2[\Sigma^{\infty}_{\Re,k}]$ is that defined by assigning to the spatiotemporal point q in the domain of ζ the value $\zeta_2(q)$. We can use the ordinary sup-norm for this space. Consider now the space of vectorial inaccuracies $\pi_2[\mathfrak{T}^{1,0}_{\Re,k}]$. We define $\xi^A_2(\zeta_2)$, the derivation of ζ_2 by $\xi^A_2 \in \pi_2[\mathfrak{T}^{1,0}_{\Re,k}]$, by $\pi_2(\xi^A(\zeta))$, where $\zeta \in \Sigma^{\infty}_{\Re,k}$ and $\xi^A \in \mathfrak{T}^{1,0}_{\Re,k}$ are such that $\zeta_2 = \pi_2(\zeta)$ and $\xi^A_2 = \pi_2(\xi^A)$. By construction, for any $\xi^A_2 \in \pi_2[\mathfrak{T}^{1,0}_{\Re,k}]$,

$$\sup_{\zeta_2 \in \pi_2[\Sigma^{\infty}_{\Re,k}]} \left\{||\xi^A_2(\zeta_2)||\right\} < \infty$$

(where $||\xi^A_2(\zeta_2)||$ is the sup-norm of that exact scalar field). The value of this supremum for ξ^A_2 is its k-norm, which we will, again, write '$||\xi^A_2||_k$', to emphasize its dependence on k. We extend this norm to tensorial inaccuracies of arbitrary indexical structure in exactly the same way as we did for the norms on inaccurate tensor-fields. This construction, in fact yields a family of norms on k-bounded, exact tensorial fields in general.

We are now in a position to characterize the generalization of compounding families to inaccurate tensors. Consider the uses a 2-compounder of tensorial inaccuracies may be put to. Since we need to know how to compound pairs of such inaccuracies with arbitrary combinations of indexical structures, we will need a separate 2-compounder for each possible combination—one, for example, for compounding a scalar and a vectorial inaccuracy, as well as one for compounding a $(3,4)$-tensorial inaccuracy with a $(4398,9)$-tensorial one. Only so much will still not suffice, for we will need separate ones for, *e.g.*, the contravection of a $(3,4)$-tensorial inaccuracy with a $(4398,9)$-tensorial one in all possible contravectional combinations, and a separate one for taking the tensor-product of them. To simplify matters a little, when considering all possible contravectional combinations of two types of tensorial inaccuracy, we will require separate compounders only for all possible resulting tensorial inaccuracies having numerically distinct indexical structures. In our example of contravecting a $(4398,9)$-tensor with $(3,4)$-tensor, for instance, we would require only two separate compounders to deal with the cases where the contravections yielded, say, a $(4394,6)$-tensor and a $(4399,11)$-tensor respectively. We will not require one for every possible way of contravecting each index on the one with each index on the other so as to yield a tensor of the resultant indexical structure.

67. In fact, all these notions are most clearly, usefully and elegantly expressed in terms of jets and their inaccurate, colored analogues, but we do not have the time or the space to rehearse such a discussion.

This leaves us still with the need for an enumerably infinite number of 2-compounders: one for each possible ordered triplet of pairs of indices such that tensors of the indexical types represented by the first two pairs in the triplet can be contravected so as to yield one of a type represented by the third pair; and one for each possible ordered triplet of pairs of indices such that the tensor-product of tensors of the indexical types represented by the first two pairs in the triplet yields one of a type represented by the third pair. Let us call such a triplet of ordered pairs an *indexically possible triplet*. Note that, in attempting to describe all the possible combinations for an n-compounder, for $n > 2$, we would need to consider not indexically possible triplets but rather indexically possible $(n+1)$-uplets. A single definition of a compounding family that attempted to cover all this ground in one go—simultaneously defining n-compounders for all n, for all indexically possible $(n+1)$-uplets—would be all but incomprehensible. I will therefore offer a definition only for a 2-compounder of tensorial inaccuracies. The extension to compounders taking any number of arguments should then be clear, though tedious to construct explicitly.

Fix a complete family of k-bounded subsets of inaccurate tensorial fields of all ranks. Let $E \equiv \{((m_i, n_i),\, (p_i, q_i),\, (r_i, s_i))\}_{i \in \mathbb{I}^\uparrow}$ be an enumeration of indexically possible triplets. We will write, for example, the second ordered pair in the n^{th} item in the enumeration as '$E(n, 2)$', and a tensor space having this indexical structure as '$\mathfrak{T}^{E(n,2)}$'.

Definition 4.4.5.1 *A* k-bounded family of 2-compounders $\mathfrak{F}_{k,2}$ *is a family of mappings* $\{\alpha_{2,i}\}_{i \in \mathbb{I}^\uparrow}$*, such that, for each* $n \in \mathbb{I}^\uparrow$*,*

1. $\alpha_{2,n} : \mathfrak{T}^{E(n,1)}_{\Re,k} \times \mathfrak{T}^{E(n,2)}_{\Re,k} \to \mathfrak{T}^{E(n,3)}_{\Re,k}$ *is continuous, surjective and totally symmetric*
2. *for every* $\lambda \in \mathfrak{T}^{E(n,1)}_{\Re,k}$ *and* $\mu, \mu' \in \mathfrak{T}^{E(n,2)}_{\Re,k}$*, if* $\alpha_{2,n}(\lambda,\, \mu) = \alpha_{2,n}(\lambda,\, \mu')$*, then* $\mu = \mu'$
3. *for every* $\lambda \in \mathfrak{T}^{E(n,1)}_{\Re,k}$ *and* $\nu \in \mathfrak{T}^{E(n,3)}_{\Re,k}$ *such that* $||\nu||_k < ||\lambda||_k$ *there exists a unique* $\mu \in \mathfrak{T}^{E(n,2)}_{\Re,k}$ *for which* $\alpha_{2,n}(\lambda,\, \mu) = \nu$
4. *for every* $\lambda \in \mathfrak{T}^{E(n,1)}_{\Re,k}$ *and* $\mu \in \mathfrak{T}^{E(n,2)}_{\Re,k}$*,* $||\alpha_{2,n}(\lambda,\, \mu)||_k < \min\{||\lambda||_k,\, ||\mu||_k\}$
5. *for every* $\lambda, \lambda' \in \mathfrak{T}^{E(n,1)}_{\Re,k}$ *and* $\mu, \mu' \in \mathfrak{T}^{E(n,2)}_{\Re,k}$ *such that* $||\lambda||_k < ||\lambda'||_k$ *and* $||\mu||_k \leq ||\mu'||_k$*,* $||\alpha_{2,n}(\lambda,\, \mu)||_k < ||\alpha_{2,n}(\lambda',\, \mu')||_k$

We will refer to a member of such a family as a *k-bounded, tensorial 2-compounder*. A collection of such families for all $n \in \mathbb{I}^\uparrow_2$ is a *k-bounded, tensorial compounding family*, $\mathfrak{F}_k$.

With this in hand, we now know how to characterize contravection and tensor-products for inaccurate tensor-fields, by analogy with equations (4.4.3.4). It is tempting straightaway to define a tensorial algebra on k-bounded, inaccurate tensorial fields, in the obvious way, but this will not quite work, for these spaces are not closed under the considered algebraic operations, as, for example, the sum of two k-bounded tangent vector-fields, say, is not itself necessarily k-bounded, nor is the tensor-product of two k-bounded tensors, or their contravection. The countably tensorial product of all these spaces is, however, convex, in the sense that, *e.g.*, for any two k-bounded vector-fields ξ^A and η^A, and any $r \in [0,\, 1]$,

$$r\xi^A + (1-r)\eta^A$$

is itself k-bounded, with the analogous statement holding for, respectively, multiplication by an inaccurate, k-bounded scalar field, contravection, contraction and the tensor-product on finite numbers of inaccurate tensors. For example, for any two k-bounded vector-fields ξ^A and η^A, and any $r \in [0,\, 1]$, the tensor-product

$$\frac{r}{k}\xi^A \otimes \frac{1}{k}\eta^A$$

is k-bounded. We will refer to operations of this form as *k-convex*. This suggests

Definition 4.4.5.2 *The* convex algebra of k-bounded, inaccurate, tensorial fields *over a differential manifold* $\mathcal{M}$ *is an ordered pair* $(\mathfrak{T}_{\mathfrak{R},k},\, \mathfrak{F}_k)$ *consisting of the* k*-convex tensor-product of all* k*-bounded, inaccurate, tensorial spaces, with the algebraic structure imposed on it by the family of* k*-convex operations, and a* k*-bounded, tensorial compounding family.*

There are some similarities with the ordinary, exact tensorial algebras, such as the following

Proposition 4.4.5.3 *For every* $m, n \in \mathbb{I}^{\uparrow}$*, an inaccurate tensor-field of indexical structure* $(m,\, n)$ *is in* $\mathfrak{T}^{m,n}_{\mathfrak{R},k}$ *if and only if it can be expressed as an inaccurately linear sum of tensor-products of* m *inaccurate* k*-bounded tangent and* n *inaccurate* k*-bounded cotangent vector-fields.*

This follows from the compactness, connectedness and convexity of the space underlying the algebra.

We can recapitulate all these definitions and arguments to construct a true, inaccurately linear algebra (*i.e.*, one closed under all algebraic operations), by restricting attention to uniformly bounded inaccurate scalar fields rather than restricting ourselves to k-bounded fields. We then characterize a set of inaccurate tangent vector-fields as those satisfying the analogue of equation 4.4.5.1 for uniformly bounded inaccurate scalar fields. Call the norm so defined the Σ-*norm* and the space of such inaccurate tangent vector-fields Σ-*bounded, normed, inaccurate tangent vector-fields*, $\mathfrak{T}^{1,0}_{\mathfrak{R}}$. In order to define algebraic operations on this space, we generalize definition 4.4.5.1 in the obvious way to handle Σ-bounded rather than k-bounded entities. It is then easy to see that $\mathfrak{T}^{1,0}_{\mathfrak{R}}$ is closed under addition, as well as under multiplication by uniformly bounded inaccurate scalar fields, and so is a true inaccurately linear space. One now marches up the ranks of indices in the standard way, using our generalized family of tensorial compounding families, leading to the Σ-*bounded, normed, inaccurately linear tensor-algebra*, $(\mathfrak{T}_{\mathfrak{R}},\, \mathfrak{F})$, an algebraically complete, inaccurately linear tensor-algebra over the space of uniformly bounded inaccurate scalar fields.

These constructions allow us now to address the issues we raised in §4.4.4, following our proposal of equation (4.4.4.4), about the physical cogency and possible significance of tensorial inaccuracies. We want to know whether we can understand the tensorial inaccuracy of an inaccurate tensor as determining a topological 4-sphere within which lies not only the determined magnitude of the inaccurate tensor, but within which as well we have reason to believe the real value of the tensor lies. For the sake of simplicity, we will work with the Σ-bounded, inaccurate, tensor-algebra. Fix $\lambda \in \mathfrak{T}^{m,n}_{\mathfrak{R}}$, with determined magnitude $\lambda_1 = \pi_1(\lambda)$ and tensorial inaccuracy $\lambda_2 = \pi_2(\lambda)$. Then the *4-sphere of possible values* for the quantity represented by λ is defined as the ball of radius $\frac{1}{2}||\lambda_2||$ in $\mathfrak{T}^{m,n}_{\mathfrak{R}}$ centered on λ. For small enough $||\lambda_2||$, where λ is, say, an inaccurate tangent vector, this may allow the real value of λ to point in the opposite direction as λ_1, but, this does not to be objectionable, in so far we this will, in general be possible only for very small vectors, where such a possibility does not seem far-fetched.

Before moving on, it will be instructive to consider in some detail the construction of the analog of the Lie derivative, as a derivation on the Σ-bounded algebra of inaccurate tensor-fields. Fix a smooth inaccurate vector-field ξ^A on C. We need first to characterize the analogue of integral curves for it. We will choose the simplest analogue to the exact case, declaring that the integral curve of ξ^A is the integral curve of the exact tangent vector-field associated with it, $\pi_1(\xi^A)$. As always, there are many ways one could do this, some more involved than others. This one, with its simplicitly and physical content, suits our purposes. We can associate with ξ^A a family of diffeomorphisms $\{\xi_h\}_{h\in\mathbb{R}^+}$ of spacetime onto itself, the "flow" of the vector-field, in the standard way (we assume without further comment that ξ^A is complete, at least in the canvas C). Each of these diffeomorphisms define a new inaccurate tangent vector-field $\xi_h \circ \eta^A$ from a given one η^A by dragging its values along the integral curves of ξ^A a given distance with respect to the parametrization of the integral curves (in this case, a distance of h). Then it is easy to see that, for any smooth inaccurate tangent vector-field, the following limit is defined without ambiguity and exists, and so defines the first component of a new inaccurate tangent vector-field $\pounds_{[\xi,\Re]}\,\eta^A$,

$$\pi_1 \circ \pounds_{[\xi,\Re]}\,\eta^A \equiv \lim_{h\to 0} \frac{1}{h}\,\pi_1\left(\eta^A - \xi_{-h}\circ\eta^A\right) \tag{4.4.5.4}$$

The same considerations as led us to choose equation (4.4.4.4) over equation (4.4.4.1) for the definition of the inaccurate directional derivative imply that we cannot define the second component of the inaccurate Lie derivative to be

$$\pi_2 \circ \pounds_{[\xi,\Re]}\,\eta^A \equiv \lim_{h\to 0} \frac{1}{h}\,\pi_2\left(\eta^A - \xi_{-h}\circ\eta^A\right) \tag{4.4.5.5}$$

as it will not depend on ξ^A. The proper limit will not be so easy to define as was that of equation (4.4.4.4), in particular what the divisor of the difference should be. As in equation (4.4.4.4), it seems likely that it should be $h + N(\xi^A,\,\eta^A)$, where $N : T_{[q,\Re]}\mathcal{M} \times T_{[q,\Re]}\mathcal{M} \to \mathbb{R}$ is a linear, normalizing function. Presumably, it will depend on some general characteristic of the way that $\xi_{-h}\circ\eta^A$ approaches η^A as h goes to zero. It is beyond the scope of this paper to consider ways of making this idea precise. We will assume, therefore, that the second component of the inaccurate Lie derivative is given by

$$\pi_2 \circ \pounds_{[\xi,\Re]}\,\eta^A \equiv \lim_{h\to 0} \frac{1}{h + N(\xi^A,\,\eta^A)}\,\pi_2\left(\eta^A - \xi_{-h}\circ\eta^A\right) \tag{4.4.5.6}$$

Note that, if one is keeping explicit track of the sup-inf inaccuracy, (4.4.5.6) would have to be modified as follows:

$$\pi_2 \circ \pounds_{[\xi,\Re]}\,\zeta \equiv \lim_{h\to 0} \frac{1}{h + N(\xi^A,\,\eta^A)}\,\pi_2\left(e^{\gamma}\left(\eta^A - \left(\xi_{-h}\circ\eta^A\right)\right)\right) \tag{4.4.5.7}$$

Again, one must keep in mind the primary difference between this inaccurate Lie derivative and the ordinary Lie derivative, to wit, that the inaccurate Lie derivative is only inaccurately linear, not fully linear. As an example of the differences consider, for $\gamma > 0$, the inaccurate Lie derivative of a constant inaccurate scalar field, $\zeta \in \Sigma_{\Re}$. For any smooth, inaccurate vector-field $\xi^A \subset T\mathcal{M}_{\Re}$, then, $\pounds_{[\xi,\Re]}\zeta$ equals the constant inaccurate scalar field whose value at every point is $(0,\,\gamma)$, not $(0,\,0)$, the latter not even being an element of $\Re$. The two are very much analogous in other ways, however. For instance, for $\zeta \in \Sigma^{\infty}_{\Re}$, $\pounds_{[\xi,\Re]}\,\zeta = \xi^A(\zeta)$, and, for another $\omega \in \Sigma^{\infty}_{\Re}$, $\pounds_{[\xi,\Re]}\,(\zeta\omega) = \zeta\pounds_{[\xi,\Re]}\,\omega + \omega\pounds_{[\xi,\Re]}\,\zeta$. It defines a Lie algebra as well. Note that the inaccurate Lie derivative of a Σ-bounded, inaccurate tensor-field may not itself be Σ-bounded, so care must be taken when one attempts to algebraically combine the result of a Lie-derivation of an inaccurate tensor-field with a Σ-bounded, inaccurate tensor-field, as the operation may not be well defined.

We can use the inaccurate Lie-derivative to extend the idea of n^{th}-order k-boundedness from inaccurate scalars field to inaccurate tensorial quantities. An inaccurate tensorial quantity λ is *uniformly k-bounded to first-order* if $\lambda \in \mathfrak{T}^{m,n}_{\Re,k}$ and

$$\sup\left\{||\pounds_{[\xi,\Re]}\,\lambda)|| : \xi^A \in \mathfrak{T}^{1,0}_{\Re,k}\right\} \;<\; \infty \tag{4.4.5.8}$$

One then continues taking suprema of Lie-derivatives to define those quantities uniformly k-bounded to higher orders.

The inaccurate, covariant derivative operator can also be defined in close analogy with that in the exact case. It will be simplest to use a construction analogous to that of the Koszul connection in defining it.[68]

68. See, *e.g.*, Spivak (1979b, ch. 6)

Definition 4.4.5.4 *Let $\eta^A, \zeta^A, \xi^A \in \mathfrak{T}^{1,0}_{\Re}$ and $\zeta \in \Sigma^{\infty}_{\Re}$. An* inaccurately linear connection *on $\mathcal{M}$ is an inaccurately linear operator $\nabla_A : \mathfrak{T}^{1,0}_{\Re} \times \mathfrak{T}^{1,0}_{\Re} \to \mathfrak{T}^{1,0}_{\Re}$ that satisfies*

1. $(\eta^N + \zeta^N)\nabla_N \xi^A = \eta^N + \nabla_N \xi^A + \zeta^N \nabla_N \xi^A$
2. $\eta^N \nabla_N (\zeta^A + \xi^A) = \eta^N \nabla_N (\zeta^A + \xi^A)$
3. $(\zeta\eta^N)\nabla_N \xi^A = \zeta(\eta^N \nabla_N \xi^A)$
4. $\eta^N \nabla_N (\zeta\xi^A) = \zeta\eta^N \nabla_N \xi^A + (\eta(\zeta))\xi^A$

where '$\eta^N \nabla_N \xi^A$' represents the value of the operation applied to $(\eta^N,\, \zeta^N) \in \mathfrak{T}^{1,0}_{\Re} \times \mathfrak{T}^{1,0}_{\Re}$.

One can demonstrate the existence of such an operator by, *e.g.*, fixing a chart, defining the analogue of an arbitrary collection of Christoffel symbols and defining ∇_A in their terms as usual. The action of this connection can be extended to inaccurate scalar fields in the usual way. This operator has the same sort of similarities and dissimilarities with the analogous, exact operator as does the inaccurate Lie derivative with its exact counterpart. Note, again, that the inaccurate covariant derivative of a Σ-bounded, inaccurate tensor-field may not itself be Σ-bounded, so care must be taken when one attempts to algebraically combine the result of such a derivative with a Σ-bounded, inaccurate tensor-field, as the operation may not be well defined.

4.4.6 Inaccurately Linear Operators

We now generalize the results of §§4.4.4 and 4.4.5 by charactering inaccurately linear operators in the abstract, and a type of stability we may demand of them. We will proceed in the standard fashion. Let $\mathbf{O}$ be an inaccurately linear operator from any normed, inaccurately linear space Λ_1 to another Λ_2. We say $\mathbf{O}$ is *bounded* just in case, over any set in Λ_1 bounded in norm, the supremum of the norms of its values is bounded: for every bounded set $L \subset \Lambda_1$ there is a c_k such that

$$\sup_{x \in L} \{||\mathbf{O}(x)||\} < c_k$$

In virtue of the inaccurate linearity of the spaces, this is equivalent to: just in case the supremum of the norms of its values on the closed ball of radius 1 in the norm on Λ_1 is itself bounded. As in the exact case, an inaccurately linear operator's being bounded implies that it is continuous with respect to the topologies induced on the domain and the range by their respective norms. Let B_1 be the ball of radius 1 (with respect to the norm) in Λ_1. The *operator-norm* of a bounded operator $\mathbf{O}$ is

$$||\mathbf{O}|| \equiv \sup_{x \in B_1} \{||\mathbf{O}(x)||\} \tag{4.4.6.9}$$

It is easy to see that, in virtue of the inaccurate linearity of the spaces, as in the exact case, this is equivalent to

$$||\mathbf{O}|| \equiv \sup_{x \in \Lambda_1 : ||x|| \neq 0} \left\{ \frac{||\mathbf{O}(x)||}{||x||} \right\} \tag{4.4.6.10}$$

We call the topology induced on the space of inaccurately linear operators by this norm the *operator-topology*. We say a subset of this space is *uniformly bounded* if it consists only of bounded operators the supremum of the bounds of which is finite.

Definition 4.4.6.1 *An inaccurately linear operator is* stable *if it has a non-trivial, uniformly bounded neighborhood, and it is ω-stable if it has a uniformly bounded neighborhood containing an open ball of radius ω in the operator-norm.*

Although we defined stability and ω-stability with respect to the operator-norm in particular, it is clear that we may define the same notions with respect to any norm we may impose on the space of bounded operators. As a related notion, we lay down

Definition 4.4.6.2 *A* bounded perturbation *of a stable operator is a second operator contained in the interior of a non-trivial neighborhood of compact closure of the first; an ω-*bounded perturbation *of an operator is a bounded perturbation of it contained in an open ball of radius ω in the operator-norm.*

There follows trivially

Proposition 4.4.6.3 *Every bounded (respectively: ω-bounded) operator has a bounded (respectively: ω-bounded) perturbation.*

It will be useful for future purposes, before leaving the topic, to record one more result. Impose the topology on $\mathfrak{T}^{0,1}_{\Re,k}$ induced by its norm. Because we have defined elements of this space as inaccurately linear functionals on ordered sets consisting of n inaccurate tangent and m inaccurate cotangent vector-fields, we can apply to them the notions of stability and ω-stability as expressed in terms of the imposed norm and topology. There follows from the convexity of the space

Theorem 4.4.6.4 *Every element of $\mathfrak{T}^{m,n}_{\Re,k}$, for every $m, n \in \mathbb{R}^{\uparrow}$, an considered as inaccurately linear functional on ordered sets consisting of n inaccurate tangent and m inaccurate cotangent vector-fields, is stable.*

4.4.7 Integrals of Inaccurate Fields, and Topologies on Their Spaces

We now turn to treat integrals on $\Re$. By pressing the same sort of analogy as we used in defining inaccurate derivations, we want to define inaccurate integrals as a species of inaccurately linear functional, $\mathbf{T} : \Sigma_{\Re} \to \Re$, satisfying a few collateral conditions. As before, we continue to deal only with physical operations unless explicitly stated otherwise. In particular, the integrals we define are those appropriate for use in physical computations, not pragmatic. By parity of reasoning, the integral, pragmatic, inaccurate operations are as trivial to define as were the algebraic ones.

Let us write, no matter how we end up defining it, the inaccurate integral of the inaccurate scalar field ζ over the canvas C as

$$\int_{[C,\Re]} \zeta \, \mathrm{d}\hat{\mu}$$

where $\mathrm{d}\hat{\mu}$ is whatever measure-like structure we end up using. By the same sort of reasoning that led to equation (4.4.4.3), we may conclude that the first component of the value of this operation ought to be the ordinary Lebesgue integral of the first component of ζ (the magnitudes of the values of the field) with respect to the ordinary Lebesgue measure defined by the volume-element ϵ_{abcd} associated with the spacetime's metric:

$$\pi_1 \circ \int_{[C,\Re]} \zeta \, \mathrm{d}\hat{\mu} \equiv \int_C \pi_1 \circ \zeta \, \epsilon_{abcd}$$

This seems all right so far. As always, the trouble enters when trying to deal with the inaccuracy.

It will not do to define the second component of the integral as the Lebesgue integral of the second component of the field, the inaccuracy, considered as an exact scalar field in its own right. If this is to be a physical operation, then we expect the second component of the result to represent the inaccuracy associated with a measurement of the first component. The Lebesgue integral of the ordinary scalar field constituted by the values of the second component of the given inaccurate scalar field—the scalar field of inaccuracies, if you will—in so far as it combines the values in an alternating process of ordinary summations and limits, does not combine them in the proper way, which in this case must involve our inaccurate, physical, algebraic operation on inaccuracies, since we want to define a functional representing a physical operation. An obvious solution suggests itself:

define a variant of the Lebesgue integral by using, rather than ordinary addition, the operations given by our compounding family to combine the inaccuracies associated with all the vanishingly small regions. This makes physical sense, as the act of combining all the inaccuracies as determined in "infinitesimal" cells throughout the region of integration, those accruing to the measurements of the values of a quantity in all those cells, comes to the application of our family of compounders on all these inaccuracies. This is, in essence, how we will proceed in the end, but getting there requires that we first deal with one somewhat delicate problem regarding the convergence of the values of our compounders applied to a sequence of sets of inaccuracies.

This is the problem. Let ζ be a simple, inaccurate field over C, *i.e.*, one taking on only a finite number of different values $(a_1,\, \chi_1),\, \dots (a_n,\, \chi_n) \in \Re$. Let C_i be the subset of C on which the value is $(a_i,\, \chi_i)$. Then the proposed analogue to the Lebesgue integral of ζ over C, using the ambient spacetime volume element, is

$$\sum_{i=1}^{n} (a_i,\, \chi_i) \int_{C_i} e_{abcd} = \sum_{i=1}^{n} \left(\int_{C_1} a_1\, e_{abcd},\, \int_{C_1} \chi_1\, e_{abcd} \right)$$

where the multiplication of each element $(a_i,\, \chi_i)$ by $\int_{C_i} e_{abcd}$ uses the pragmatic group operation on $\Re$, and the summation over the inaccuracies uses our physical, algebraic operation. Thus, the value of the inaccuracy is

$$\alpha_n \left(\int_{C_1} \chi_1\, e_{abcd}\,,\, \dots \int_{C_n} \chi_n\, e_{abcd} \right)$$

This value has no ambiguity in its computation.

Consider the next step in defining the Lebesgue integral on more complex fields, extending this sum to countably simple fields, which take on only a countable number of different values, $\{(a_i,\, \chi_i)\}_{i \in \mathbb{I}^+}$. In this case, using the same notation, the value of the integral is

$$\lim_{n\to\infty} \sum_{i=1}^{n} (a_i,\, \chi_i) \int_{C_i} e_{abcd}$$

The value of the inaccuracy in this case is

$$\lim_{n\to\infty} \alpha_n \left(\int_{C_1} \chi_1\, e_{abcd}\,,\, \dots \int_{C_n} \chi_n\, e_{abcd} \right) \tag{4.4.7.1}$$

Our definition of a compounding family does not guarantee that this limit is unambiguously defined. Consider two different orderings of the countable number of values, $\{(a_i,\, \chi_i)\}_{i \in \mathbb{I}^+}$ and $\{({a'}_i,\, {\chi'}_i)\}_{i \in \mathbb{I}^+}$. We have no way of knowing whether or not

$$\lim_{n\to\infty} |\alpha_n\, (\chi_1,\, \dots \chi_n) - \alpha_n\, ({\chi'}_1,\, \dots {\chi'}_n)\, | = 0$$

If this does not hold for every pair of enumerations of the countable family of values, however, then the limit of the countable sum has no unambiguous definition, since its value will depend on the order in which the inaccuracies are "fed into it", as it were. We do know, however, that, for any particular ordering of our countable set of inaccuracies, the limit (4.4.7.1) does converge. Indeed, it is bounded from below by zero and from above by $\chi_{\inf} \displaystyle\int_C \epsilon_{abcd}$, where $\chi_{\inf}$ is the infimum of the set of inaccuracies. (The upper bound follows from item 4 in definition 4.4.3.2, and the fact that it converges and does not endlessly oscillate from item 5.) The general template of constructions in the theory of Lebesgue integration points to an appealing (though by no means the only) way forward from here. Fix a countably simple, inaccurate scalar field ζ over the canvas C, taking its values from $A = \{(a_i,\, \chi_i)\}_{i \in \mathbb{I}^+}$ on, respectively, the subsets $\{C_i \subset C\}_{i \in \mathbb{I}^+}$. Let $\mathfrak{A}$ be the class of all enumerations of the elements in A, and $({\chi'}_1, \ldots {\chi'}_n)$ be the ordered n-tuplet consisting of the inaccuracies of the first n elements of the enumeration $A' \in \mathfrak{A}$. Then

$$\sup_{A' \in \mathfrak{A}} \left\{ \lim_{n \to \infty} \alpha_n \left(\int_{{C'}_1} {\chi'}_1 \, e_{abcd} \, , \ldots \int_{{C'}_n} {\chi'}_n \, e_{abcd} \right) \right\} \tag{4.4.7.2}$$

exists.

It is tempting to take this as the value of the second component of the integral of ζ, but I do not think it is yet quite right. Let's say, for example, that we are integrating a continuous distribution of electric charge over a region of spacetime, to compute the total charge contained therein. To accord with the principles laid down so far, we will represent the charge-density by an inaccurate scalar field, defining a variable density of inaccuracy, if you will, representing at a point the inaccuracy associated with the determination of the magnitude of the charge density in some vanishingly small subset of our region containing that point (perhaps the variability in the inaccuracy comes from the fact that the measurements become more inaccurate in proportion to the distance of the small region from the probe measuring the charge). Taking the integral of this charge-density corresponds, physically, to adding up the values of the determined magnitudes of the charge in as many vanishingly small subsets of the region as one can. To compute the integrated inaccuracy using the formula (4.4.7.2) would imply that the inaccuracy associated with measuring a charge-distribution increases in proportion to the volume the charge occupies, if the magnitudes and inaccuracies themselves remain unchanged as one proportionately increases the volume occupied by the charge-distribution. As we remarked earlier, however, we are defining a physical operation, and so we expect the second component of the result to represent the inaccuracy associated with a measurement of the first component. In this light, it seems to me rather that, if one applies the same techniques and instruments of measurement to two charge-distributions differing only in the volume each occupies, then the inaccuracy will be roughly the same for both, perhaps even a little less, in general, for the determination of the larger charge, in so far as one may be able to take more measurements, using the same techniques and instruments, in the larger volume than in the smaller one.

To account for this, I propose that, in defining our limits and their supremum, we take the volume-weighted average to compute the integrated inaccuracy. Denote the volume of the region C by '$v[C]$', *i.e.*, $v[C] = \int_C \epsilon_{abcd}$.

Definition 4.4.7.1 *The* inaccurate integral *of the countably simple, inaccurate scalar field* $\zeta \in \Sigma_{\Re}$ *over the region* C *is fixed by the equations*

$$\begin{aligned} \pi_1 \circ \int_{[C,\Re]} \zeta \, d\hat{\mu} &= \int_C \pi_1(\zeta) \, \epsilon_{abcd} \\ \pi_2 \circ \int_{[C,\Re]} \zeta \, d\hat{\mu} &= \frac{1}{v[C]} \sup_{A' \in \mathfrak{A}} \left\{ \lim_{n \to \infty} \alpha_n \left(\int_{{C'}_1} {\chi'}_1 \, e_{abcd} \, , \ldots \int_{{C'}_n} {\chi'}_n \, e_{abcd} \right) \right\} \end{aligned} \tag{4.4.7.3}$$

One can now give a rigorous treatment of $\mathrm{d}\hat{\mu}$ as a kind of Stieltjes measure, properly modified so as to accord with the structure of inaccurate fields, and so treat the integral itself as a modified type of Lebesgue-Stieltjes integral, in close analogy with the standard techniques, extending its action to all the inaccurate, integrable scalar fields in $\Sigma_{\Re}$ so constructed.[69] We will write the space of inaccurate, integrable scalar fields as '$\mathcal{L}_1[C, \Re]$', the space of inaccurate, square-integrable scalar fields as '$\mathcal{L}_2[C, \Re]$', *etc.*, in order to distinguish them from the analogous spaces of exact scalar fields, which we denote as usual by '$\mathcal{L}_1[C]$', *etc.*

69. See Curiel (2011) for the technical details.

It will be convenient to impose a topology on $\mathcal{L}_1[C, \Re]$. There are several options to choose from, including the so-called strong and weak topologies (the analogues, for spaces of continuous operators, of the compact-open and the finite-open topologies for spaces of ordinary, continuous fields), that defined by uniform convergence (with respect to the norm on $\Re$),[70] and that defined by the natural $\mathcal{L}_1$-norm:

$$||\zeta||_1 \equiv \int_{[C,\Re]} ||\zeta|| \, \mathrm{d}\hat{\mu} \tag{4.4.7.4}$$

The topology defined by this last norm is coarser than that of uniform convergence ("more sets can be open when you account for the behavior of fields on sets of measure zero"), though if one considers the topologies induced by each on the space of $\mathcal{L}_1$ fields mod disagreement only on sets of measure zero, then they are the same. In fact, when employing the $\mathcal{L}_1$-topology, we will always assume the space at issue to be composed of equivalence classes of fields agreeing up to sets of measure zero.[71] Both are coarser than the strong and weak topologies ("more sets can be open when you account for the behavior of fields on proper subsets of its domain"). One can as well impose the sup-norm on $\Sigma^b_\Re$, the space of inaccurate scalar fields bounded with respect to the norm on $\Re$. Since the topology of $\Re$ satisfies the first axiom of countability, moreover, the sup-norm topology on $\Sigma^b_\Re$ is equivalent to the topology of uniform convergence (restricted to $\Sigma^b_\Re$). We can impose all these topologies on $\Sigma^b_\Re$, $\Sigma^0_\Re$, $\Sigma^n_\Re$, and $\Sigma^\infty_\Re$ in the usual way, by considering each as an open set in the respective larger family and using the restriction topology. We will require in this paper only the $\mathcal{L}_1$-topology and the sup-norm topology (when we restrict attention to $\Sigma^b_\Re$ and its open subsets).[72] It is useful to note that $\mathcal{L}_1[C, \Re]$ is a Banach space with respect to its norm, as is $\Sigma^b_\Re$, and that $\mathcal{L}_2[C, \Re]$ can be given, in the standard way, the structure of an inaccurate Hilbert space (*i.e.*, one that is in all ways like an ordinary Hilbert space, only having an inaccurately linear rather than a linear structure on its elements), by defining the norm to be

$$||\zeta||_2 \equiv \left(\int_{[C,\Re]} ||\zeta^2 \mathrm{d}\hat{\mu}\, || \right)^{\frac{1}{2}}$$

and the inner product

$$\langle \zeta,\, \eta \rangle \equiv \int_{[C,\Re]} \zeta \, \eta \, \mathrm{d}\hat{\mu}$$

Note that the inner-product, being a physical operation, takes its values in $\Re$ rather than in $\mathbb{R}$.

We are now in a position to state the following proposition, which sums up almost all the properties of inaccurate scalar fields required for their easy application to the problems we will treat.

Proposition 4.4.7.2 *$\zeta \in \Sigma_\Re$ is, respectively, $\mathcal{L}_1$, $\mathcal{L}_2$, bounded, continuous, n-times differentiable, or smooth if and only if each of its components, considered as an ordinary, exact scalar field in its own right, is, respectively, so with respect to the germane exact structure.*

The proof is straightforward, so I skip it.

Finally, our norm on $\Re$ can be used as well to define exact differential and integral operations on inaccurate scalar fields, which we will put to good use later. The definition of each such action

70. One ought to keep in mind that, given two sequences of inaccurate scalars, for the sequence formed by the norms of the differences of the scalars in the two sequences, taken in order, it must be the case that the magnitudes of the scalars in each sequence approach each other and that their inaccuracies separately approach the sup-inf inaccuracy, in order for the two sequences to converge to a common value.

71. Strictly speaking, equation (4.4.7.4) defines only a semi-norm, since it can happen that $||\zeta||_1 = 0$ even when ζ $emphNicomacheanEthics0$. With our present stipulation, that we deal only with equivalence classes of fields modulo disagreement on sets of measure zero, this becomes a true norm.

72. The strong and weak topologies are used in Curiel (2011) to construct inaccurate Sobolev spaces.

follows the same template. To apply an exact operator to an inaccurate field, one takes the norm of the value of the field at each point to form an exact scalar field and then applies the operator to the constructed field. For example, the action of an ordinary, exact tangent vector-field ξ^a on an inaccurate scalar field ζ is given by

$$\xi^a(||\zeta||)$$

that of the exact integral by

$$\int_C ||\zeta||\,\epsilon_{abcd}$$

and so on.

4.4.8 Colorings

We are now in a position to complete the definition of the kind of mathematical field required for use in modeling physical fields so as to conform to the requirements of a regime. Let $\mathcal{C}$ be the decoupage of a canvas C.

Definition 4.4.8.1 *An* inaccurate coloring *on* $\mathcal{C}$ *is a mapping* $\theta : \mathcal{C} \to \Re$.

An inaccurate coloring, then, is an inaccurate field with the decoupage as its domain.[73] An *infimal, inaccurate coloring* is one restricted to $\mathcal{C}^{\text{inf}}$, representing an instance of the finest possible specification of initial data conforming to the regime of a theory for the initial-value formulation of that theory. Similarly, an *exact coloring* is a mapping from $\mathcal{C}$ to $\mathbb{R}$ (*i.e.*, an ordinary, exact scalar field on the manifold defined by $\mathcal{C}$). When we speak of a 'coloring' without qualification, we should be understood to mean an inaccurate coloring. Let $\Theta_\Re$ be the space of colorings, and Θ that of exact colorings. Because $\mathcal{C}$ is compact and Hausdorff (by dint of the fact that C is compact and Hausdorff), all its subsets are of compact closure, and so all colorings on it have compact support. Let $\Theta^b_\Re$ be the space of bounded colorings (with respect to the norm on $\Re$), $\Theta^0_\Re$ that of continuous colorings (with respect to the topologies defined on $\mathcal{C}$ and $\Re$), and $\Theta^\infty_\Re$ that of smooth colorings (with respect to the differential structures on the spaces).[74]

In virtue of the differential structure naturally accruing to $\mathcal{C}$ as an 8-dimensional smooth manifold, inaccurately linear tangent and co-tangent vectors, as well as higher-order inaccurately linear tensors of arbitrary indexical structure, *etc.*, can be defined over colorings in the same way as was done for such structures over inaccurate scalar fields. To distinguish them from such structures over inaccurate scalar fields, we will use the adjective *colored* to qualify them, speaking, *e.g.*, of *colored, inaccurate tangent vector-fields*. We will represent these structures using, again, the Penrose-Geroch abstract-index notation, with upper-case Greek letters as indices, *e.g.*, ξ^Ω and ζ_Ω for a tangent and cotangent field respectively. Similarly, the *colored, inaccurate covariant derivative* of the coloring θ, for instance, will be denoted '$\nabla_\Psi\theta$'. The similar tensorial and affine structures over exact colorings will be denoted using lower-case Greek letters as indices, *e.g.*, ξ^ω and ζ_ω for a colored tangent and cotangent field respectively. The inaccurate integral of a coloring as well is defined in the same way as that of an inaccurate scalar field, with the difference consisting only in the measure used to compute each. We will not go into the technical details of the construction of the measure on decoupages, for which see Curiel (2011, appendix A). All the properties proved for the analogous

73. In Curiel (2011), for strictly technical reasons, a coloring is defined only on subsets of $\mathcal{C}$ such that the intersection of C with the spacetime region defined by the union of the scraps in that support (each scrap considered simply as a region of spacetime) is open and contains none of the boundary of C. We will not need to take account of this detail here. As always, for complete technical details, see Curiel (2011).

74. Since colorings purport to represent actual, inaccurate data conforming to the regime of a theory, it could be objected that many real fields, such as Maxwell fields, are never confined to compact support. At first blush, this idealization seems no worse (and no better) than the ones of this sort standardly employed in physics. In fact, I think it is better, in so far as we demand only that these fields be inaccurately zero, not exactly zero.

structures on inaccurate scalar fields carry over intact to those on colorings. Thus, in the same vein as proposition 4.4.7.2, one has

Proposition 4.4.8.2 *$\theta \in \Theta_{\Re}$ is, respectively, in $\mathcal{L}^1[\mathcal{C}, \Re]$, $\mathcal{L}^2[\mathcal{C}, \Re]$, Θ^b,* et al.*, if and only if each of its components, considered as an exact coloring in its own right, is, respectively, so.*

We can immediately extend the definitions of all the various k-bounded and Σ-bounded norms and structures defined in §4.4.5 to the analogous colored structures. We will, as usual, qualify these structures with 'colored' when context will not suffice to disambiguate the sense.

One class of colored, inaccurate tensorial and affine objects is of such importance and utility that we will explicate their construction. A smooth *colored, inaccurate semi-Riemannian metric* of Lorentz signature, say, $(+, -, -, -)$, is a symmetric, indefinite, invertible tensor $g_{\Psi\Omega}$, such that there exists a tetrad $\{\overset{0}{\xi}{}^{\Psi}, \overset{1}{\xi}{}^{\Psi}, \overset{2}{\xi}{}^{\Psi}, \overset{3}{\xi}{}^{\Psi}\}$ orthonormal in the sense that

1. $\pi_1 \circ (g_{\Psi\Omega} \overset{0}{\xi}{}^{\Psi} \overset{0}{\xi}{}^{\Omega}) = 1$ (*inaccurately timelike*)
2. $\pi_1 \circ (g_{\Psi\Omega} \overset{\rho}{\xi}{}^{\Psi} \overset{\rho}{\xi}{}^{\Omega}) = -1$, for $\rho \in \{1, 2, 3\}$ (*inaccurately spacelike*)
3. $\pi_1 \circ (g_{\Psi\Omega} \overset{i}{\xi}{}^{\Psi} \overset{j}{\xi}{}^{\Omega}) = 0$, for $i, j \in \{0, 1, 2, 3\}$ and $i \neq j$
4. $\pi_2 \circ (g_{\Psi\Omega} \overset{i}{\xi}{}^{\Psi} \overset{j}{\xi}{}^{\Omega}) = \gamma$, for $i, j \in \{0, 1, 2, 3\}$

Note that the indefiniteness of the metric appertains only to the first component of colored, inaccurate tangent vector-fields. If one wanted to have a more detailed and thorough treatment, one could demand, for example, that some smeared out average of $g_{\Psi\Omega} \overset{0}{\xi}{}^{\Psi} \overset{0}{\xi}{}^{\Omega}$ equal 1, but we do not need to go into such detail here; the reader is invited to supply such details on his or her own, following the schemata of methods employed earlier. One now defines the associated, inaccurate Levi-Cevita derivative operator, ∇_{Ψ}, in the usual way, as that unique torsion-free, affine connection with respect to which the metric is inaccurately zero. The proof that such a connection exists follows exactly the same line of reasoning as in the exact case (see, *e.g.*, the proof of theorem 2.2 in Kobayashi and Nomizu (1963, vol. 1, ch. 4)).

One oddity about such metrics must be pointed out: the act of lowering or raising an index of a tensor by its use tends to reduce the inaccuracy associated with the quantity represented by that tensor. I believe this makes some (perhaps not much, but some) physical sense. Say one lowers the index of an inaccurate tangent vector in order to contravect it with one of a given tensor, to represent the taking of the physical component of an index of that tensor at a point, in the spatiotemporal direction the inaccurate tangent vector determines. To do this, one must set up measuring devices along the line determined by the vector to take the directional derivative of the aspect of the quantity represented by the tensor's index, which process ought to reduce the uncertainty in the knowledge of the line along which the vector points.

It is interesting to note as an aside, moreover, that, although this treatment of colored, inaccurate metrics looks at first glance as though it could be used to render an inaccurate treatment of the spacetime metric and affine structure, it in fact cannot, in so far as we had to assume the existence of the exact spacetime metric in all these arguments and constructions, *e.g.*, in defining infimal decoupages. Indeed, it does not seem possible to use this kind of scheme to attempt to reconcile observed gravitational measurements of metric structure and curvature, with their associated, inevitable inaccuracies, with rigorous solutions to the Einstein field equation itself, as the method outlined here requires, in general, a metric, at the least, for its employment—one needs already in place what one would be attempting to approximate—carving up spacetime to make it finite, in order to approximate spacetime itself. Indeed, this is why we have not attempted to take into account the inaccuracy of spatiotemporal determinations themselves in formulating the notion of a regime. The issues raised are too hard to be dealt with here. Of course, to a certain extent, one faces the same problem writ small in applying the method to any fields on a relativistic spacetime, as the fields one is attempting to approximate themselves form the flesh and bone of the spacetime's metric structure, in virtue of the Einstein field-equation. The methods work only in so far, then, as one ignores the contribution of the fields one is modeling to this metric structure. Consequently, this method itself has a limited regime of applicability, as it were: it cannot be applied to fields the intensity of which makes untenable the excluding of their contribution to the metric structure. The reasoning behind this conclusion, I believe, points to serious, generally ignored questions about the definability, in general relativity, of observable quantities in regions of intense curvature.

4.5 Physical Theories

[*** IT IS PRECISELY THE POSSESSION OF A REGIME THAT DIFFERENTIATES OTHERWISE FORMALLY IDENTICAL THEORIES—it is the representation of the fact that we distinguish among physical systems that are from a certain formal point of view dynamically identical. Otherwise, one could not tell what sort of physical system one was modeling. IT'S HOW A THEORY BECOMES A *PHYSICAL* THEORY, ABOUT SOME PARTICULAR PART OF THE PHYSICAL WORLD. ***]

From §4.4, we now have (a sketch of) mathematical structures in hand in the terms of which we can model the behavior of physical systems in such a way as to incorporate directly into the model itself the sorts of constraints and conditions a regime may place on the application of a particular theory to that sort of modeling, and in the same terms of which we may articulate mathematically exact and rigorous theories. We will attempt, in this section of the paper, to use those structures to construct a single, unified model of the practices and of the subject-matters of the theoretician and the experimentalist. We must deal with several issues before we will be in a position to shoot directly for that goal. In particular, the discussions of sections 4.3 and 4.4 raise a poignant question: what becomes of the initial-value formulation of the partial-differential equations comprised by a theory with a regime? From the mathematical point of view, the partial-differential equations of the theory, in modeling a system in the sense of theoretical physics, are formulated in terms of quantities modeled by exact scalar and tensorial fields on spacetime—cross-sections of an exact linear bundle over the spacetime manifold itself—and not by fields of bounded, connected, compact intervals of $\mathbb{R}$ over the closures of convex, normal, open regions of spacetime, in the terms of which, as I contend, the arguments and results of the experimentalist may be framed.

The penultimate goal of this section is to understand how one can construct a well set initial-value formulation for partial-differential equations over colored, inaccurate fields in such a way as to take account of the demands possession of a regime imposes on the equations comprised by a theory.[75] In order to get there, we will need to clarify what these demands are, and how the structures developed in the previous section encode them. I begin by briefly sketching in §4.5.1 some of the details of the ways that the mathematical structures introduced in §4.4 can be used to construct physical theories whose components directly model the restrictions and conditions its regime, as spelled out in §4.3, may demand of it, as opposed to a theory employing the ordinary, exact structures of mathematical physics, which will necessarily have its regime appended as an entity external to the rest of the theory. Next, in §4.5.2, I study possible ways our inaccurate, colored fields and the exact, spatiotemporal fields customarily employed in physics may relate to each other, with the aim of fixing a canonical, physically significant relation between the two. This will put us in a position to consider in §4.5.3, in a purely formal way, how to define an initial-value formulation for partial-differential equations on colored, inaccurate fields, and what it may mean for

75. Integral equations are beyond the scope of this paper.

one to be well set, using as our guides the analogous notions in the theory of exact partial-differential equations. We will also consider the possible relations, from a formal perspective, between these two sorts of equation and their respective sorts of solution. In §4.5.4, in the light of these discussions, we will re-work the purely formal notion of a well set initial-value formulation for partial-differential equations over colored, inaccurate fields worked out in §4.5.3 in order to take account of the demands possession of a regime imposes on the equations comprised by a theory.

This analysis will lead us to the ultimate goal of this section, a more precise, partial characterization of a physical theory, given in §4.5.5, which will provide the terms in which we may at last, in §4.5.6, articulate the primary contention of this paper, for which much of the paper up to that point may be considered a constructive proof, that of the consistency of the joint practice and subject-matter of the theoretician and the experimentalist.

4.5.1 Exact Theories with Regimes and Inaccurate, Colored, Kinematically Constrained Theories

By the end of §4.3.3, we had more or less worked out what it meant to constrain by a kinematical regime the applicability of a theory of mathematical physics as ordinarily practised, *viz.*, a theory that represents the quantities it treats using exact fields on spacetime and that comprises only exact partial-differential equations over those fields. There is no clear way, using only that machinery, to represent in a single, unified structure the practice, on the one hand, of the theoretician in abducting exact, rigorous theories from the inaccurately determined data provided by the experimentalist, about which she must judge whether or not they capture and express adequately the essential form of the patterns inherent in the data, and, on the other hand, that of the experimentalist in reckoning from the exact, rigorous theories provided by the theoretician models of actual experiments, about which he must judge whether or not they adequately model his experimental arrangements, and, if so, whether their predictions conform to the inaccurately determined data he gathers from those experiments. The construction of such a representation, however, is the project we have set ourselves in this paper. In order to move towards its accomplishment, we will now explicate how the mathematical machinery developed in §4.4 provides a framework in the terms of which we can frame a unified, consistent representation of the practice of both.

Recall from definition 4.3.3.1 that a kinematical regime imposes the following kinds of conditions on the admissible exact values of the quantities modeled by a theory. First, there are algebraic and differential conditions that the values of the fields representing not only the quantities modeled by the theory but also those representing the relevant environmental quantities must satisfy. At a minimum these conditions include uniform upper and lower bounds on the values these fields can take, as well as on the values of their partial-derivatives up to some fixed order. For tensorial quantities, these bounds are imposed on the kinematical norms we constructed for the fields, and for their derivations, in §4.4.5. The regime dictates that the spatiotemporal region in which the measurements or observations take place conforms to the metrical conditions it imposes. It demands that the preparation and the measurement of these quantities proceed by way of one of a set of fixed interactions mediated by one of a set of experimental techniques, all of a type appropriate for the given scheme of conditions. Finally, it requires a family of algorithms for computing the intervals of inaccuracy of the determined values of the quantities based on the actual conditions of the experiment (the values of the environment's quantities, the exact spatiotemporal character of the laboratory, *etc.*).

To make all this a little more concrete, consider the following, highly schematic description of the way an experimentalist might go about observing the dynamical evolution of a known type of system in order to compare the results of the observation with the predictions of a known physical theory that models that sort of system. He begins by sketching a crude schema of the type of experiment he wants to perform, containing just enough information for him to prepare a suitable laboratory and experimental arrangement of appropriate apparatuses in the laboratory for its performance; [76] he then prepares within this arrangement a token of the type of system he will observe, adjusting its initial state as nicely as the experiment requires and as available techniques and equipment and his knowledge allow, rendering it amenable to the observations he will make of it, in the context of the arrangement; he formulates a model of the arrangement of his proposed experiment, including the system, in the terms of the theory, representing all the germane quantities of the given, actual system using the set of exact fields on spacetime, perhaps some scalar, perhaps some tensorial, that the theory uses to represent those quantities; he then attempts to determine the values of quantities those fields model by employment of one of a variety of techniques and experimental arrangements suitable for the purpose, taking into account as best he could in this determination the family of ranges of possible inaccuracy in the determination of those values of that field as accomplished by the chosen method; after settling in some way or other on the set of exact magnitudes of the fields to be used in the model, he observes the dynamical evolution of the system, measuring the values of the various quantities along the way using some acceptable set of experimental methods, once again trying to take account of the possible inaccuracies while determining these magnitudes; at the same time, he constructs and solves, in the context of the model he has formulated of the arrangement, the initial-value formulation of the partial-differential equations comprised by the theory, using the initially determined values of the quantities as initial data; next, he compares, on the one hand, the final values of the quantities as determined by observation (account having been taken in this determination, as always, of the possible inaccuracies in measurement) with, on the other, the values predicted by the theory in the form of the solution to the equations' initial-value formulation; finally, he decides whether the inevitable deviance of the observed from the predicted values falls within the acceptable margin of error for such observations and calculations, as determined by some method appropriate for the task, in whatever way that propriety may be gauged; if they do, he can conclude that the outcome of the experiment accorded with the predictions of the theory; if they do not, then he must attempt to determine whether this discrepancy amounts to a contravention of the theory requiring its modification in some way or other, or whether the discrepancy can be accounted for by an inadequacy in the performance or in the modeling of the experiment, or in the calculation of the inaccuracies or in that of the acceptable deviances of predicted from observed values that could be rectified by a repeat performance.

To make these ideas vivid, imagine that the experimentalist has a set of thirteen canonical books containing all (and only) the information he needs to plan, model, perform and analyze experiments for all types of physical systems treated by a particular theory with its regime. [77] The volumes are as follows.

[*** A better scheme? remove all purely non-syntactic elements from the first list; segregate all semantical elements into the 13th volume. Also, represent the volumes in symbols, as otherwise it's too difficult to keep in mind? Or, *à la* Peirce, leave it in words? ***]

1. The first book contains: an enumeration of physical quantities, both kinematic and dynamic; and an enumeration of types of physical system, the space of states of each of which can be parametrized by the (values of the) dynamic physical quantities in the first enumeration.
2. The second volume contains: an enumeration of a set of physical quantities, both kinematic and dynamic; and an enumeration of types of environment, each of which bears all the quantities in the first enumeration.
3. The third contains an enumeration of ordered pairs consisting of a type of system listed in

76. Because this is only a schematic description, one should not take literally the seemingly temporal verbs, conjunctions and adverbial phrases it employs, such as 'begin', 'then', 'at the same time', and so on; if one likes, they coordinate only the logical relations of the different bits of the description.

77. The descriptions of these volumes is, of necessity, schematic and sketchy in the extreme.

the first volume and a type of environment listed in the second volume; this enumeration of ordered pairs, moreover, is such that each type of system in the first volume appears in at least one of the ordered pairs and each type of environment listed in the second volume appears as well in at least one of these pairs (*i.e.*, projection on each component of the ordered pairs is surjective).

4. The fourth contains a Σ-bounded, colored, normed, inaccurately linear tensor-algebra.
5. The fifth contains an enumeration of the same cardinality as the enumeration of quantities in the first volume, each enumerand of which is a tensorial subspace of the algebra in the fourth volume; each tensor-space, moreover, is of an indexical structure appropriate for the representation of the physical quantity at the same ordinal position in the enumeration of the first volume as this tensor-space occupies in the present enumeration.
6. The sixth contains an enumeration of the same cardinality as the enumeration of quantities in the second volume, each enumerand of which is a tensorial subspace of the algebra in the fourth volume; each tensor-space, moreover, is of an indexical structure appropriate for the representation of the physical quantity at the same ordinal position in the enumeration of the second volume as this tensor-space occupies in the present enumeration.
7. The seventh volume contains: an enumeration of physical quantities (partially) characterizing subsets of a spacetime; an enumeration of the same cardinality as this first enumeration, each enumerand of which is a tensorial subspace of the algebra in the fourth volume; each tensor-space, moreover, is of an indexical structure appropriate for the representation of the physical quantity at the same ordinal position in the first enumeration of this volume as this tensor-space occupies in the present enumeration.
8. The eighth contains an enumeration of ordered triplets, each consisting of a constant $k > 0$ (not necessarily the same for the first component of every enumerand), a k-bounded, tensorial compounding family defined over the elements of the algebra in the fourth volume, and an ordered pair from the enumeration in the third volume.
9. The ninth contains an enumeration of ordered pairs, the first component of which is a formal set of differential and algebraic conditions formulated in terms of the tensor-product of all the tensor-spaces listed in the fifth and sixth volumes (formal in the sense that we do not specify a particular k-bounded, tensorial compounding family for the algebraic and differential operations on the elements of the algebra), and the second component of which is an ordered pair from the enumeration in the third volume.
10. The tenth contains an enumeration of ordered pairs, the first component of which is a formal set of algebraic conditions formulated in terms of the tensor-product of all the tensor-spaces listed in the fifth, sixth, and seventh volumes (formal in the sense that we do not specify a particular k-bounded, tensorial compounding family for the algebraic and differential operations on the elements of the algebra), and the second component of which is an ordered pair from the enumeration in the third volume.
11. The eleventh contains an enumeration of ordered triplets, the first component of which is a type of experimental apparatus, the second a schematic technique for using that type of apparatus in experiments, and the third component of which is an ordered pair from the enumeration in the third volume.

The twelfth volume is more complex. It contains an enumeration of ordered octuplets, the components of which (in order) are

1. an ordered pair from the third volume
2. an ordered triplet from the listing in the eleventh volume, having as its third component the ordered pair in the first component of this octuplet

3. an ordered triplet from the listing in the eighth volume, having as its third component the ordered pair in the first component of this octuplet
4. the tensor-product of the k-bounded subspaces of all the tensor-spaces enumerated in the fifth volume, where k is given by the first component of the ordered triplet in the third component of this octuplet
5. the tensor-product of the k-bounded subspaces of all the tensor-spaces enumerated in the sixth volume, where k is given by the first component of the ordered triplet in the third component of this octuplet
6. the tensor-product of the k-bounded subspaces of all the tensor-spaces enumerated in the seventh volume, where k is given by the first component of the ordered triplet in the fourth component of this octuplet
7. an ordered pair from the ninth volume, having as its third component the ordered pair in the first component of this octuplet
8. an ordered pair, the first component of which is an ordered pair from the tenth volume, having as its third component the ordered pair in the first component of this octuplet; the second component of the ordered pair consists of a particular relativistic spacetime, $(\mathcal{M}, g_{ab})$.

The contents of these first twelve books provide a complete, (mostly) syntactic articulation of the regime of the theory the experimentalist employs in modeling those sorts of experiments.

The thirteenth book provides a semantical model of the syntactics of the first twelve volumes, as follows. [78]

1. The first enumeration of the first volume lists the types of physical system modeled by the theory, and the second enumeration lists the types of physical quantities shared by those systems that are treated by the theory, in virtue of which the same theory can model all these system, though they (the physical systems) differ in ways significant enough for us to declare them of different types. [79]
2. The first enumeration in the second volume lists the types of environment in which the systems enumerated in the first volume may manifest themselves in a form amenable to treatment by the theory, and the second enumeration lists those quantities borne in common by those environments the values of which are relevant in the determination of the applicability of the theory in modeling a system existing in one of these types of environments.
3. The ordered pairs of the third volume represent those combinations of particular types of systems and environments that do manifest themselves together and are in fact amenable to modeling by the theory.
4. The tensor-algebra of the fourth volume is the one used for all mathematical modeling of these types of systems and environments, and the relevant properties of regions of spacetime in which combinations of the two may manifest themselves.
5. The tensor-spaces of the fifth volume are the ones whose tensor-product is used to represent states of the types of physical systems listed in the first volume.

78. I suppose that we could have completely segregated the syntactic from the semantic elements in this mythology, by making no mention until the thirteenth volume of physical systems, physical quantities, environments, measuring apparatuses, experimental techniques, and so on, leaving only the formal, mathematical elements in the first twelve volumes. I don't see what would be gained, either with regard to physical comprehension or philosophical understanding, by such a maneuver, which strikes me as artifical and contrived.

79. I would love to know whether there is any basis over and above the preferences our psychological constitution more or less enforces on us for nominating physical systems to be of different "types" when, with respect to theory, they share the same dynamical structure. Think of the Newtonian equation representing the dynamical evolution of a simple harmonic oscillator, and then of how many systems of seemingly different "types" find their appropriate model in that equation. Think as well of the fact that, at the level of quantum field theory, the idea of the simple harmonic oscillator plays a fundamental role in several ways, if not most, of articulating the structure of the theory.

6. The tensor-spaces of the sixth volume are the ones whose tensor-product is used to represent states of the types of environments listed in the second volume.

7. The first enumeration in the seventh volume lists those properties of regions of spacetime the values of which are relevant in the determination of the applicability of the theory in modeling a system existing in one of these types of environments in a spacetime region with given values for the listed properties; the second lists the tensor-spaces whose tensor-product is used to represent the properties of spacetime regions listed in the first enumeration of the volume.

8. The first component of a triplet in the eighth volume, k, represents the minimal bound on values of the scalar quantities (and on the kinematical norms defined on the tensorial quantities) demanded by the regime of a theory; the second component is a k-bounded, tensorial compounding families used to impose convex, k-bounded, inaccurate, tensorial, algebraic and differential structures on the germane subspaces of the algebra in the fourth volume, so as to represent the way that the inaccuracies of all the relevant quantities combine when experiments performed on the system *cum* environment indicated by the third component of the triplet are modeled using the theory.

9. A set of algebraic and differential formulæ constituting the first component of an ordered pair listed in the ninth volume defines the conditions that the values of the quantities of tokens of the type of physical system and tokens of the type of environment indicated in the ordered pair constituting the second component must jointly satisfy, when one of the systems manifests itself in one of the environments, in order to be amenable to modeling by the theory.

10. The algebraic formulæ constituting the first component of an ordered pair listed in the tenth volume defines the conditions that the values of the quantities of a region of spacetime, in conjunction with those of tokens of the type of physical system and tokens of the type of environment indicated in the ordered pair constituting the second component, must jointly satisfy, in order for those systems to be able to manifest themselves in those environments in such a region of spacetime, so as to be amenable to modeling by the theory.

11. The first two components of the ordered triplets in the eleventh volume consist of experimental apparatuses and methods for employing them that experimentalists can use to perform experiments on the given type of physical systems *cum* environment indicated in the ordered pair constituting the third component of the triplet, in a way amenable to modeling by the theory.

The twelfth volume, not surprisingly, requires a more involved semantical accounting. Each octuplet represents, broadly speaking, a family of systems amenable to modeling by the theory, in accordance with its regime, along with a schematic representation of the way the regime informs and constrains this modeling.

1. The first component of the ordered pair in the first component of an octuplet indicates the type of system at issue, and the second a particular type of environment in which a token of that type of system may appear in such a way as to be amenable in principle to modeling by the theory in accordance with its regime.

2. The second component indicates a type of instrumentation the experimentalist may use, according to the associated techniques, to probe a token of that type of system in a token of that type of environment so that the entire experiment is amenable in principle to modeling by the theory in accordance with its regime.

3. Given an experiment of this type that an experimentalist proposes to perform, the third component of the octuplet fixes the mathematical structure—a convex, k-bounded, inaccurately linear, tensorial algebra—she will use to model the experiment, where k is chosen so as to

enforce the minimal kinematical conditions on the values of the quantities for states of the combined system and environment—that they (and some subset of their partial-derivatives) be uniformly bounded from above and from below. This algebra enforces these minimal conditions by not allowing the representation of any state of the system *cum* environment that does not satisfy them; moreover, the k must be such that all states represented by elements of the algebra have values for their quantities well defined with respect to experimental probing by the chosen instrumentation, as applied using the given technique.

4. The fourth, fifth, and sixth components represent, respectively, the spaces of states of the system, the environment, and regions of spacetime in which the experiment may be performed. [80]

5. The seventh component of the octuplet picks out a subspace of the combined space of states of the system and the environment, those the values of whose quantities jointly satisfy all the given algebraic and differential conditions; states represented by elements of this space satisfy the remainder of the conditions (besides the minimal one captured by the imposition of the k-bound) that the values of the system and the environment must jointly satisfy in order to conform to the regime of the theory, in so far as the system is probed using the experimental apparatus applied using the associated techniques given in the second component of the octuplet.

6. The spacetime indicated in the eighth component of the octuplet represents the world in which the envisioned type of experiment will occur. The algebraic conditions are such that only a canvas can satisfy them; they serve, moreover, to pick out a subspace of the space of canvases on the given spacetime, those in which it is both the case that a token of the type of system *cum* environment whose state can be modeled by an element of the subspace determined by the elements of the seventh component of this octuplet manifests itself, and the case that the values of the relevant spatiotemporal properties of this canvas, along with the values of the quantities of the system *cum* environment manifested in the canvas, all jointly satisfy the algebraic conditions. These are the kinematically appropriate laboratories for these types of experiments according to the theory's regime, as spelled out by these thirteen volumes. This family of laboratories, finally, serves to refine further the subspace determined by the elements of the seventh component of this octuplet, to those systems *cum* environments appearing in one of these laboratories. The systems represented by elements of this subspace are precisely those, in the given spacetime, that do in fact conform to the regime of the given theory, and are such that experiments performed on them using the given experimental apparatus applied using the associated techniques can be consistently modeled by the theory in accordance with its regime.

We thus have two senses in which a physical theory may possess, and conform to, a regime. The theory and its regime may, on the one hand, be formulated in the terms of the ordinary, exact structures of mathematical physics and so have its regime appended as a separate entity that all the relevant, exact quantities must be made to conform to, as a separate mathematical condition on them. This is the mode of representation used in §4.3 when we first characterized a regime. We will call such a representation of a theory an *exact theory with regime*. On the other hand, the theory may be formulated from the start in terms of inaccurate, colored structures that incorporate directly the strictures of its regime, in the way sketched just above. We will call such a theory an

80. There is an awkwardness here in the formal presentation of these structures. We are demanding that the space representing the properties of regions of spacetime relevant to determining whether or not a given experiment conforms to the regime of a theory have the same k-bound as the spaces representing the states of, respectively, the environment and the system. It would be preferable to have the space representing these properties be only Σ-bounded, but to characterize the tensor-product of a Σ-bounded space and a k-bounded space would take us too far afield. The same, indeed, goes for the space of states of the environment, which one may expect to have a k-bound different from that of the space of states of the system. See Curiel (2011) for complete details.

inaccurate, colored, kinematically restricted theory, or just an *inaccurate theory*, for short. We will say that an inaccurate theory is *kinematically equivalent* to an exact theory with regime if all the kinematic constraints encoded in its inaccurate, colored structures embody the regime of the exact theory. We will try to make these ideas precise in the next several sections.

Before moving on, I must record a qualm about this discussion—I hesitate to include the first two volumes (and so several of the others), those enumerating the physical systems and environments, in my proposed set of mythological volumes. They raise perhaps a larger number of difficult issues and questions they do not address than that of simple or difficult ones they do. What does a "type" of physical system come to? How are such types differentiated, if not by the very theories that successfully model them? Ought a type be a catalogue of all existing physical systems or only of possible ones, or only a listing of properties a system should have, necessary or sufficient, to constitute a token of the type? In any event, how ought one distinguish individual physical systems, one from another, for surely one can consider in a natural way the same sum total of fields of quantities in the same patch of spacetime as part of more than one "system", depending on the joints one chooses to carve at? And so on. In the event, it is in large part on account of these very questions and others like them that I decided to include the two volumes. I see no way of satisfactorily answering any of them, and all others like them, once and for all. Indeed, I would argue that, viewed *sub specie æternitatis*, these questions have no sense. The only hope I see for their satisfactory address lies in formulating them with regard to a particular, larger set of issues one is trying to resolve, with well defined goals and clearly delineated methods acceptable for the use in achieving them. I thus include them, in part, to underscore again the thoroughly pragmatic character of all attempts to understand and to employ scientific theories. Someone of a more pessimistic stripe could say with some justice that its inclusion underscores how little one would have accomplished in comprehending *physical* theory were one to have accomplished even to a high degree of success the project I have set myself in this paper.

4.5.2 Idealization and Approximation

In order to make contact with theoretical physics, we need a method for associating an exact scalar field with a coloring by a relation substantive enough to use in comparing models of physical systems as represented by our constructions with those of physical systems as represented by the types of theories employed in physics as ordinarily practiced, along with the strictures of an externally imposed regime. We will associate the two by means of a construction—an algorithm that, given a coloring as input, yields an exact scalar field that may be considered an idealized, exact model of the same quantity attached to the same physical system the coloring inaccurately models, satisfying the kinematical constraints encoded in the coloring. This algorithm will of necessity have a bipartite character, for it will not only transform a field on an infimal decoupage of a canvas to a field on the spacetime points composing the canvas itself, it will also transform the field from one valued in $\Re$ to one valued in $\mathbb{R}$—it will need to transform both the domain and the range of the field. There are several ways one may envisage implementing such a procedure. Purely for the sake of simplicity, we will construct a two-stage algorithm, transforming the domain and the range each in its own operation.

Before sketching the characterization of the operators we will use, it will be instructive to consider a possible method of constructing one for the domain, mapping $\Theta_{\Re}$ onto $\Sigma_{\Re}$, that will in the end not serve the purpose, though it will, in the event, shed useful light on the physical content of the machinery developed so far. Fix a canvas C along with its decoupage $\mathfrak{C}$. Let us say that the inaccurate scalar field ζ is *in harmony with* (or *harmonious with*) the coloring θ if, for all $S \in \mathfrak{C}$,

$$\theta(S) = \int_{[S,\Re]} \zeta \, \mathrm{d}\hat{\mu}$$

One might hope to associate a harmonious field with a given coloring, if that coloring satisfies certain conditions. In general, however, this will not be possible in a physically viable way. To see this, recall that a measurable, exact scalar field f is said to be *absolutely continuous* in a region C if, given any $\epsilon > 0$, there exists a $\delta > 0$ such that, given any finite set $\{S_i\}_{i=1,\ldots,n}$ of mutually disjoint subsets of C for which

$$\sum_{i=1}^{n} \int_{S_i} \epsilon_{abcd} < \delta$$

then it is the case that

$$\sum_{i=1}^{n} \int_{S_i} f\,\epsilon_{abcd} < \epsilon$$

Analogously, an inaccurate set function θ, in this case a coloring, is said to be *absolutely continuous* with respect to a σ-finite measure μ on the σ-ring $\mathfrak{M}$ if, given any $\epsilon_1, \epsilon_2 > 0$, there exists $\delta > 0$ such that, if $\{S \in \mathfrak{M} : \mu(S) < \delta\}$, then $\pi_1 \circ \theta(S) < \epsilon_1$ and $\pi_2 \circ \theta(S) < \epsilon_2$. The Radon-Nikodým Theorem[81] has as an immediate consequence the fact that a coloring possesses an inaccurate scalar field in harmony with it if and only if that coloring is absolutely continuous. Not every coloring, however, not even every continuous or even smooth coloring, need be absolutely continuous. Being absolutely continuous implies that the coloring must, in general, take on smaller values for scraps of smaller volume, but nothing requires this of a coloring.

The physical content of being absolutely continuous is easily illustrated. Let ζ be harmonious with the coloring θ. Then for $S \in \mathfrak{C}$, because ζ is continuous (since absolute continuity implies continuity), there is, by the mean-value theorem, a $q \in S$ such that

$$v[S]\pi_1(\zeta(q)) = \pi_1 \circ \int_{[S,\mathfrak{R}]} \zeta \mathrm{d}\hat{\mu}$$

where we use the pragmatic, multiplicative group-operation, and so

$$\pi_1(\zeta(q)) = \frac{\pi_1(\theta(S))}{v[S]}$$

Let $T \in \mathfrak{C}$ be another set such that, for the same q,

$$v[T]\pi_1(\zeta(q)) = \pi_1 \circ \int_T \zeta \epsilon_{abcd}$$

(such a T can always be found), and so

$$\pi_1(\zeta(q)) = \frac{\pi_1(\theta(T))}{v[T]}$$

Thus, a constraint on the definition of θ is that, for all $S, T \in \mathfrak{C}$ that share a mean-value point,

$$\frac{\pi_1(\theta(S))}{v[S]} = \frac{\pi_1 \circ (\theta(T))}{v[T]}$$

or, equivalently,

$$\frac{\pi_1(\theta(S))}{\pi_1(\theta(T))} = \frac{v[S]}{v[T]}$$

In effect, this says that absolutely continuous inaccurate scalar fields are inaccurate scalar densities: θ will scale in a way that approximates to being in inaccurately linear proportion to the volumes

81. See, *e.g.*, Halmos (1950, ch. VI, §31).

of the scraps it takes values in. It is difficult to see why such a severe constraint should be placed on a θ that is to represent possible initial data for a physical theory. Temperature, for instance, is a true scalar and does not satisfy this condition. All the same, it is useful to know that colorings representing scalar densities ought in fact to be absolutely continuous, at a minimum.

We turn now to characterize operators that will serve the purpose. Consider the physical content of transforming a coloring into an exact scalar field. A single measurement made under actual, laboratorial conditions consists of necessity of a sort of smeared-out average of a multitude of more miniscule and varied interactions. A thermometer, for example, does not stay stationary with respect to the caloric mass being probed, certainly not, in any event, as measured in units of the order of magnitude of the mean free-path of the basic constituents of the caloric mass (that is to say, basic with respect to the theory at issue), which, by definition, will be of the same order of magnitude as the size of the infimal scraps determined by the kinematic regime of the theory at issue. Rather, it bobs and jiggles constantly, sampling, as it were, the temperature over most if not all the region immediately proximate to the point of interest during the time it takes for the system to equilibrate. On this picture, the value of the coloring on "almost every" infimal scrap containing a given point $q \in C$ ought to contribute to the value of the constructed scalar field at q. Still, even though the thermometer in its jiggling samples "almost every" scrap containing q, one may expect that the largest contributions to the final reading will come from those scraps in which q lies closest to the spatiotemporal center, or at least furthest from the boundary, in some sense or other. Note that exactly these sorts of considerations as well illuminate some of the sources of the inevitable inaccuracy in measurements and observation in physics.

To render these considerations precisely, we begin by defining operators that perform the yeoman's work, acting only on the domains and ranges of the fields at issue, which we will then employ to define the required mappings among the fields themselves. We require two kinds, one that maps scraps onto spacetime points, and so maps decoupages onto canvases, and another that maps $\Re$ onto $\mathbb{R}$. With these operators in hand, we will have two obvious methods of constructing operators that map inaccurate colorings on a decoupage $\mathcal{C}$ to exact scalar fields on its canvas C: we may define a pair of operators mapping, respectively, $\Theta_{\Re} \to \Sigma_{\Re}$ and $\Sigma_{\Re} \to \Sigma$, and then chain the two together; or we may define a pair mapping, respectively, $\Theta_{\Re} \to \Theta$ and $\Theta \to \Sigma$, and then chain the two together. The fact that two methods offer themselves will provide a clue as to natural conditions to impose on the operator we ultimately define to map $\Theta_{\Re} \to \Sigma$.

Definition 4.5.2.1 *A* shrinker *is a surjective, continuous, open and closed mapping* $\mathbf{S}$ *of a decoupage* $\mathcal{C}$ *onto its associated canvas* C*, such that, for every* $S \in \mathcal{C}$*,* $\mathbf{S}(S) \in \breve{S}$*.*

For any scrap $S \in \mathcal{C}$, we will write 'p_s' for the value of $\mathbf{S}$ at the scrap S, and we will refer to it as the scrap's *center*.

Definition 4.5.2.2 *An* exactor *is a surjective, continuous, open and closed mapping* $\mathbf{E}$ *of* $\Re$ *onto* $\mathbb{R}$*, such that, for every* $(a,\, \chi) \in \Re$*,*

$$a - \frac{1}{2}\chi \;<\; \mathbf{E}(a,\, \chi) \;<\; a + \frac{1}{2}\chi$$

We will refer to the image of an inaccurate scalar $(a,\, \chi)$ under an exactor, written 'a_χ', as its *exactitude*. It immediately follows that smooth shrinkers and exactors are submersions.

To show that shrinkers exist, I explicitly constructed a useful one in Curiel (2011). Though the fine details are too involved to go into in any depth here, the construction, roughly speaking, involves a method for explicitly shrinking the scraps of an infimal decoupage in a smoothly parametrized way to points of spacetime. In order to reflect the considerations drawn out in our discussion just above of thermometry, it addressed one desideratum in particular: settling on a spacetime point to serve as p_s, its center as we have termed it, contained in a given scrap that lies, in some technical sense, furthest from its boundary and closest to its "true spatiotemporal center". It relies on the approximation, mentioned above in §4.3.2, that we used in making the space of decoupages finite-dimensional from the full, infinite-dimensional decoupage. Using the parameters employed in that approximation, I defined a kind of "center of mass" of a scrap as that point that satisfied a set of conditions on the maximum and minimum values of those parameters. This mapping, moreover, smoothly varies in the spacetime points as one moves smoothly around the decoupage. I emphasize the point that, so far as I can see, there is no canonical or preferred way of constructing a shrinker. Indeed, the definition itself could be altered in any of a number of ways while remaining true to the spirit behind it.

Displaying examples of exactors is far easier. On the face of it, it would appear to be trivial, consisting of nothing more than the shrinking of the second component of an inaccurate scalar, *i.e.*, the extent of the inaccuracy itself, to zero, and, indeed, π_1 satisfies all the conditions and so is an exactor. Selecting the "proper" exactor, however, what ever criteria we plump for in coming to the judgment, is not so straightforward. Nothing requires that we hold the first component, the magnitude, fixed as we shrink the inaccuracy. We chose the magnitude to be the mid-point of the range of inaccuracy for the sake of convenience and simplicity, and because it satisfied a few physical, heuristic (*i.e.*, hand-waving) arguments. The idealized value of the quantity as represented by the exact scalar can in fact be any point in the entire interval of inaccuracy. One thus has a continuum of exactors, shrinking the same inaccurate scalar down to any of a continuously varying family of exact scalars.

We will use shrinkers and exactors to define the operators that map fields to fields. Because shrinkers are not injective and operate on the domains of the fields, a little footwork remains before we can define operators based on them, which we do first. The operators based on exactors, which we treat after the ones based on shrinkers, are easier to deal with; it will not matter that they are not injective.

To define a mapping from a space of fields on a decoupage to one on a canvas using a shrinker, one cannot simply declare that the value of the field on the canvas at a given point be the value of the scrap that gets mapped to that point, as there will be, in general, many scraps that get mapped to the same point. The discussion of the thermometer suggests that, given a shrinker, we require a way of distilling a value for a field at a spacetime point q in a canvas from the values of the coloring at all the scraps in the pre-image of q under the shrinker. We also demand that the relations among the value of the field at q and those at the scraps in its pre-image reflect the fact that the value at q is supposed to be an idealization of some sort of the inaccurate values of the coloring on the scraps in its pre-image. There are several ways one can try to make this idea precise. The one we work with has the virtues of simplicity and manifest physical content. Write the power-set of $\mathcal{C}$ as '$\mathfrak{P}[\mathcal{C}]$'. Given a shrinker $\mathbf{S}$, define $\mathfrak{P}_s[\mathcal{C}] \subset \mathfrak{P}[\mathcal{C}]$ to be the family of all and only those $\mathcal{S} \in \mathfrak{P}[\mathcal{C}]$ for which there is a $q \in C$ such that $\mathcal{S} = \mathbf{S}^{-1}[q]$. Roughly speaking, elements of $\mathcal{S}$ are those maximal families of scraps having a single point as their total intersection.

Definition 4.5.2.3 *Given a shrinker* $\mathbf{S}$*, an* $\mathbf{S}$-distiller *is a mapping* $\mathbf{d}_s : \Theta_{\Re} \times \mathfrak{P}_s[\mathcal{C}] \to \Re$ *such that, for* $\theta \in \Theta_{\Re}$ *and* $q \in C$*,*

$$\inf_{S \in \mathbf{S}^{-1}[q]} \{||\theta(S)||\} \;\leq\; ||\mathbf{d}_s(\theta,\, \mathbf{S}^{-1}[q])|| \;\leq\; \sup_{S \in \mathbf{S}^{-1}[q]} \{||\theta(S)||\}$$

Similarly,

Definition 4.5.2.4 *Given a shrinker* $\mathbf{S}$*, an* exact $\mathbf{S}$-distiller *is a mapping* $\hat{\mathbf{d}}_s : \Theta \times \mathfrak{P}_s[\mathcal{C}] \to \mathbb{R}$ *such that, for* $\theta \in \Theta$ *and* $q \in C$*,*

$$\inf_{S \in \mathbf{S}^{-1}[q]} \{|\theta(S)|\} \;\leq\; |\hat{\mathbf{d}}_s(\theta,\, \mathbf{S}^{-1}[q])| \;\leq\; \sup_{S \in \mathbf{S}^{-1}[q]} \{|\theta(S)|\}$$

Again, in order to show that distillers exist, I explicitly constructed one in Curiel (2011), based on the shrinker I constructed (mentioned above), albeit the constructed distiller was defined only on $\mathcal{L}_1[\mathfrak{C}, \Re]$ (more precisely, on the inaccurate Sobolev spaces based on it), not on all of $\Theta_\Re$, which, however, is all we require. The same construction serves to show that exact distillers exist as well. In brief, given an $\mathcal{L}_1$-coloring θ over a canvas, and a point q in that canvas, the value of the constructed distiller as applied to $(\theta, \mathbf{S}^{-1}[q])$ consists of the limit of the value of a kind of weighted average of the values of the coloring over the scraps in $\mathbf{S}^{-1}[q]$ as the scraps shrink to q.[82]

We have enough under our belt now to characterize the operators this section has worked towards. We first treat those that map the space of colorings into that of inaccurate scalar fields; we can then immediately extend the definition to operators mapping the space of exact colorings into that of exact scalar fields. The basis for these operators will be shrinkers and distillers. Given a shrinker $\mathbf{S}$ and an $\mathbf{S}$-distiller, we will say that the field $\zeta \in \Sigma_\Re$ is the $\mathbf{d}_s$-*distillate* of $\theta \in \Theta_\Re$, if it is such that, for all $q \in C$, $\zeta(q) = \mathbf{d}_s(\theta, \mathbf{S}^{-1}[q])$. The *exact* $\mathbf{d}_s$-*distillate*, a field in Σ derived from a field in Θ, is defined in the obvious way, using definition 4.5.2.4 rather than 4.5.2.3. It would be convenient for these operators to have such nice properties as mapping $\mathcal{L}_1[\mathfrak{C}, \Re]$ to $\mathcal{L}_1[C, \Re]$, $\Theta^b_\Re$ to $\Sigma^b_\Re$, and so on. It is indispensable that they "act as linearly as they can", which in this case means that they ought to be, respectively, inaccurately and exactly linear.

Definition 4.5.2.5 *A* lens *is an inaccurately linear bjiective* $\mathbf{L} : \mathcal{L}_1[\mathfrak{C}, \Re] \to \mathcal{L}_1[C, \Re]$ *such that*

1. $\mathbf{L}$ *is bounded and stable in the operator-norm*
2. *there exists an ordered pair* $(\mathbf{S}, \mathbf{d}_s)$ *consisting of a shrinker* $\mathbf{S}$ *and an* $\mathbf{S}$*-distiller* $\mathbf{d}_s$ *such that, for all* $\theta \in \mathcal{L}_1[\mathfrak{C}, \Re]$, $\mathbf{L}[\theta]$ *is the* $\mathbf{d}_s$*-distillate of* θ
3. *the restriction of the action of* $\mathbf{L}$ *to* $\mathcal{L}_2[\mathfrak{C}, \Re]$, $\Theta^b_\Re$, $\Theta^0_\Re$ *and* $\Theta^\infty_\Re$ *is, respectively, an inaccurately linear bijection into* $\mathcal{L}_2[C, \Re]$, $\Sigma^b_\Re$, $\Sigma^0_\Re$ *and* $\Sigma^\infty_\Re$

Note that, in the last item, since we are dealing with the spaces $\mathcal{L}_2[C, \Re]$, $\Sigma^b_\Re$, *et al.*, defined as restrictions of $\mathcal{L}_2[C, \Re]$, we assume they have the $\mathcal{L}_1$-topology. We will say that $\mathbf{L}$ is *derived from* $\mathbf{d}_s$. Let us call the value of a coloring under a lens its *focus*.

We can, almost without comment, modify this discussion to characterize a mapping from the space of exact colorings into that of exact scalar fields.

Definition 4.5.2.6 *An* exact lens *is a linear bijection* $\widehat{\mathbf{L}}_s : \mathcal{L}_1[\mathfrak{C}] \to \mathcal{L}_1[C]$ *such that*

1. $\widehat{\mathbf{L}}$ *is bounded and stable in the operator-norm*
2. *there exists an ordered pair* $(\mathbf{S}, \mathbf{d}_s)$ *consisting of a shrinker* $\mathbf{S}$ *and an* $\mathbf{S}$*-distiller* $\mathbf{d}_s$ *such that, for all* $\theta \in \mathcal{L}_1[\mathfrak{C}]$, $\widehat{\mathbf{L}}[\theta]$ *is the exact* $\mathbf{d}_s$*-distillate of* θ
3. *the restriction of* $\widehat{\mathbf{L}}$ *to* $\mathcal{L}_2[\mathfrak{C}]$, Θ^b, Θ^0 *and* Θ^∞ *is, respectively, a linear bijection into* $\mathcal{L}_2[C]$, Σ^b, Σ^0 *and* Σ^∞

Let us call the value of an exact coloring under the latter mapping its *exact focus*. One has available for exact lenses all the constructions used to extend lenses to higher-order structures.

I showed in Curiel (2011) that lenses exist by proving that the constructed shrinker and its derived distiller defined one. The case of most importance for us will be that in which the original coloring is defined on the infimal decoupage of a canvas that satisfies the requirements of the regime

82. Strictly speaking, this description is misleading. A coloring is not in general defined for decoupages containing arbitrarily small scraps, and so we cannot compute anything based on the values a coloring assumes on a given scrap that shrinks to a point, once that scrap has become too small with respect to the metrical conditions imposed by the region at issue in its approach to the point. I circumvented this problem in Curiel (2011) by defining a way of computing a number associated with an arbitrarily small scrap based on yet another weighted average of the values of the coloring over all scraps that the coloring assumes values for and that have non-null intersection with the given arbitrarily small scrap.

of a theory. All results proved in Curiel (2011) involving lenses continue to hold when restricted to this case, since the use made of the decoupage in proving them depends only on the compactness of its elements (considered as subsets of spacetime) and on their sizes being (in a certain technical sense) uniformly bounded from below, both of which hold for infimal decoupages. I emphasize the point that, so far as I can see, there is no canonical or preferred way of defining such a mapping. Indeed, given a lens $\mathbf{L}$, one can always define a new one by multiplying $\mathbf{L}$ by a smooth inaccurate scalar field bounded by some $\epsilon > 0$ that depends on the particular properties of $\mathcal{C}$ and $\mathbf{L}$.

The construction of an exact scalar field from an inaccurate scalar field, and, correlatively, an exact coloring from a coloring, is somewhat simpler than the construction of an inaccurate scalar field from a coloring, in that the domain of the field in each case remains the same. We will use the same considerations as guided our definition of a lens to define the process of deriving an exact scalar field from an inaccurate one (and so for that of deriving an exact coloring from a coloring). In this case we demand that the operator yield an exact scalar field respecting the kinematic constraints encoded in the inaccurate field, to wit, that the exact value of the quantity fall within the interval of possible inaccuracy.

Definition 4.5.2.7 *A* polarizer *is a linear bijection* $\mathbf{P} : \mathcal{L}_1[C, \mathfrak{R}] \to \mathcal{L}_1[C]$ *such that*

1. $\mathbf{P}$ *is bounded and stable in the operator-norm*
2. *there exists an exactor* $\mathbf{E}$ *such that, for every* $\zeta \in \mathcal{L}_1[C, \mathfrak{R}]$ *and every* $q \in C$, $\mathbf{P}[\zeta](q) = \mathbf{E}(\zeta(q))$
3. *the action of* $\mathbf{P}$ *restricted to* $\mathcal{L}_2[C, \mathfrak{R}]$, $\Sigma^b_{\mathfrak{R}}$, $\Sigma^0_{\mathfrak{R}}$ *and* $\Sigma^\infty_{\mathfrak{R}}$ *is, respectively, a linear bijection into* $\mathcal{L}_2[C]$, Σ^b, Σ^0 *and* Σ^∞

We will say that $\mathbf{P}$ is *derived from* $\mathbf{E}$. Its being derived from an exactor ensures that the determined value lies within the interval of possible inaccuracy. Given an inaccurate scalar field ζ and a polarizer $\mathbf{P}$, we call the image of ζ under $\mathbf{P}$ its *polarization*. Similarly,

Definition 4.5.2.8 *A* colored polarizer *is a linear bijection* $\widehat{\mathbf{P}} : \mathcal{L}_1[\mathcal{C}, \mathfrak{R}] \to \mathcal{L}_1[\mathcal{C}]$ *such that*

1. $\widehat{\mathbf{P}}$ *is bounded and stable in the operator-norm*
2. *there exists an exactor* $\mathbf{E}$ *such that, for every* $\theta \in \mathcal{L}_1[\mathcal{C}, \mathfrak{R}]$ *and every* $S \in \mathcal{C}$, $\mathbf{P}[\theta](S) = \mathbf{E}(\theta(S))$
3. *the action of* $\mathbf{P}$ *restricted to* $\mathcal{L}_2[\mathcal{C}, \mathfrak{R}]$, $\Theta^b_{\mathfrak{R}}$, $\Theta^0_{\mathfrak{R}}$ *and* $\Theta^\infty_{\mathfrak{R}}$ *is, respectively, a linear bijection into* $\mathcal{L}_2[\mathcal{C}]$, Θ^b, Θ^0 *and* Θ^∞

Given a coloring θ and a colored polarizer $\widehat{\mathbf{P}}$, we call the image of θ under $\widehat{\mathbf{P}}$ its *colored polarization*.

The simplest polarizer, $\mathbf{P}_\pi$, is given by direct application of π_1: $\mathbf{P}_\pi[\theta](q) = \pi_1(\theta(q))$. To construct another, we need only, then, choose some map $\sigma : \Sigma_{\mathfrak{R}} \times C \to \mathbb{R}^+$ such that the mapping $\mathbf{P}_\sigma : \Sigma_{\mathfrak{R}} \to \Sigma$ defined by

$$\mathbf{P}_\sigma[\theta](q) = \sigma(\theta, q)\pi_1(\theta(q)) \tag{4.5.2.1}$$

satisfies definition 4.5.2.7. Indeed, since $\mathcal{L}_1[C, \mathfrak{R}]$ is a Banach space, the implicit function theorem guarantees that, for any polarizer $\mathbf{P}$, there exists a $\sigma : \mathcal{L}_1[C, \mathfrak{R}] \times C \to \mathbb{R}$ such that $\mathbf{P}$ can be represented in the form given by equation (4.5.2.1). In this case, σ is the *implicit polarizer* of $\mathbf{P}$. The mapping from $\mathcal{L}_1[C, \mathfrak{R}]$ to $\mathcal{L}_1[C]$ with perhaps the clearest physical content is that defined by our norm on $\mathfrak{R}$, under which, for $\zeta \in \mathcal{L}_1[C, \mathfrak{R}]$, $\zeta(q) \mapsto ||\zeta(q)||$. This, however, strictly speaking, is not a polarizer, since some of the values of a $\mathcal{L}_1[C, \mathfrak{R}]$ field may not have well defined norms; moreover, there is no guarantee that, when the norm is defined, its value lies within the interval of possible inaccuracy. If we want to avail ourselves of this mapping, we cannot address the issue by simply restricting our definition of a polarizer to $\Sigma^b_{\mathfrak{R}}$, for then we could not extend our constructions to the relevant, inaccurate Sobolev spaces. We therefore will employ a small cheat when we use this polarizer: we will "smooth out" each $\zeta \in \mathcal{L}_1[C, \mathfrak{R}]$ at just the points at which its values have no defined norm, by convolving it with a small, smooth normalizing factor in neighborhoods of the singular points smaller (in a technical sense) than the infimal scraps of the decoupage. Since the set of such singular points will always be of measure zero, this affects nothing of substance. We will, furthermore, apply a smooth normalizing factor (perhaps depending on the value it is applied to) to the mapping, to ensure that the resulting value always lies within the interval of possible inexactness.

Chaining a lens and polarizer together yields a mapping from $\mathcal{L}_1$-colorings to exact $\mathcal{L}_1$-scalar fields, $\mathbf{P} \circ \mathbf{L} : \mathcal{L}_1[\mathcal{C}, \mathfrak{R}] \to \mathcal{L}_1[C]$. Likewise, chaining a colored polarizer and an exact lens together, $\widehat{\mathbf{L}} \circ \widehat{\mathbf{P}}$, yields a mapping from $\mathcal{L}_1$-colorings to exact $\mathcal{L}_1$-scalar fields. We now want criteria to impose on such possible mappings to pick out the physically relevant ones. The fact that I see no way of imposing a preference for one of these ways of chaining over the other suggests one such natural condition, I think the one of importance at a brute level, as it were; as we take account of the demands a regime places on these structures, more constraints will naturally suggest themselves. Let us say that $\mathbf{L}$ and $\widehat{\mathbf{L}}$ are *co-focused* if they are derived from the same $\mathbf{S}$-distiller, and that $\mathbf{P}$ and $\widehat{\mathbf{P}}$ are *co-polarized* if they are derived from the same exactor.

Definition 4.5.2.9 *An* idealizer *is a linear, bijective mapping* $\mathbf{I} : \mathcal{L}_1[\mathcal{C}, \mathfrak{R}] \to \mathcal{L}_1[C]$ *such that there exists an ordered quadruplet* ($\mathbf{L}$, $\mathbf{P}$, $\widehat{\mathbf{L}}$, $\widehat{\mathbf{P}}$) *consisting of a lens* $\mathbf{L}$*, a polarizer* $\mathbf{P}$*, an exact lens* $\widehat{\mathbf{L}}$*, and a colored polarizer* $\widehat{\mathbf{P}}$*, for which*

1. $\mathbf{L}$ *and* $\widehat{\mathbf{L}}$ *are co-focused*
2. $\mathbf{P}$ *and* $\widehat{\mathbf{P}}$ *are co-polarized*
3. *for* $\theta \in \mathcal{L}_1[\mathcal{C}, \mathfrak{R}]$

$$\mathbf{I}[\theta] = \mathbf{P} \circ \mathbf{L}[\theta] = \widehat{\mathbf{L}} \circ \widehat{\mathbf{P}}[\theta]$$

There follows from definitions 4.5.2.5–4.5.2.9

Theorem 4.5.2.10 *For an idealization* $\mathbf{I}$*,*

1. $\mathbf{I}$ *is bounded and stable in the operator norm*
2. *the action of* $\mathbf{I}$ *restricted to* $\mathcal{L}_2[\mathcal{C}, \mathfrak{R}]$*,* $\Theta^b_{\mathfrak{R}}$*,* $\Theta^0_{\mathfrak{R}}$ *and* $\Theta^\infty_{\mathfrak{R}}$ *is, respectively, a linear bijective mapping into* $\mathcal{L}_2[C]$*,* Θ^b*,* Θ^0 *and* Θ^∞

By dint of the fact that the lenses and polarizers I constructed in Curiel (2011) (sketched above), when composed, form an idealizer, there follows

Theorem 4.5.2.11 *Idealizers exist.*

We will call the inverse of an idealizer an *approximator*. As it is easy to see, all the same properties hold for approximators as do for idealizers.

It is not difficult to extend these operators to ones on inaccurate tensorial and affine spaces. We have already done most of the heavy lifting, in §4.4.5. I will restrict myself to stating the definitions and perhaps making a remark or two about the details. I will do so only for lenses and polarizers; the treatment of exact lenses and colored polarizers should then be clear. We will also state definitions only for the Σ-bounded, inaccurate, colored tensorial algebra $\widehat{\mathcal{T}}_{\mathfrak{R}}$ rather than for a convex, k-bounded one $\widehat{\mathcal{T}}_{\mathfrak{R},k}$. The extension of the definitions to the convex, k-bounded case is largely straightforward; the only delicacy lies in keeping track of which combinations of operators and entities are and are not permitted, which must be done manually, as it were, since such spaces are not algebraically closed. Let us write '$\mathcal{T}^{m,n}_{\mathfrak{R}}|_{c}$' for the restriction to the canvas $C \in \mathcal{M}$ of the (non-colored) inaccurate space of (m, n)-tensorial fields on $\mathcal{M}$, and '$\mathcal{T}^{m,n}_{\mathfrak{R}}|_{q}$' for the fiber of that space over $q \in C$, *etc.*

In so far as shrinkers do not care about the character of any space one may define over the scraps and points it deals with, we can use definition 4.5.2.2 in this case without alteration. Everything else, though, requires a bit of reworking.

Definition 4.5.2.12 *A* tensorial exactor *is a surjective, continuous, open and closed mapping* $\mathbf{E} : \mathcal{T}_{\mathfrak{R}} \times \mathcal{M} \to \mathcal{T}$ *(the space of exact tensor-fields on spacetime), such that, for every* $(m, n) \in (\mathbb{I}^{\uparrow})^2$ *and* $q \in \mathcal{M}$,

1. $\mathbf{E}[\mathcal{T}_{\mathfrak{R}}^{m,n} \times q] = \mathcal{T}^{m,n}|_q$
2. *for* $\lambda \in \mathcal{T}_{\mathfrak{R}}^{m,n}|_q$, $\mathbf{E}(\lambda, q)$ *is in the 4-sphere of possible values contained in* $\mathcal{T}^{m,n}|_q$ *defined by the magnitude* $\pi_1(\lambda)$ *of* λ *and the* Σ*-norm of its inaccuracy* $\pi_2(\lambda)$

As with the definition of the scalar exactor, the second condition captures the requirement that the idealized exact value ought to lie within the bounds of possible inaccuracy that λ determines.

Definition 4.5.2.13 *Given a shrinker* $\mathbf{S}$, *an* (m, n)-tensorial $\mathbf{S}$-distiller *is a mapping* $\mathbf{d}_s : \widehat{\mathcal{T}}_{\mathfrak{R}}^{m,n} \times \mathfrak{P}_s[\mathcal{C}] \times C \to \mathcal{T}_{\mathfrak{R}}|_c$ *such that, for* $\lambda \in \widehat{\mathcal{T}}_{\mathfrak{R}}^{m,n}$ *and* $q \in C$,

1. $\mathbf{d}_s(\lambda, \mathbf{S}^{-1}[q], q) \in \mathcal{T}_{\mathfrak{R}}^{m,n}|_q$
2. $$\inf_{S \in \mathbf{S}^{-1}[q]} \{||\lambda(S)||\} \;\leq\; ||\mathbf{d}_s(\lambda, \mathbf{S}^{-1}[q], q)|| \;\leq\; \sup_{S \in \mathbf{S}^{-1}[q]} \{||\lambda(S)||\}$$

Note that we use here the Σ-norm, $||\lambda(S)||$, for elements of $\widehat{\mathcal{T}}_{\mathfrak{R}}^{m,n}$.[83] Note also that, in so far as there is no natural isomorphism between inaccurate tensor-spaces over different scraps of a decoupage, as there is none for exact tensor-spaces over different points of spacetime, we must include explicitly in the domain of the distiller the points of the canvas over which the fields are defined.

Definition 4.5.2.14 *A* tensorial lens *is an inaccurately linear bijection* $\mathbf{L} : \widehat{\mathcal{T}}_{\mathfrak{R}} \to \mathcal{T}_{\mathfrak{R}}$ *such that*

1. $\mathbf{L}$ *is bounded and stable in the operator-norm*
2. *there exists an ordered pair* $(\mathbf{S}, \mathbf{d}_s)$ *consisting of a shrinker* $\mathbf{S}$ *and an* $\mathbf{S}$*-distiller* $\mathbf{d}_s$ *such that, for all* $\lambda \in \widehat{\mathcal{T}}_{\mathfrak{R}}$, $\mathbf{L}(\lambda)$ *is the* $\mathbf{d}_s$*-distillate of* λ
3. *the restriction of the action of* $\mathbf{L}$ *to a k-bounded, convex, tensorial sub-algebra of* $\widehat{\mathcal{T}}_{\mathfrak{R}}$ *is an inaccurately k-convex bijection into* $\mathcal{T}_{\mathfrak{R}}$

Similarly,

Definition 4.5.2.15 *A* tensorial polarizer *is a linear bijection* $\mathbf{P} : \mathcal{T}_{\mathfrak{R}} \to \mathcal{T}$ *such that*

1. $\mathbf{P}$ *is bounded and stable in the operator-norm*
2. *there exists an exactor* $\mathbf{E}$ *such that, for every* $\lambda \in \mathcal{T}_{\mathfrak{R}}$ *and every* $q \in C$, $\mathbf{P}[\lambda](q) = \mathbf{E}(\lambda(q))$
3. *the restriction of the action of* $\mathbf{P}$ *to a k-bounded, convex, tensorial sub-algebra of* $\mathcal{T}_{\mathfrak{R}}$ *is an inaccurately k-convex bijection into* $\mathcal{T}$

Finally, one has

Definition 4.5.2.16 *A* tensorial idealizer *is a linear, bijective mapping* $\mathbf{I} : \widehat{\mathcal{T}}_{\mathfrak{R}} \to \mathcal{T}$ *such that there exists an ordered quadruplet* $(\mathbf{L}, \mathbf{P}, \widehat{\mathbf{L}}, \widehat{\mathbf{P}})$ *consisting of a tensorial lens* $\mathbf{L}$, *a tensorial polarizer* $\mathbf{P}$, *an exact tensorial lens* $\widehat{\mathbf{L}}$, *and a colored tensorial polarizer* $\widehat{\mathbf{P}}$, *for which*

1. $\mathbf{L}$ *and* $\widehat{\mathbf{L}}$ *are co-focused*
2. $\mathbf{P}$ *and* $\widehat{\mathbf{P}}$ *are co-polarized*
3. *for* $\lambda \in \widehat{\mathcal{T}}_{\mathfrak{R}}$
$$\mathbf{I}[\lambda] = \mathbf{P} \circ \mathbf{L}[\lambda] = \widehat{\mathbf{L}} \circ \widehat{\mathbf{P}}[\lambda]$$

83. To be more precise, we are using the Σ-norm defined on individual points of a tensorial field rather than that defined on sections of the field itself. The restriction of the latter to the former is straightforward.

Not only can we construct in this way all smooth, exact scalar and tensorial fields out of colorings, but, because we can use the smooth colorings to approximate to analytic colorings, we can, in the limit, construct all tensorial and scalar fields ordinarily used in theoretical physics to model physical fields. It would have been awkward, indeed, if we could not have done so, as, rightly or wrongly, [84] analytic fields are the stock in trade of the physicist, both theoretically and experimentally, and it would have been a severe difficulty if they could not have been used. Even had we not been able to recover them, however, I do not think that would have shown the method I use here to be wholly unjustified. The game, after all, is not in the end to recapitulate the functions that are in fact used in the practice of physics; it would have sufficed for our purposes merely to have shown that one can construct *some* fields or other that would appear to be acceptable and sufficient for use in physical investigations, since, again, the aim of this paper is only to show that *some* such logically sound reconstruction is possible.

One may have hoped that imposing "natural" conditions on possible procedures of focusing and polarizing would have sufficed for picking out a unique idealizer, or even, conversely, that imposing "natural" conditions on possible procedures of idealization would have sufficed for picking out a unique set of lenses and polarizers. I see no way of doing this. The process of idealization and approximation in physical science seems to me to be irremediably a pragmatic work of art, guided by a pragmatic artist's intuition, informed by his æsthetical predelictions and constrained by the needs of particular investigations.

4.5.3 An Inaccurate, Well Set Initial-Value Formulation

Idealizers and approximators are the tools we will use to try to make sense of the relation of kinematical equivalence between an inaccurate and an exact theory. Before we can apply them directly, however, we need to clarify the nature of the relations that may hold, on the one hand, between partial-differential equations formulated over inaccurate, colored fields and their solutions, to, on the other, partial-differential equations formulated over exact fields on spacetime and their solutions. It is not yet clear, for instance, what it may mean to require that the inaccurate, colored quantities of the inaccurate theory satisfy the "same" differential and algebraic contraints on their values as do their counterparts in a kinematically equivalent exact theory.

In this section, therefore, we will treat the initial-value formulation of partial-differential equations on colored, inaccurate fields. We will do so at first from a formal point of view, without worrying about the physical content of the constructions. We will then attempt to use the machinery developed in the previous section to explore possible ways of translating such equations into partial-differential equations on exact fields and of relating the solutions of the inaccurate equations to those of the constructed exact ones, and vice-versa. Defining an initial-value formulation for colored, inaccurate equations and characterizing what it may mean to say that it is well set do not pose any serious problems. Exploring the relations of these equations and their solutions to those of the constructed ones, however, pose significantly more involved technical issues, which we will address in this paper with only the most minimal of sketches; they are addressed at length in Curiel (2011).

First, recall the classical notion of a well set initial-value formulation for partial-differential equations over fields on a relativistic spacetime; while the following characterization does not exactly match any other I have seen in the literature, it articulates what we require in this paper while

84. My vote: wrongly. It's not as though we can solve in closed form any mildly versimilitudinous equation of motion for any vaguely realistic model of a physical system anyway. We can't even solve the Schrödinger equation for the isolated Hydrogen atom, but rather are forced to rely on perturbative expansions in a power-series of spherical harmonics and the like. I feel strongly that the focus, in teaching physics, on overly idealized, simplified equations with perspicuous, analytic solutions can easily lull people into a false sense of what it is and is not possible to accomplish with physical theory.

remaining true enough to standard accounts to come to very much the same thing.

Definition 4.5.3.1 *An exact initial-value formulation is an ordered quadruplet consisting of*

1. *a differentiable manifold*
2. *an exact partial-differential equation over a family of exact fields on the manifold*
3. *a non-trivial subset of an appropriate hypersurface in that spacetime (the* domain of initial data)
4. *a specification of values on the given portion of the hypersurface for each field in the given set and for the appropriate number of their derivatives if any (the* initial data*)*

Given an initial-value formulation, its *domain of dependence*, roughly speaking, is the maximal subset of the manifold on which the dynamical evolution of the system may be uniquely determined by the evolution, as modeled by a solution to the partial-differential equation, of initial data off the domain of initial data.[85]

Definition 4.5.3.2 *An exact initial-value formulation is* well set *if*

1. *there exists a unique solution to the equation in the domain of dependence satisfying the initial data*
2. *that solution is stable (in a certain technical sense) under small perturbations of the initial data*

We now turn to constructing the analogous definition for colored, inaccurate fields. In order to consider partial-differential equations, we must have derivative operators to formulate them with. The constructions of §§4.4.4 and 4.4.6 provide us with these, and we may easily use them to write down inaccurate, colored partial-differential equations willy-nilly. Next, in order to have access to classical results on the solutions to exact partial-differential equations in defining a well set initial-value formulation over inaccurate fields—or, more precisely, to be able to use all the classical results on exact structures for giving reasonably simple proofs of the analogous results for inaccurate structures—, I defined in Curiel (2011) distribution-like functionals on test-spaces of inaccurate fields, and showed they had all the required properties needed for use in constructing the analogue of Sobolev spaces of inaccurate fields. Members of such spaces are the natural candidates to serve as initial data for an initial-value formulation of inaccurate partial-differential equations, as well as the natural spaces in which solutions to the equations may be found. The definition of the proper domain for initial data requires construction of an appropriate region of a manifold to serve as the analogue of the domain if initial data, which I called the *convex hull of an infimal decoupage*.

Definition 4.5.3.3 *An inaccurate, colored initial-value formulation* $\mathfrak{i}$ *is an ordered quadruplet consisting of*

1. *a differentiable manifold whose elements are canvases of an ordinary manifold*
2. *an inaccurate partial-differential equation over a set of colorings on a decoupage of a canvas in that manifold*
3. *a convex hull of an infimal decoupage on a canvas in the manifold (the* the domain of initial data*)*
4. *a specification of values, on the given domain of initial data, for each coloring in the given set and for the appropriate number of their derivatives, if any, as determined using the fixed, derivative operator (the* initial data*)*

85. See Geroch (1970a) for a thorough discussion of this notion in general relativity.

One then defines the *domain of dependence* of the domain of initial data in a way naturally analogous to that in the theory of exact partial-differential equations, as the maximal decoupage on which the dynamical evolution of the system may be uniquely determined by an initial-value formulation. We now define a well set, inaccurate initial-value formulation in more or less the same way as was done in the exact case.

Definition 4.5.3.4 *An inaccurate initial-value formulation is* well set *if*

1. *there exists a unique solution to the equation in the domain of dependence satisfying the initial data*
2. *that solution is stable (in a certain technical sense) under small perturbations of the initial data*

It is clear how to apply the schema of these definitions to the case of an exact coloring, and to that of an inaccurate field on an ordinary manifold, as opposed to a coloring.

In the same way as in the ordinary case, one can classify these partial-differential equations (at least the quasi-linear ones, which are the only ones we consider in this paper) into hyperbolic, parabolic and elliptic, and show that the solutions to those of parabolic or elliptic type are analytic fields, whereas solutions to those of hyperbolic type possess the same properties as in the exact case, most importantly that they have well set initial-value formulations, that their characteristic wave-fronts propagate with finite speeds, and that discontinuities in initial data propagate into the solutions.[86] One has, for example,

Theorem 4.5.3.5 *Only inaccurate, colored, hyperbolic partial-differential equations have well set, inaccurate initial-value formulations on decoupages in any given spacetime. No inaccurate, colored parabolic or elliptic partial-differential equation has a well set, inaccurate initial-value formulation.*

This theorem does not rule out the possibility that inaccurate parabolic or elliptic partial-differential equations have well set, inaccurate boundary-value or mixed problems (analogous, *e.g.*, to the classical Neumann or Dirichlet problem). In this paper, however, we are interested only in the initial-value formulation of partial-differential equations. Note as well that even hyperbolic equations may have no well set initial-value formulations for initial data posed on certain regions of a spacetime. Possession of a well set inaccurate initial-value is relative to the spacetime and the particular regions of the spacetime on which the equations are formulated,

In order to state the sought-after results of this section, the primary remaining problem concerns the construction of exact equations from inaccurate ones and the relating of the solutions of the one to those of the other in a meaningful way. In a naive sense, we already know how to translate inaccurate, colored partial-differential equations, posed in terms of an inaccurate, colored, covariant derivative operator, into partial-differential equations on exact scalar fields in a more or less direct way, by fixing an idealizer and applying it individually to all the fields used in formulating the inaccurate partial-differential equation on colorings, simultaneously transforming the inaccurate covariant derivative operator into its idealized counter-part. So much is straightforward. On the face of it, however, nothing guarantees that the solution of the idealized partial-differential equation bears any substantial relation to that of the original, inaccurate one. Without a strong relation with clear physical significance between the two, the sort of analysis of a regime I advance would fall on its face. Indeed, we demand that the strongest relation possible hold between them: that the process of idealization commute with the solution of partial-differential equations. In other words, we demand that, if we first solve the inaccurate partial-differential equation and then idealize the

86. It is this last property that plays a decisive role in our ability to individuate collections of otherwise undifferentiated and commingled physical fields into discrete, separate physical systems identifiable over extended periods of time—the boundary of a body may be thought of as a discontinuity propagating in the solution to the governing system of hyperbolic equations of dynamic evolution in the region.

solution, we end up with the same exact scalar field as we would have, had we first idealized the partial-differential equation itself and then solved that idealized equation.

To show that this is indeed the case, I constructed in Curiel (2011) a method of interleaving among the different stages of idealization, treated as the composition of the actions of a lens and a polarizer, processes for approximating partial-differential equations and their solutions with sequences of convergent, finite processes, one such process for each of the possible species of partial-differential equation at issue, those on inaccurate and exact colorings and those on inaccurate and exact fields. For an inaccurate, colored partial-differential equation, for example, the process consists of approximating the decoupage by a sequence of finite, ever-denser, lattice-like structures the elements of which are spacetime points, and approximating a coloring on the decoupage by its restriction to the elements of the lattice-like structures. Given an inaccurate, partial-differential equation over a coloring on the decoupage, one writes the analogous algebraic, finite-difference equation on each member of the sequence, and proves that it has a solution. The construction then ensures that, in the limit as the finite lattice-like structures converge, in a rigorous sense, to the infimal decoupage, the solutions to the algebraic, inaccurate, finite-difference equations converge to a solution of the original, inaccurate partial-differential equation on the decoupage (assuming it has at least one, which we do assume hereafter without further comment). I proceeded to show how to intercollate these approximative methods with the processes of focusing and polarization, from which procedure it followed that the space of inaccurate colorings, the equations over them and the solutions to those equations converge, respectively, to the space of exact fields, partial-differential equations over them and their solutions, in the joint limits as one takes the finite-difference equations into partial-differential equations, as one takes the decoupage into a canvas and as one takes the inaccurate quantities into exact ones; this convergence, moreover, takes place in such a way that, given an inaccurate partial-differential equation over colorings, a well set initial-value formulation of that equation and a solution for that fixed initial data, the solution to the idealized partial-differential equation, as derived from the approximative process, is in fact the idealized solution to the original inaccurate, colored partial-differential equation.

I modeled the approach after that standardly employed in approximating exact partial-differential equations,[87] though the differences between inaccurate, colored and exact structures demand not insignificant differences between the methods of approximation. In standard accounts, one starts with exact Banach spaces of fields over an ordinary, differential manifold (*e.g.*, Sobolev spaces of some appropriate order), and discretizes these spaces by restricting the values of the fields to finite lattices on the manifold. In other words, in a certain sense, the discretized spaces may be considered restrictions of regions of the manifold, or, if you like, the discretized spaces may be considered the images of a family of finite, injective mappings of a region of the manifold onto itself, such that in the limit as the lattices get bigger and bigger the spaces of discrete fields become dense (in a certain technical sense) in the exact Banach spaces. In this paper, on the other hand, we are starting with a manifold built up from *subsets* of an ordinary manifold (in this case, the decoupage of scraps built up from a canvas in the spacetime manifold), on which are defined Banach spaces of fields that may be considered generalizations of ordinary Banach spaces of fields. One then considers particular ways of "shrinking" the elements of these manifolds and restricting the generality of these fields via exacting approximations in such a way that, in a certain limit, there results a collection of points dense in the original subset of the ordinary manifold, and a family of fields that may be considered a dense subset of the ordinary Banach space of fields over that subset. In the process, moreover, we must transform as well the range of the fields involved in the equations.

The results may be summed up in

Theorem 4.5.3.6 *Fix an idealizer and a hyperbolic, colored, inaccurate partial-differential equation with a well set initial-value formulation on a decoupage over a canvas in a manifold. If one first*

87. See, *e.g.*, Wloka (1987).

solves the equation and then applies the idealizer to the inaccurate, colored solution, one arrives at the same exact field on the canvas as if one had first applied the idealizer to the inaccurate, colored partial-differential equation to derive an exact partial-differential equation on the canvas and then solved that idealized equation.

This theorem does not assert that the idealized solution mentioned in the theorem is independent of the idealizer chosen. It asserts only that the exact solution, for a fixed idealizer, is unambiguously determined.

This method of addressing the problem, by intercollating finite, discrete, convergent approximations among the stages of idealization, has other virtues as well. The approximative methods developed allow one to state and prove in a natural way results on the stability of solutions to both inaccurate and exact partial-differential equations and the stability of relations among them. This expressive power will come in handy in §4.5.4 below, in formulating what it may mean for the initial-value formulation of the partial-differential equations comprised by a physical theory to be well set in a physically relevant sense.

4.5.4 A Physically Well Set Initial-Value Formulation

In §4.5.3, we sketched the definition of a well set, inaccurate initial-value formulation from a strictly formal point of view, without relation to physical theory. In this section we will attempt to take account of the constraints a physical theory's possession of a regime places on the initial-value formulation of a partial-differential equation, in so far as the partial-differential equation is part of that physical theory and the initial-value formulation conforms to the regime.[88] Because we now are dealing, in contradistinction to the strictly formal case of §4.5.3, with partial-differential equations modeling the dynamical evolution of physical systems, we require a substantive relation between an inaccurate initial-value formulation's being well set and the physical well-setness, as it were, of exact initial-value formulations we construct from the inaccurate one. To work out an appropriate one, we will need to state with a sufficient amount of rigor and precision what it means for an exact theory and an inaccurate theory to be kinematically equivalent. With the results of §4.5.3 at our disposal, we are now in a position to do this. An inaccurate theory and an exact theory are kinematically equivalent if there exists an idealizer and its related approximator such that,

1. applying the idealizer to the partial-differential equations the inaccurate theory comprises yields the partial-differential equations comprised by the exact theory
2. applying the approximator to the algebraic and differential constraints on the values of the physical quantities treated by the system yields a system of inaccurate, colored algebraic and differential constraints automatically satisfied by all the relevant entities in the inaccurate theory

It must also be the case that the metrical conditions encoded in the collection of infimal decoupages on which the fields of the inaccurate theory may be defined agree with those articulated in the exact theory. Roughly speaking, an inaccurate theory and an exact theory are kinematically equivalent if one can transform each into the other by applying an idealizer or approximator to everything in the theory the operator can act on. One has, as the summation of much of §4.4 and of this section up to here,

Theorem 4.5.4.1 *The relation of being kinematically equivalent is stable under "small" perturbations of the given idealizer.*

88. To do full justice to this subject, we would need also to consider ways of moving from finite sets of inaccurate data to inaccurate fields, as required for a more complete model of the interaction between the practice of the experimentalist and that of the theoretician. That, however, is beyond the scope of this paper. See Curiel (2011) for a brief sketch of how such an account might go, were we to essay one.

A complete argument for this claim is tedious and not worth the effort the details require, at least for the purposes of this paper.[89] It makes intuitive sense, though. An idealizer takes an inaccurate structure and renders from it an exact structure satisfying the kinematic constraints encoded in the inaccurate structure. There are, in the event, many exact structures that will satisfy these constraints, given the looseness of fit provided by the inaccuracies of the regime. Perturbing the given idealizer slightly enough will, because idealizers are ω-stable, yield one of these other exact theories.

Now, we turn to the characterization of a physically well set initial-value formulation for exact partial-differential equations comprised by an exact theory with regime, before moving on to the case of the inaccurate, colored initial-value formulation of partial-differential equations in an inaccurate theory. Fix an exact physical theory with its kinematical regime $\mathfrak{K} = (\mathfrak{e}, \mathcal{E}, \mathfrak{k}, \mathfrak{m}_k, \mathcal{K})$, its exact partial-differential equations $\mathfrak{E}$ and a well set initial-value formulation $\mathfrak{i}$ for those equations.

Definition 4.5.4.2 $\mathfrak{i}$ conforms *to* $\mathfrak{K}$ *if*

1. *the domain of initial data is a* $\mathfrak{k}$*-appropriate subset of a* $\mathfrak{k}$*-appropriate observatory*
2. *the values of* $\mathfrak{e} \cup \mathcal{E}$ *satisfy* $\mathfrak{k}$ *in the domain of initial data*
3. *the preparation and measurement of the initial data occur in conformity with some subset of* $\mathcal{K}$

We will also say that the initial-value formulation *kinematically conforms* to the regime. Although this definition may appear (somewhat) innocuous, it imposes a severe restriction on the initial-value formulation of a physical theory, in so far as that initial-value formulation will conform to the kinematical regime of the theory, in an important way: the domain of initial data must be compact. The remarks in §4.3.3, after the definition of a kinematically admissible observatory, foreshadowed this fact.

Definition 4.5.4.3 *An initial-value formulation of an exact theory with a regime is* physically well set *if it is well set and it conforms to* $\mathfrak{K}$*.*

We can now offer the primary definition of this section, and state the primary results.

Definition 4.5.4.4 *A* physically well set, inaccurate initial-value formulation *of an inaccurate theory consists of a well set, inaccurate, colored initial-value formulation and an idealizer, such that*

1. *the exact initial-value formulation yielded by applying the idealizer to the inaccurate one is physically well set in the exact theory with regime that is equivalent to the given inaccurate theory*
2. *the process of idealizing the equation is uniformly* ω*-stable, in the sense that for every* $\omega > 0$ *there exists a* $\delta > 0$ *such that the idealization of a* δ*-perturbation of a solution to the inaccurate equation yields an* ω*-bounded perturbation of a solution to the idealized equation*
3. *the process of idealizing the equation is stable, in the sense that for every* $\omega > 0$ *there exists a* $\delta > 0$ *such that using a* δ*-perturbation of the original idealizer to construct the exact equation and its solutions yields a process of idealizing that is uniformly* ω*-stable*

The requirement that the relations among data and theory not depend in an essential way on the choice of idealizer suggests the last two conditions. There follows from this definition, in particular from the compactness of the domain of initial data required for an initial-value formulation to be physically well set with regard to a theory, in conjunction with theorem 4.5.3.5,

89. As ever, see Curiel (2011) for the details.

Theorem 4.5.4.5 *Only inaccurate, colored, hyperbolic partial-differential equations have physically well set, inaccurate initial-value formulations on decoupages over canvases in a given spacetime. In particular, no inaccurate parabolic or elliptic partial-differential equation has a physically well set, inaccurate initial-value formulation.*

This result says, in effect, that parabolic and elliptic partial-differential equations are acceptable for the modeling of physical systems *to the extent* that we are willing to accept the risk of almost-null thunderbolts disrupting our experiments, as it were. Since we in fact live our entire lives having tacitly, for the most part, accepted this risk (at least, those of us without psychoses or acute neuroses of some stripe), this seems on the face of it a reasonable assumption to make in the theater of physical science. This attitude, however, while surely good enough in the actual performance of experimental physics, cuts no ice in attempting to understand the relation of theoretical physics to that performance. According to theoretical physics, if the prediction of an initial-value formulation of a parabolic or elliptic partial-differential equation does not hold good when the domain of initial data forms a proper subset of a spacelike hypersurface, correlations between the initial data given and the values of quantities far distant on a spacelike hypersurface containing the domain of initial data may bear the blame—if one had from the start enlarged the domain of initial data to have included that distant, in the event relevant data in the initial data, the prediction would have held good. This dilemma—bad equation or bad data?—does not arise, at least not in this form, when one uses hyperbolic partial-differential equations. If one has collected the data as carefully as possible and the prediction still does not hold good, then the problem lies with the partial-differential equation itself in the hyperbolic case. One never has this iron-clad (aluminum-foil clad?) guarantee with parabolic or elliptic partial-differential equations.

This theorem, with its suggestion that a partial-differential equation may have a physically well set initial-value formulation in one spacetime but not in another, raises an interesting question: what it may mean, in a physical sense, to speak of the "same" equation posed on different spacetimes. Take, for example, the relativistic Navier-Stokes equations. It is not clear to me, on its face, what one means in referring to *the* relativistic Navier-Stokes equations *simpliciter*. Such a system must be formulated on a particular manifold having a particular topology and differential structure, using a particular affine structure. Presumably, in speaking of the *same* equations on two different spacetimes—the Navier-Stokes equations on each—we often intend something like the following: we mean that system on each that has the syntactic form of the system (4.2.2.3)–(4.2.2.6), with the appropriate derivative operator used in each case. So much causes no trouble, as it poses only a problem of orthography. The problem arises in the attempt to understand each of these systems as equations of the same *semantic form*, as it were.

Let me try to clarify what I am gesturing at. Think of the difficulties early investigators in general relativity faced in trying to generalize a system of equations from Minkowski spacetime to a generic curved spacetime by way of an application of the principle of equivalence. How ought one introduce the "coupling of the curvature" to the terms of the equation, to get the proper form? One cannot rely on mere syntactic equivalence of the equations (substitution of '∇_a' for '∂_a'), if one demands only that the equations in curved spacetime "reduce to" those in Minkowski spacetime "in the limit as curvature vanishes", as all terms involving curvature will be identically zero in Minkowski spacetime. Equation (4.2.2.3), for instance, could be written either as

$$\nabla_m(\nu\xi^m) = 0$$

or as

$$\nabla_m(\nu\xi^m) + R_{abcd}R^{abcd} = 0$$

in a spacetime with non-trivial curvature, and still reduce to

$$\partial_m(\nu\xi^m) = 0$$

in Minkowski spacetime. Determining the correct form in spacetimes with non-trivial curvature depends essentially on a determination of the way the fluid's particle-number density *physically* depends on the curvature.

In the same way, in trying to determine which system of equations on two different spacetimes model the same sorts of physical systems, it will not suffice to demand only that the equations have the same syntactic form in the absence of a physical investigation. For all we know, the sorts of systems that we want, for various reasons, to identify as being "the same" in two different spacetimes couple to the curvature in a way that does not manifest itself except in spacetimes whose Riemann tensors have some *outré* property, or the systems may depend on the topology of the spacetime manifold, in the sense that non-trivial terms depending on, *e.g.*, its Euler characteristic must be included in the equations of the system—indeed, the systems may exhibit behavior that is, for one reason or another, best modeled using the terms of almost any mathematical structure one can imagine accruing to a model of the system. It is only by observation that such issues can be settled.

How do Navier-Stokes fluids behave in areas of non-trivial curvature, and in spacetimes with non-trivial Euler characteristic?—which is to ask, what system of equations most accurately models their behavior, within the proper regime? The answer to that question arbitrates questions as to the proper form for the Navier-Stokes equations. This, then, raises the questions of how one identifies a "Navier-Stokes fluid", if not by susceptibility of modeling by whatever it is we settle upon as the Navier-Stokes equations—and here, I think, is the place where causal just-so stories find their place, in a limited way, in physics. "We know it's a Navier-Stokes fluid because it is the concomitant, in the expected place, at the expected time, with the expected result, of this sort of coupling with this experimental apparatus. . . "—and if it fails to obey either the "regular" or a recognizably altered version of the Navier-Stokes equations, then we're off and running, back to the races. This is, in general, no mean feat, especially when one is attempting to account for the features of physical phenomena that may "depend", in some sense or other, on global properties of a spacetime. One cannot move to a region of non-zero Euler characteristic in a spacetime that has a zero one! [90]

This discussion goes some way, I hope, towards explaining why I felt it necessary to include in the eighth component of the octuplets in volume twelve of our mythologizing books described in §4.5.1, and in the definition of well set initial-value formulations in this and the previous section, explicit reference to the spacetimes on which all these structures are imposed. Some spacetimes will not admit $\mathfrak{m}_k$-appropriate observatories for some kinematical regimes, so it seems not out of place to specify from the start which spacetimes one has in mind with regard to the application of one's theory. It would be interesting to formulate some precise questions along these lines, to attempt to determine, for example, what sorts of constraints common spacetime models (Schwarzschild, FRWL, *et al.*) impose on the admissibility of potential models of laboratories.

4.5.5 Maxwell-Boltzmann Theories

The word 'theory' in physics and in philosophy has various, sometimes partially overlapping, meanings. I will focus here on its usage in physics. In physics, for instance, one may speak of classical Navier-Stokes theory, of the theory of stellar structure, of quantum field theory and of thermodynamical theory. Roughly speaking, these examples descend in order of clear delineation of subject matter, clear delineation of physical phenomena to which they are applicable, clear delineation of experimental techniques used to probe those phenomena, and clear delineation of generally accepted mathematical structures and illustrative techniques used in solving problems in their re-

90. I believe considerations of this sort show why David Lewis's and Saul Kripke's conceit of some sort of *a priori* identification of physical systems "across possible worlds" is fair nonsense—or, if you will, given the uncritical acceptance often accorded the idea, unfair nonsense.

spective domains.[91]

My use of 'theory' in this paper does not exclude those traditional uses; rather, it generalizes them. For my purposes, a theory is any more or less formal structure that contains both a system of partial-differential equations and a dynamical regime for the application of those partial-differential equations in the quantitative modeling of physical phenomena. Such a structure may include bits and pieces of other structures more commonly conceived of as integral theories, mixing and matching as it chooses, as happens in the modeling of actual experiments, so long as the sum total has a single, consistent dynamical regime. An example would be the use, in the elements of the regime of a theory, of just so much of quantum field theory as required for a Planckian treatment of electromagnetic radiation when modeling the measurement of systems having temperatures above 1063° Celsius.[92]

In my usage, the intended sense of 'theory' is in many ways an ambiguous, even a nebulous concept—it is not *a priori* clear, for instance, whether a theory that treats joint electromagnetic and thermodynamic phenomena for which the Callendar equation suffices ought to constitute a theory different from one modeling essentially the same phenomena at similar and at lower temperatures, requiring the use as well of the van Dusen equation. I think this ambiguity is, in any event, no worse than that accruing to the standard usages. More to the point, however, I think this ambiguity underscores in a salutary way an important fact about physical theory: how one delineates a particular set of structures for the modeling of a more or less well delineated family of phenomena—the only substantive issue, I think, one can dispute about concerning the application of the term 'theory'—is a profoundly pragmatic procedure. There is nothing *a priori* about it. Whether, for instance, the two structures differing almost only by the inclusion of the van Dusen equation constitute different theories will depend on one's purposes in using or analyzing the structures of the theories, and why one cares in the first place about distinguishing various structures, more or less formal, as different theories.

This point relates to the Carnapian one about the pragmatic character of the selection of a linguistic framework. It differs from it in that here, one may say, one is attempting to decide how to *differentiate* linguistic frameworks, one from another, rather than to distinguish among for various purposes and select one from an already existing family of frameworks. I should emphasize that my use of 'pragmatic' in this paper is not technical in the slightest. It bears no particular relation to any of its variegated uses by philosophers from Peirce to James to Carnap to Goodman to Quine and beyond. I (am trying to) use it in its diurnal, pedestrian—its pragmatic, if you will—sense: that pertaining to the choosing of courses of action by the weighing of alternatives and striking what one hopes is a satisfactory balance among all the competing and confluent objectives.

To guard against a possible misconstrual of my arguments and conclusions, I want to take a brief pause in the flow of this section to make clear that I think these arguments and conclusions serve directly to controvert positions such as those advocated in, *e.g.*, Cartwright (1999); they in no way support them. In particular, I am talking about Cartwright's arguments and conclusions to the effect that science—scientific theories—consist of *nothing more* than families of more or less disparate, unrelated schemata of ways of modeling particular kinds of experiments. For instance, there's the schema using the quantum S-matrix formalism to model the scattering of fundamental particles, which bears (on her view) no particular relation to the schema using standard perturbative

91. I would not want to cavil about my proposed ordering—I'd be more than happy to rearrange it in the face of even mildly convincing argument. I care only that the point I am trying to make is clear.

92. I am glossing over an important distinction, that between, on the one hand, what theoretical apparatus forms part of a theory *per se*, and, on the other, what forms part of the concomitant but tacit theoretical apparatus required for a theory of measurement of the quantities modeled by the theory in experiments testing or employing the theory at issue—the apparatus forming part of the theory *per accidens*, if you will—this is a complex question, which I will not be able to address here.

In any event, though it may appear superficially similar, this issue is not directly related to the distinguishing in definition 4.3.3.1 between $\mathfrak{c}$, the quantities directly treated by the theory, and $\mathcal{E}$, the environmental factors relevant to the applicability of the theory.

techniques (say, expansion in spherical harmonics) to model, at the quantum level, the dynamic evolution of a Hydrogen atom in a static electric field. On the contrary, I believe my arguments show the profound, inextricable connections among such different theoretical models—the idea of a regime of applicability gets off the ground in the first place only to the extent that one can bring to bear on each other superficially disparate theoretical structures, as in the application of the Planckian treatment of electromagnetic radiation to thermodynamical thermometry. One must have already in hand a well worked out theoretical apparatus with understood ramifications into other such structures going far beyond an enumeration of highly schematized mathematical models in order to ascertain with confidence the applicability of a given theoretical structure to a particular experimental arrangement, just as one must have already in hand the practical experience of many performances of particular kinds of experiments in order to conclude that their outcomes accord with or contravene the predictions of those theories, and whether the fault, in the case of contravention, lies with the experimental arrangement or performance, or perhaps rather indicates the presence of some novel phenomena not accounted for by the given theory. [93] If anything, this paper serves as an argument for a sort of Carnapian pluralism—one chooses a "framework" (theory with a regime) based on *pragmatic* criteria—simplicity, ease of use, facility for physical insight, elegance, what have you. It is a striking, brute fact about physics, perhaps the most singular fact about physics as a human enterprise—a fact that deserves puzzling over—that, in almost all known cases, there is a *single* candidate that jointly satisfies all these interests to a greater degree, for almost all active investigators in the field, than any competitor.

To pick up the thread of the argument of this section, I contend that, in order to be viable as a *physical* theory, a theory must have, at least in principle, a regime of applicability (or something very like it) allowing for physically well set initial-value formulations in some (model of) spacetime or other. Before trying for more precision in the definition, I want to point out that, even as it stands, it can make itself useful. Many comments and questions about physical theories themselves and about their inter-relations can be formulated naturally in its terms. Take, for example, the puzzling case of the Bernoulli Principle, which has as a consequence that a body of water moving with a uniform, constant velocity has a static, hydrodynamic pressure less than that of an otherwise identical body of water at rest. This principle manifestly controverts the more fundamental and dearly held Principle of Gallileian invariance: consider two channels parallel to and at rest with respect to each other; one, A, contains water still in relation to its banks; the other, B, contains water moving uniformly in relation to its banks. Is it the case, then, that, the water in A, and thus the 2 channels, are at rest and the water in B is moving? That the water in B is actually still, but the water in A as well as the two channels are in motion? Or that all four are in motion, so arranged as to give the described relation among them? According to Gallileian invariance, no experiment we could perform should differentiate these possibilites from each other. In fact, however, actual experiments *do* differentiate them—a shower curtain's motion inward toward the stream of water when a shower is first turned on provides a simple example, as does the fact that a blocked garden-hose will burst whereas one whose water is under the same motive force but is flowing will not. In the terms of the idea of a regime of applicability, we would describe the situation by saying that, in the hydrodynamical regime (for gross measurements of fluid velocity), no well-defined quantity will manifest any behavior that is not Galileian invariant, excepting only this, the hydrostatic pressure under conditions in which Bernoulli's Principle finds application. In another regime, a finer-grained one, we expect there will be defined only Galileian invariant quantities, in terms of which one can show why satisfaction of the Bernoulli Principle appears to be a violation of that invariance, but is in fact not. [94] This suggests that, no matter what else is the case, equality (in some sense) of

93. To guard against a further misunderstanding, let me emphasize as well that I am speaking here only of the case of the application of an already well entrenched theory, not the testing of a novel one, which is a far more difficult case to straighten out.

94. One may, for instance, give an "explanation" of the phenomena as follows. At the hydrodynamic scale, that at

regimes is a necessary condition for the identity of two seemingly different theories.[95]

One may ask, "To which theory in particular does the Bernoulli Principle belong, if it contravenes Galileian invariance, one of the foundations of Newtonian mechanics?" I do not think the question is of interest and perhaps not even, as it stands, sensical. Again, the taxonomy of theories can be argued over as one likes, or even the referent of 'theory', without changing the fact that it makes sense to articulate a system of partial-differential equations and an interpretation of them (which includes a regime) such that the equations under that interpretation adequately model the physical systems whose behavior conforms to that described by the Bernoulli Principle—"model it" in the sense that the theory models not only the particular family of phenomena demarcated by the Bernoulli Principle, but some non-negligible other class of families of hydrodynamic phenomena as well (*e.g.*, laminar flow).[96] This last caveat is an attempt to respect the cantilevered, intermixed, even polygamous character of physical theories (or, if you prefer, of the components of Physical Theory).

Accepting all these difficulties, ambiguities and caveats, I propose the following as a characterization of part of what it is to be a (highly idealized) representation of a physical theory.

Definition 4.5.5.1 *A* Maxwell-Boltzmann theory $\mathfrak{T}_{MB}$ *is a physical theory that includes (at least) an ordered triplet* $(\mathfrak{B}, \mathfrak{I}, \mathfrak{V})$ *such that*

1. $\mathfrak{B}$ *is a set of thirteen canonical volumes, whose contents are as described in §4.5.1*
2. $\mathfrak{I}$ *is a connected set of idealizers, bounded with respect to the operator norm*
3. $\mathfrak{V}$ *is a family of initial-value formulations of the partial-differential equations the theory comprises, each one physically well set with regard to the exact initial-value formulation yielded by application of any of the idealizers in* $\mathfrak{I}$

To be a physical theory, I claim, a theory must be capable in principle of being made a Maxwell-Boltzmann theory. I use the qualifier "Maxwell-Boltzmann" to gesture at the fact that many of the seeds of this notion are already contained in a Maxwell-Boltzmann statistical treatment of thermodynamical phenomena. Accepting, then, definition 4.5.5.1, at least provisionally, one has as its most obvious consequence

Theorem 4.5.5.2 *A Maxwell-Boltzmann theory whose family of initial-value formulations contains at least one member must comprise only hyperbolic partial-differential equations.*

This is one way to make precise the claim that, so far as physical theory goes, only hyperbolic partial-differential equations have well set initial-value formulations. It also explicates in a precise way at least part of the privilege of the role played in physics by hyperbolic partial-differential equations.

4.5.6 The Consistency of Theory and Experiment

We are finally in a position to articulate the primary claim of this paper: in so far as one accepts that the models I have proposed adequately represent logical forms, as it were, of the common

which the Bernoulli principle finds application, each body of water moves uniformly. At molecular scales, *i.e.*, those scales at which the hydrodynamic regime breaks down, the vector-field representing the accelerations of the molecules in the one body differs dramatically in kind from that representing the accelerations of the molecules in the other, in such a way as to yield, when averaged out, a lower hydrostatic pressure in the one than in the other. Because acceleration is a Galileian invariant quantity, we thus recover our dearly held principle.

95. This example suggests another possible virtue of the idea of a regime, no matter how one fleshes it out: the idea of a reduction of one theory to another may naturally be framed in its terms.

96. The question of what makes phenomena "hydrodynamic" over and above susceptibility to satisfactory modeling by theories that we nominate 'hydrodynamic' is fascinating—or, I should rather say, the question whether one can ascribe any meaning to the term relevant to physics over and above this susceptibility. It is also too difficult to address here.

playground of the theoretician and the experimentalist, of their toys and rides, and of the games they play with each other—or at least in so far as one accepts that my proposals show that it is possible that something much better in their spirit can be constructed to represent these things—, one may conclude that there is no inherent contradiction between the practice and the subject-matter of the theoretician on the one hand and the experimentalist on the other, in so far as the entire paper up to this point has served as a constructive proof of this claim. The proof of the claim, moreover, does not depend on the existence of any actual theory or any actual, extended interplay between theoreticians and experimentalists that my models adequately represent. It is enough that such things may be logically modeled within the same, consistent schema, in a manner close enough to the actual proposal of theories and the performance of experiments to have some seeming to it.

Before leaving this subject, I want to emphasize one last time the lack of pretense that any thorough rigor accrues to this claim and its proof. The goal will have been reached if we have achieved a modicum of rigor, in those parts of the subject that can bear it, and slightly more clarity that that in the whole.

4.6 The Soundness of Physical Theory

So far, we have treated theories only in so far as they may be kinematically appropriate for the modeling of physical phenomena, without regard to how well or how poorly their models fare in the fineness and accuracy of their predictions about those phenomena. We turn now to consider these latter issues. We begin in §4.6.1 by examining what it may mean to claim that a theoretical prediction made by a Maxwell-Boltzmann theory does or does not agree with the experimental determination of the value of a physical quantity. This discussion naturally leads to a characterization of the self-consistency of a Maxwell-Boltzmann theory, in a particular form related to the idea of being able to test a theory in an unambiguous way that conforms to its kinematical regime. This will allow us, in §4.6.3, to characterize what it may mean to say that such a theory is *sound*, in the sense of modeling to a desired degree of accuracy the phenomena it purports to treat. We will focus on characterizing the elements a theory must possess in order for one to be able to judge whether or not it is sound—its *regime of dynamical soundness*—and on the sorts of properties those elements must have for one to conclude that the theory is in fact sound. This will lead, in §4.6.4, to a discussion of a peculiar form of under-determination necessarily attendant on any sound physical theory, one which I will be able to summarize with a precise, formal statement.

We will not treat any of these matters with anything near the same degree of thoroughness that we have up till now attempted in investigating the requirements for the kinematical adequacy of a theory. It is, perhaps, a game for another time.

4.6.1 The Comparison of Predicted and Observed Values

In studying the interplay between theory and experiment with regard to the kinematical requirements on physical theory, the most important and difficult issues involved accomodating the inevitable inaccuracies attendant on the determination of the values of quantities in any physical investigation. Now, in studying the issues that must be grappled with in attempting to determine whether a theory is adequate for the modeling of a class of physical systems in the sense of yielding sufficiently accurate predictions about the dynamical behavior of those systems, we must try to accomodate the related inevitability of the deviance of predicted from observed values in the modeling of thoroughly understood physical systems by application of even the most well founded of theories. It is difficult to imagine, for instance, a theory better comprehended and more well founded than Newton's theory of gravitation—it is what we use, after all, to calculate the trajectories of successfully executed manned flights to the moon. It's hard to get more successful than that in science.

Even so, errors inexorably occurred in the calculation of those trajectories, for reasons of widely varying types, theoretical, empirical and pragmatic. Our goal at the moment is to characterize with only the broadest of brush-strokes how this inevitable deviance of predicted from observed values may be treated within the framework we have developed so far.

As with the inevitable ranges of inaccuracy accruing to the experimental determination of the value of any physical quantity, we want to demand that a theory itself, in conjunction with other theories—those treating the measuring instruments employed in a particular experiment, for example—provide means for calculating ranges of admissible deviance of, on the one hand, the predictions of the theory for a particular system from, on the other, the results of measurements made during the course of the actual dynamical evolution of that system. *Admissible* here means nothing less and nothing more than that any measurement not according with the prediction to within that range of deviance ought require that one re-calculate the prediction, attempting to include the influence of factors not yet accounted for, or else that one similarly re-calculate the range of admissible deviance, or else that one repeat the experiment with a finer grain of control over the experimental circumstances, until the difference between the measurement and the prediction does fall within with the range of admissible deviance, or else that one count the experimental evidence as a contravention of the theory. Determining which conclusion to draw, and so which course of action to attempt, in any given case is one of those peculiar games that often cannot be played by either the experimental or the theoretical physicist alone, but will require the active participation of both. [97]

It is important to be clear on how this differs from the calculation of the inaccuracies in the determination of the values of those quantities. Calculating the inaccuracies of a measurement involves no normative judgment; its result is a description of a brute, factual matter. Judging whether or not the value determined by the experiment accords sufficiently well with the value predicted by the theory is a thoroughly normative affair; it is not a description of a brute, factual matter, is not, indeed, a description of a physical state of affairs at all, but rather an assessment of the soundness of our knowledge of the physical world. When I say the former involves "no normative judgement", I do not mean to imply that the aesthetic and pragmatic considerations I have emphasized all along as being in play in the mutual application of theory and experiment are not normative. They are. I mean rather to say only that calculations of inaccuracy are not normative statements. They express no judgments about fineness or suitability or acceptability or what-have-you. Judging the admissibility of a certain deviation of predicted from observed values, however, is inescapably normative, in so far as the physical world does not provide for us criteria to judge the degree of soundness of our knowledge of it in any particular case. There are similarities between the two cases as well. Although judging the soundness of the prediction is a normative affair, determining the deviance itself of the predicted from the observed value is a brute, factual matter along the same lines as the computation of the inaccuracy accruing to the measurement. There are, therefore, two distinct steps in judging the soundness of a theory's predictions as compared to experimental observations. First, one must determine what the deviance is of the prediction from the observation. Next, one must judge whether or not it is admissible. In this section, we will consider only the former issue, postponing the latter one until §4.6.3.

On its face, the idea of *the* deviance of predicted from observed values is not a clear one. Consider trying to calculate such a thing for a scalar quantity. In predicting the value of that quantity for a particular system under particular circumstances, one will produce not an exact scalar but rather an inaccurate scalar, representing the spread of possible values for the scalar within the range of possible inaccuracy of the experiment being modeled. Measuring the quantity will also yield not

97. The discovery in the 1950s of the non-conservation of parity in processes mediated by the weak nuclear force provides a vivid example of the necessity of the participation of both sides. For a lively and gripping account of the episode, see Yang (1961).

an exact but an inaccurate scalar, and, inevitably, a different one. We can compute their algebraic difference readily enough, using any of our three types of operations—physical, psychological or pragmatic—but none of them seems quite right for the job. The pragmatic one will not do, since we are not treating these values as mere numbers, but rather as the representation of the value of a physical quantity. On the other hand, the physical operation will not do either, for these are not representations of physical quantities whose physical combination we are trying to represent by the use of an algebraic operation; these are rather different types of representations of the value of the same physical quantity. Recall that we selected the form for our physical operations based on an analysis of the way that inaccuracies combine and propagate in calculations involving the values of physical quantities, in so far as those combinations of values represent the kinematical and dynamical relations and interactions of the quantitities; in particular, our analysis relied on the fact (or, if you like, the assumption) that these sorts of inaccuracies tend to cancel each other out and so decrease over time. In order for this argument to work, we must assume that, in a typical case, the inaccuracies of all the quantities distribute themselves in a more or less Gaussian form around, respectively, the more or less stable mean of each. It makes no sense to say that the two inaccuracies we are considering here will "tend to cancel each other out over time", because the inaccuracy as determined by the theoretically predicted value does not arise from the physical interaction of actual physical systems, the very variability of which over time allows us to treat inaccuracies as we do. The predicted inaccuracy arises from something like a representation of the Platonic form of the observation—it never changes—not from the actual, physical circumstances of the experiment as it is being performed. We must, it seems, come up with some other way of comparing the predicted and the measured values.

[*** The point: the comparison of the two types of values is not an algebraic operation at all; it is wholly topological ***]

We want to define a way of comparing inaccurate fields, for use as a criterion in determining, without ambiguity, whether some set of measured values, with their associated inaccuracies, falls within the ranges of allowed deviances from the predicted values. Because an approach based on algebraic operations faces formidable difficulties, we will attempt a route with a more topological bent. As with almost every structure proposed in the construction of our model, there is more than one way to do it, and some of the choice must be made on pragmatic grounds, influenced by the demands of the enterprise at hand and guided by taste and predilections. I choose one that seems to me to have clear physical significance in a wide variety of applications, and that is simple to comprehend and simple to apply.

Definition 4.6.1.1 *Two inaccurate scalar fields are* consonant *if*

1. *they share the same support*
2. *at every point of their support, the real intervals representing, respectively, the values of each have a non-trivial intersection*

Of two consonant, inaccurate scalar fields, the first dominates *the second if, at every point of their support, the inaccuracy of the first is greater than that of the second, and their intersection includes the magnitude of (*i.e.*, the midpoint of the interval representing) the first.*

For many pairs of consonant, inaccurate scalar fields, neither will dominate the other. This definition can be extended directly to inaccurate tensor-fields.

Definition 4.6.1.2 *Two inaccurate tensorial fields are* consonant *if*

1. *they share the same support*
2. *at every point of their support, the 4-spheres of possible magnitudes representing, respectively, the values of each have a non-trivial intersection*

Of two consonant, inaccurate tensor-fields, the inaccuracies of the first dominate *those of the second if, at every point of their support, the radius of the 4-sphere of possible values of the first, as determined by the* Σ*-norm (or* k*-norm, as applicable), is greater than that of the second, and their intersection includes the magnitude of (*i.e., *the center of the 4-sphere of) the first.*

The relation of dominance gets us closer to what we want, but does not by itself suffice, as it provides no quantitative measure of the differences between the inaccurate values. We can, however, use it to state a plausible necessary condition for a predicted value to deviate admissibly from the observed value: the deviation is admissible only if the field of experimentally observed values dominates the field of predicted values.

There is, as always, a wide selection to choose from in imposing a quantititative measure on the relation of dominance, which can then be used to formulate sufficient conditions on the two fields for the deviance of the values of the one from those of the other to be admissible. We choose the following for all the usual reasons. Fix two inaccurate, colored tensorial fields, κ and λ, such that κ dominates λ. At every scrap S of their shared support $\mathcal{S}$, let $\delta_s(\kappa(S), \lambda(S))$ be the absolute magnitude of the difference of the diameters of their respective 4-spheres of possible values.

Definition 4.6.1.3 *The* dominance *of* κ *over* λ, $\Delta_\kappa(\lambda)$*, is the volume-weighted average of the integral of* δ_s *over* $\mathcal{S}$*:*

$$\Delta_\kappa(\lambda) \equiv \frac{\int_{\mathcal{S}} \delta_s(\kappa(S),\, \lambda(S))\, d\mu(S)}{\int_{\mathcal{S}} d\mu(S)}$$

dμ is the measure on decoupages we defined in Curiel (2011) and referred to in §4.4.8 above. We can now characterize the deviance of predicted from observed values in terms of dominance. Given the predicted inaccuracy in the value of a field after dynamical evolution of a certain sort, and given the range of possible inaccuracy accruing to the determination of that value dictated by the kinematical regime under the experimental circumstances after the evolution, we say these two have a *deviance* equal to the dominance of one over the other, if one of them does in fact dominate the other. Clearly, then, the admissible deviances of a theory can be represented, at least in a purely formal fashion, by a family of dominances. One now can determine whether or not a given prediction accords with a given experimental result by comparing the actual deviances of the predicted from the observed values with the admissible deviances.

4.6.2 Consistent Maxwell-Boltzmann Theories

Before we discuss what is involved in laying down criteria for the admissibility of deviances, let's assume for the moment that we have a set of such criteria in hand, expressed as a family of dominances. There now arises the issue of the self-consistency, in a certain sense, of the physical theory. Let's say, for example, that an experimentalist is modeling a proposed experiment using a Maxwell-Boltzmann theory, including the family of dominances. She constructs and solves a physically well set initial-value formulation modeling the experiment, including the calculation of the inaccuracies accruing to the determination of the magnitudes at the end of the experiment, according to her solution. She then finds, to her surprise and discomfiture, that these calculated inaccuracies are all greater than the admissible deviances of predicted from observed values for that type of system under those experimental circumstances. It would seem, in such a case, that one could not determine whether an observation that seemingly conformed to the theory actually did so; the inaccuracy in the determination of the value is so great that the experimental agreement may be purely artifactual and not a true indicator of the soundness of the theory.

[*** clarify this muddle ***]

In more traditional terms, one might describe the problem as follows. Any two sets of exact initial data falling within the interval of possible inaccuracy for that system at the moment the experiment commences [98] have equal claim to represent the idealized, exact state of the system. Given two sets of initial data differing only slightly from each other, moreover, less than the possible inaccuracy, the two respective solutions to the set of exact partial-differential equations may, in general, evolve to be further and further apart, in a variety of technical senses, as time passes. Say, then, for a fixed observatory, we begin with two $\mathfrak{K}$-appropriate sets of exact initial data for the idealized partial-differential equations of a theory, each set representing the initial state of the same spatiotemporal region of the observatory, the two differing from each other by no more than the theory's interval of possible inaccuracy. Say, moreover, that solving the equations for the evolution of the system for each of the two sets yields respective exact solutions that, after a finite period of time, differ from each other by more than the admissible range of deviance of observed values of the quantities from predicted ones. Under these circumstances, it would not seem to make sense to ask whether the outcome of an experiment beginning with exact initial data within the possible range of inaccuracy encompassing those two sets conformed to or controverted the predictions of the theory, in so far as one's choices among differing but equally acceptable sets of exact initial data yield respective exact results that cannot possibly all be consonant with the predictions of the theory. [99]

In the terms of the machinery developed in this paper, we would say that, for the models of this experiment provided by the theory, the solutions to the physically well set initial-value formulations derived from two different idealizers will eventually be so different that at most one of them could be correct. This possibility arises from the fact that the solution to a well set initial-value formulation of an inaccurate partial-differential equation is itself an inaccurate field. It would seem that, if we are to have the capacity to judge whether or not the predictions of a theory soundly model physical phenomena, then we require that the inaccuracies of solutions to the theory's physically well set initial-value formulations dominate the ranges of possible inaccuracy accruing to the determinations of the values of those solution as dictated by the kinematical regime, and does so, moreover, by at least the given dominance. In the absence of any further requirements, nothing guarantees this.

[*** consistency can only be had for finite temporal intervals, since any solution will achieve arbitrarily large inaccuracies eventually, after a long enough period of evolution—or is that true? are there solutions that yield asymptotically bounded inaccuracies? I bet there are ***]

This possibility raises two questions. Given a Maxwell-Boltzmann theory and a set of criteria for determining the admissibility of observed deviances, are there any physically well set initial-value formulations that yield solutions of a kind as to be meaningfully compared with observational data? And is there any way to guarantee that a theory will possess such physically well set initial-value formulations, for instance by imposing requirements on the form of its comprised partial-differential equations? The latter question is beyond the scope of this paper. We can, however, give a simple way to characterize theories in answer to the former question. Let us denote the set of admissible

98. Where, recall, "moment" really means "a likely quite brief, yet finite, temporal interval".

99. I will not have room to treat another fascinating topic along these lines. So far, I have treated (albeit implicitly) "noise" as a source only of inaccuracy in the outcomes of experiments. Noise, however, can play other, indeed beneficial roles. In modeling and measuring the firing of action potentials in large populations of neurons, for instance, one *depends* on the presence of a small amount of background noise in order to potentiate the firings. More precisely, in the presence of subthreshold inputs, a small amount of noise (which doesn't have to be white—it can be colored, *i.e.*, temporally correlated, such as white noise passed though a first-order linear filter) causes threshold crossings and can reveal information about the structure of the otherwise subthreshold signal. Because the firing of a neuron is a time-delayed event occurring only after a discrete threshold in the potential is achieved, rather than a continuous and continuously responsive process, therefore, one can show that, under certain circumstances, the presence of noise markedly increases the probability of a firing-event during any given temporal interval. In so far as we believe these conditions to obtain *in vivo*, it is hypothesized that the character of the ambient noise in the brain plays an integral role in determining the firing patterns a neuronal population will exhibit under fixed stimulus. This is an example of a more general phenomenon known as 'dithering', familiar to the engineering community since the 1950s. See, for example, Knight (1972), Gammaitoni (1995), Wiesenfeld and Moss (1995), and Hunter, Milton, Thomas, and Cowan (1998).

I thank John Hunter for bringing to my attention these sorts of peculiar roles noise can play in experiements, as well as for directing me to the citations.

deviances we assume the theory to possess by '$\mathcal{D}$'. A physically well set initial-value formulation $\mathfrak{i}$ of a Maxwell-Boltzmann theory *respects* $\mathcal{D}$ if the inaccuracies of the fields derived as its solution dominate the kinematical inaccuracies accruing to the determination of the values of those fields, with a dominance at least as great as that given by $\mathcal{D}$. This gives us the required condition for consistency.

Definition 4.6.2.1 *A Maxwell-Boltzmann theory is* consistent *with a family of dominances* $\mathcal{D}$ *if every physically well set initial-value formulation in the theory respects* $\mathcal{D}$.

This suggests that, in order to be consistent, a Maxwell-Boltzmann theory must have its number of canonical volumes increased by two: a fourteenth volume consisting of an enumeration of ordered pairs, each consisting of a family of formal dominances and an octuplet from the twelfth volume, and a supplement to the thirteenth volume to render the semantical interpretation of these dominances as admissible deviances. We now turn, in the next section, to discussion of the sorts of content we may want in this semantical supplement.

4.6.3 The Dynamical Soundness of a Physical Theory

To begin to get a grip on these issues, consider again, for a given theory, the system of equations the theory comprises. Recall from the discussion of Navier-Stokes fluids in §4.2.3 that such a system may fail in so far as it is applied to regions in which it does not provide adequate predictions—in other words, the initial-value formulation of that system breaks down when formulated over such regions—even though the quantities treated by the theory are well defined over such regions. Now, what counts as an adequate prediction by a theory will vary from application to application in the following sense. Say we are to measure the temperature of a given type of physical system under two different sets of initial conditions, using a type of thermometer different in the one case from the other, and we then compare the results of the measurements with the predictions of our theory. Whereas we may find it admissible, say, 5% of the time, for the actual, measured value of the temperature to deviate 3% from the predicted value in the first case, we may find such a deviance to be wholly inadmissible with any non-negligible frequency in the second. In order to make these sorts of judgements, we require a set of methods for calculating, for particular types of systems under particular conditions, admissible deviances, from the values predicted by the theory, of the values of the quantities measured using particular types of probes, along with a statement of probability indicating one's level of confidence that the actual error will lie within the range of deviance settled on.

The measurement of temperature once again provides an excellent, concrete example of this phenomenon. Consider a thermometer immersed in an environment the temperature of which one has reason to believe is increasing at a constant rate. Because the equilibration of the thermometer with its environment always takes a finite amount of time, at any given instant the thermometer's reading will be a measure not of the environment's temperature at that very instant, but rather at an instant in the more or less immediate past. One naturally wonders about the time of response of the thermometer to the change in temperature—how many seconds behind the actual environmental temperature is the thermometer's reading? In order to treat this question at a somewhat elementary level, let us make the following assumptions. All the heat transferred to the thermometer is to be by convection, and all this heat is then retained in the thermometer. Thus, the rate of the transfer of heat through the convective layer immediately surrounding the thermometer exactly equals the total rate at which the thermometer absorbs heat. The equation of the evolution of the thermometer's temperature can then be expressed by combining Newton's law of cooling with Black's equation of heat capacity:

$$\frac{\mathrm{d}\theta}{\mathrm{d}t} = \frac{\chi A}{wc}(\theta_e - \theta)$$

where θ is the thermometer's temperature at time t, θ_e is the environment's temperature at time t, χ is the coefficient of convective heat-transfer between the environment and the thermometer, A is the surface area of the thermometer through which heat is transferred, w is the weight of the thermometer and c is its specific heat capacity.[100] $\tau \equiv \left(\frac{\chi A}{wc}\right)^{-1}$ has the dimensions of time, and is known as the *characteristic time constant* of the system. Assuming the simple initial relationship $\theta = \theta_e - Rt$, where R is the rate of temperature increase, then, after some elementary manipulation and integration, one deduces that, in the limit $t \gg \tau$, the relationship settles down to $\theta = \theta_e - R\tau$.[101] In other words, if the thermometer has been immersed in the environment for a long enough period of time, then the characteristic time constant is the length of time between the environment's being at a certain temperature and the thermometer's indication of that temperature.

In order to determine the characteristic time constant of the system, therefore, it suffices to immerse the thermometer, initially at a fixed uniform temperature, into an environment the rate of change of the temperature of which is constant and known. Two obvious problems now arise: in order to determine the rate of change of the environment's temperature, one must have a thermometer whose characteristic time constant for that environment is either known already or, at least, is known to be negligibly small; and to do this, the coefficient of convective heat-transfer must be known. This latter presents a particularly difficult challenge, for it follows from Nusselt's equation of heat transfer by forced convection that determination of this coefficient within reasonable bounds of uncertainty depends on somewhat detailed knowledge of the mass velocity of the environment relative to the thermometric surface, which is in general a complicated, 3-dimensional flow. Other factors, which must be taken into consideration when going beyond our elementary assumptions, are known to influence the characteristic time constant as well, including, *inter alia*, the Mach number of the environment, the size of the temperature change being considered, the rate of axial conduction of heat from the environment to the thermometer, the intensity and quality of the ambient radiation, and the turbulence in the environment. For certain kinds of systems, for example, a fourfold increase in the total temperature change can lead to a 25% variation in τ, and a 1.5% change in the intensity of turbulence can, by dint of influencing the coefficient of convective heat-transfer, change τ by up to 25% as well.[102] As this example illustrates, the specification of admissible deviances must be made with regard, *e.g.*, to temperature measurements of particular types of systems under certain kinds of conditions, not generically for all thermodynamical temperature measurements *simpliciter*.

These sorts of consideration suggest the following way of making these ideas precise.

Definition 4.6.3.1 *A* regime of dynamical consistency *(or* dynamic regime*, for short) for a Maxwell-Boltzmann theory is an ordered pair* $(\mathfrak{d}, \mathcal{D})$ *such that*

1. $\mathfrak{d}$ *is a set of algebraic and differential constraints on the physical quantities of the systems* cum *environments treated by the theory*
2. $\mathcal{D}$ *is a family of dominances defined in terms of those quantities consistent with the theory*

This allows for, finally,

Definition 4.6.3.2 *A Maxwell-Boltzmann theory is* dynamically sound *if it has a regime of dynamic consistency that accords with experiment.*

100. Of course, all these latter quantities, when measured to a sufficient accuracy, are themselves functions of the temperature, among other relevant quantities, but we will bracket this point for a moment, assuming that the coefficient of convective heat-transfer, the area, *etc.*, are constant.
101. See, *e.g.*, Benedict (1969, §11.2, pp. 144–5).
102. *Ad loc.*

4.6.4 Theoretical Under-Determination

[*** The fact that any given *measurement* is compatible with any of an infinite number of different theoretical propositions (*i.e.*, the ascription of an exact real number, or field of real numbers, to a point or region of spacetime) is well known (see, *e.g.*, Duhem, *The Aim and Structure of Physical Theory*, on the difference between what he calls practical facts and theoretical facts); the theorem I offer is more far-reaching: any given set of measurements, no matter the cardinality one allows for the set, is compatible with an infinite number of different dynamical structures, in a certain sense all continuous with each other—that is to say, with an infinite number of different mathematical theories of the same phenomena. ***]

Let us say that we have a dynamically sound theory in hand. We have used the machinery of canvases and decoupages, and of inaccurate fields and colorings, to represent the inaccuracy inevitably inhering in the modeling of experiments and the data they generate, as well as in defining the admissibility of deviances of predicted from observed values; nevertheless, it is still the case that, when we want to make contact with physics as practiced today, we must idealize. This raises the question: what idealizer will we choose? Even though the idealization of an inaccurate structure picks out a unique exact structure, there will in general, in virtue of the stability of idealizers, be many that yield exact structures so close to each other, in a certain technical sense, as to be indistinguishable with respect to the regime of the theory in play. In this section, we will make these considerations precise and draw out a few of their implications.

Consider again the proposed hyperbolic theories of relativistic, dissipative fluids discussed in §4.2, assuming for the sake of argument that they possess the structure of sound Maxwell-Boltzmann theories. Because the fineness of the observation and measurement of terms in the hyperbolic systems is circumscribed by the regime's possible inaccuracy, one will not be able to distinguish in a finite temporal interval any two solutions of the system differing from each other in an appropriate sense by no more than allowed by this inaccuracy during that interval. More to the point, one will not be able to distinguish a solution to one hyperbolic system from that of another, comprised by otherwise identical theories, so long as, again, those two solutions differ from each other by no more than the possible inaccuracy allowed by their shared kinematical regime. This fact naturally suggests the question, whether, given a hyperbolic system and a kinematical regime, there exists another hyperbolic system such that the set of solutions to the first system corresponding to any set of admissible initial-data continued for a finite temporal interval differs by no more than the possible inaccuracy allowed by that kinematical regime, for the same set of initial-data during the same temporal interval.

The perhaps surprising answer is that one can give an almost trivial proof to a mathematically precise statement of the question; the proof depends, however, on the hyperbolicity of the partial-differential equations at issue. As a consequence, given any sound two Maxwell-Boltzmann theories agreeing in their regimes and differing only with respect to the systems of partial-differential equations they comprise, both of which satisfy the conditions of the theorem, one will have no grounds for concluding on purely observational grounds that one of the theories is to be preferred over the other. In any event, one should again *not* take this as an argument for any sort of anti-realism or instrumentalism, *à la* Cartwright (1999).

Before stating the primary result of this section, the mentioned theorem, we need to lay down a few more definitions. First, the *supremal spacelike diameter* $\sigma_{\sup}[O]$ of an open subset of spacetime of compact closure O is defined by

$$\sigma_{\sup}[O] \equiv \sup\left\{ \int (|\gamma^m \gamma^n g_{mn}|)^{1/2} \mathrm{d}s : \gamma \in \mathfrak{s}_O \quad \& \quad \gamma^a = \frac{\mathrm{d}\gamma}{\mathrm{d}s} \right\}$$

A second one depends on the fact that, as shown in Curiel (2011, §5.5.2), given an inaccurate initial-value formulation, there is a more or less natural way to single out an ordinary spacelike hypersurface contained wholly within the spacelike, convex hull of an infimal decoupage on which the inaccurate initial data is fixed. Consider the family of all ordinary timelike geodesics orthogonal to this surface that intersect the future boundary of the ordinary domain of dependence of that hypersurface. The *maximal time of Cauchy development* for this inaccurate initial-value formulation, then, is the supremum of the intervals of proper time from the hypersurface to the boundary along the geodesics in the family.

There follows, as a direct consequence of theorems 4.5.3.6, 4.5.4.1, and 4.5.5.2,

Theorem 4.6.4.1 *Given any sound Maxwell-Boltzmann theory and any $\sigma, \tau > 0$, there exists a second sound Maxwell-Boltzmann theory distinct from the first having the same exact theory with regime as its idealization, if one restricts attention to physically well set initial-value formulations such that the supremal diameter of the domain of initial data is less than σ, and its maximal time of Cauchy development is less than τ.*

In effect, one has the freedom to change the system of partial-differential equations comprised by any Maxwell-Boltzmann theory, so long as one makes corresponding adjustments in one's choice of a family of idealizers, and produce a second, mathematically distinct theory that is observationally indistinguishable from the first: no possible set of observations could favor one over the other.

This theorem differs in significant ways from traditional results on the under-determination of theory by data. For starters, this result bears solely on the mathematical structure of the theory—its syntax, as it were—not on its interpretation, as in, say, disputes over the observational indistinguishability of that old chestnut, the Copernican and the Ptolemaic systems. It mandates, nevertheless, a far more pragmatic, almost inductive form of indistinguishability than the traditional one, in so far as it declares an indistinguishability tied to our actual practices of investigating nature by the vehicle of physical science—you tell me how long you want indistinguishability for, and I'll arrange it. Even when the issue is not one of strict observational indistinguishability, all such results I know of rely on the quantity of data being finite or even countable, for any rigorous conclusions to be drawn. No such restriction is placed here. No other result I know of, moreover, allows one to quantify the allowable deviances of different theories from each other depending on the spatiotemporal extent one wants to allow for observation and measurement. Finally, and perhaps most importantly, the whole tenor of this result differs from traditional ones, in so far as it allows one to have already in hand a theory whose predictive power is as accurate as allowed for by the nature of the quantities at issue and the character of the techniques available for their study—quite literally, as accurate as possible—and still declares that this does not suffice for the unique fixing of a mathematical stucture for a theory, *not even up to isomorphism of any sort.* In traditional arguments about the observational indistinguishability of two theoretical structures, such as with the Copernican and the Ptolemaic systems, one can always show that the essential kernel of the respective mathematical structures are isomorphic in the relevant way. In fact, I would argue that any proof of traditional indistinguishability requires such an isomorphism. This conclusion, if you like, underscores the necessarily approximative character of the enterprise of physical science.

4.7 The Theory Is and Is Not the Equations

Hertz famously said, "Maxwell's theory is his equations." Well, yes and no.

In one sense, the sense I think he meant, Hertz was assuredly correct. Maxwell's theory at the time stirred up a storm of controversy in large part because physicists of the day did not know how to think about the theory in the terms with which they were accustomed. The theory manifestly modeled a type of system of an oscillatory nature, and yet the theory said not a word about what was "doing" the oscillating. Where there are waves, the thinking went, there is a physical medium waving (the "luminiferous æther"); any putative theory that predicts waves and yet does not identify the medium, does not tell us how to investigate the properties of the medium, does not tell us how to envision the medium, is no physical theory. The same sort of thinking lay behind much of the contemporaneous controversy surrounding Newton's publication of his theory of gravity, and in particular his bold claim, "Hypotheses non fingo." I think he was correct in that bold claim, in the same sense that Hertz was correct in his assessment of Maxwell's theory. The theory in both cases, in the sense I am currently discussing, just *is* a more or less formal structure representing the patterns of behavior we have managed to extract from and impose on masses of observational data. The sterling virtue of this form of representation, moreover, is the capacity it lends us both to make predictions about the future behavior of physical systems of a certain sort, and to reason in a fairly precise way about quite general features of those sorts of physical systems. We do not require of this structure, in order to use it for these purposes, that it embody the tenets of any system of beliefs about the nature of the world that goes beyond the empirical, beyond, that is to say, what can be observed experimentally.

As I say, in that sense I think Hertz (and Newton) was correct. There is still a sense, however, in which I think the statement is not correct, though, I must stress, the sense is different enough that it has no bearing on the correctness of Hertz's intent, as I see it, in stating it.[103] To explicate that sense, consider the recent controversy, misguided in my opinion, about the priority of the "discovery" of general relativity: Hilbert or Einstein? If one thinks the theory of general relativity "is the Einstein field-equation", then I suppose there is a strong case for Hilbert. In a certain sense, I agree with this sentiment: knowing the Einstein field-equation, one, in a *recherché* sense to be sure, but still in a definite sense, knows how to model and comprehend all the phenomena putatively treated by the theory. What more could one ask of a theory?

As I have strenuously attempted to demonstrate in this paper, however, in a substantive and profound sense, the equation is most assuredly *not* the theory: one also needs all the collateral knowledge, both theoretical and practical, *not* contained in the equation, in order to apply the equation to the modeling and comprehension of all the phenomena putatively treated by the theory. To put the matter more vividly: the equation *as a result of* a (profound) investigation of the physical phenomena at issue, of all the empirical data and attempts to model that data heretofore, a teasing apart and characterization of the maximally common structure underlying the system of relations that obtains among them (supposing there is such a thing)—the sort of investigation that Newton and Maxwell and Einstein did and Hilbert did not accomplish—the equation as a result of that is the theory.[104]

There is no theory without experimental data to comprehend, in all its extravagant and fertile inaccuracy, without the capacity to get the laboratory into the theory. It was Einstein the physicist, not Hilbert the mathematician, who calculated the precession of Mercury's perihelion using a system of intuitively well founded approximations (Schwarzschild had not yet discovered his solution), who modeled a real physical system in the terms of his proposed theory and demonstrated that the theory had the (or: a) structure proper for the modeling of the phenomena. This is not the same thing as demonstrating that the theory is accurate in its predictions about the phenomen; I think it is all too common, however, to conflate these two, by focusing on prediction as the be-all, end-all of physics. The two—modeling and prediction—do not come to the same thing, as the distinction between the

103. I suspect—or at least hope—based on my reading of Hertz that he would have agreed with what I am saying here.

104. Compare the remarks of Born (1943, p. 9), very much in harmony with the spirit of this discussion: "...none of the notions used by the mathematicians, such as potential, vector potential, field vectors, Lorentz transformations ...are evident or given *a priori*. Even if an extremely gifted mathematician had constructed them to describe the properties of a possible world, neither he nor anybody else would have had the slightest idea how to apply them to the real world. The problem of physics is how the actual phenomena, as observed with the help of our sense organs aided by instruments, can be reduced to simple notions which are suited for precise measurement and used for the formulation of quantitative laws."

kinematic regime and the domain of soundness shows. [105]

Not only are theory and experiment consonant with each other, they are mutually inextricable—not, however, as equals. Theory plays Boswell to the subtle and tragic clown of experiment's Johnson.

105. My version of the mechanist philosophy? No explanations except relational ones?

References

Anile, A., D. Pavón, and V. Romano (1998). The case for hyperbolic theories of dissipation in relativistic fluids. arXiv:gr-qc/9810014v1).

Anscombe, G. E. M. (1971). Causality and determination. In *Metaphysics and the Philosophy of Mind*, Volume 2 of *The Collected Philosophical Papers of G. E. M. Anscombe*, pp. 133–147. Minneapolis: University of Minnesota Press, 1981. Originally delivered as Anscombe's inaugural lecture for her professorship at Cambridge University in 1971.

Arnowitt, R., S. Deser, and C. Misner (1962). The dynamics of general relativity. See Witten (1962).

Ashtekar, A. (1991). *Lectures on Non-Perturbative Canonical Gravity*, Volume 6 of *Advanced Series in Astrophysics and Cosmology*. Singapore: World Scientific.

Belot, G. (1996, September). Why general relativity *does* need an interpretation. *Philosophy of Science 63, Supplement: Proceedings of the 1996 Biennial Meetings of the Philosophy of Science Association. Part* I*: Contributed Papers*, S80–S88. Stable URL: `http://www.jstor.org/stable/188514`.

Benedict, R. (1969). *Fundamentals of Temperature, Pressure and Flow Measurements*. New York: John Wiley & Sons, Inc.

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

Bondi, H., M. van der Burg, and A. Metzner (1962). Gravitational waves in general relativity: VII. Waves from axi-symmetric isolated systems. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences 269*, 21–52.

Born, M. (1943). *Experiment And Theory in Physics*. New York: Dover Publications, Inc. The work is a slightly expanded form of a lecture given to the Durham Philosophical Society and the Pure Science Society, King's College, Newcastle-upon-Tyne, May 21, 1943. This edition is a 1956 reprint of the original Cambridge University Press printing of 1943.

Bosshard, B. (1976). On the b-boundary of the closed Friedmann model. *Communications in Mathematical Physics 46*, 263–268.

Bosshard, B. (1979). On b-boundaries of special space-time models. *General Relativity and Gravitation 10*, 963–966.

Boyd, R. (1991). On the current status of scientific realism. In *The Philosophy of Science*, pp. 195–222. Cambridge, MA: The MIT Press. Originally published in *Erkenntnis* 19(1983):45–90.

Bunge, M. (1979). *Causality and Modern Science* (third revised ed.). New York: Dover Publications, Inc.

Burgess, G. (1928, July–December). The international temperature scale. Technical Report RP 22(635), The Journal of Research of the National Bureau of Standards.

Callendar, H. (1887). On the practical measurement of temperature. *Proceedings of the Royal Society of London 178*, 160.

Cartwright, N. (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.

Clarke, C. (1973). Local extensions in singular space-times. *Communications in Mathematical Physics 32*, 205–214.

Clarke, C. (1975). Singularities in globally hyperbolic space-times. *Communications in Mathematical Physics 41*, 65–78.

Clarke, C. (1993). *The Analysis of Space-Time Singularities*. Number 1 in Cambridge Lecture Notes in Physics. Cambridge: Cambridge University Press.

Colombeau, J. (1992). *Multiplication of Distributions*. Number 1532 in Lecture Notes in Mathematics. Berlin: Springer.

Curiel, E. (1996). General relativity has no proper Hamiltonian formulation. Unpublished manuscript.

Curiel, E. (1999). The analysis of singular spacetimes. *Philosophy of Science 66*(S1), 119–145. A more recent version, with corrections and emendations, is available at http://strangebeautiful.com/phil-phys.html.

Curiel, E. (2000a). The constraints general relativity places on physicalist accounts of causality. *Theoria 15*(1), 33–58.

Curiel, E. (2000b). Is energy a well-defined quantity in the theory of general relativity? Unpublished manuscript.

Curiel, E. (2009). On tensorial concomitants and the non-existence of a gravitational stress-energy tensor. arXiv:0908.3322v3 [gr-qc].

Curiel, E. (2011). A formal model of the regime of a physical theory, with applications to problems in the initial-value formulation of the partial-differential equations of mathematical physics. Unpublished.

Curiel, E. (2014). The geometry of the Euler-Lagrange equation. Unpublished manuscript, most recent version available at http://strangebeautiful.com/phil-phys.html.

Ducasse, C. (1926). On the nature and the observability of the causal relation. In E. Sosa and M. Tooley (Eds.), *Causation*, Oxford Series in Philosophy, pp. 125–36. Oxford: Oxford University Press. Originally published in *Journal of Philosophy* 23(1926):57–68.

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

Eddington, A. (1923). *Mathematical Theory of Relativity* (Second ed.). Cambridge: Cambridge University Press.

Ehlers, J. and W. Kundt (1962). Exact solutions of the gravitational field equations. See Witten (1962), pp. 49–101.

Einstein, A. (1916). The foundation of the general theory of relativity. In *The Principle of Relativity*, pp. 109–164. New York: Dover Press, 1952. Published originally as “Die Grundlage der allgemeinen Relativitätstheorie”, *Annalen der Physik* 49(1916).

Einstein, A. (1984). *The Meaning of Relativity* (fifth ed.). Princeton: Princeton University Press. Trans. E. Adams, E. Straus and S. Bargmann. First edition published in 1922.

Ellis, G. and B. Schmidt (1977). Singular space-times. *General Relativity and Gravitation 8*(11), 915–953. doi:10.1007/BF00759240.

Eu, B. (2002). *Generalized Thermodynamics: The Thermodynamics of Irreversible Processes and Generalized Hydrodynamics*. Dordrecht: Kluwer Academic Publishers.

Fine, A. (1982). Some local models for correlation experiments. *Synthese 50*, 279–94.

Gammaitoni, L. (1995). Stochastic resonance and the dithering effect in threshold physical systems. *PRE 52*(5), 4691–4698.

Geroch, R. (1966). Singularities in closed universes. *Physical Review Letters 17*, 445–447. doi:10.1103/PhysRevLett.17.445.

Geroch, R. (1968a). Local characterization of singularities in general relativity. *Journal of Mathematical Physics 9*, 450–465.

Geroch, R. (1968b). What is a singularity in general relativity? *Annals of Physics 48*, 526–540.

Geroch, R. (1969). Spinor structure of space-times in general relativity I. *Journal of Mathematical Physics 9*, 1739–1744.

Geroch, R. (1970a). Domain of dependence. *Journal of Mathematical Physics 11*(2), 437–449.

Geroch, R. (1970b). Spinor structure of space-times in general relativity II. *Journal of Mathematical Physics 11*, 343–8.

Geroch, R. (1973, December). Energy extraction. *Annals of the New York Academy of Sciences 224*, 108–117. Proceedings of the Sixth Texas Symposium on Relativistic Astrophysics. doi:10.1111/j.1749-6632.1973.tb41445.x.

Geroch, R. (1981). *General Relativity from A to B*. Chicago: University of Chicago Press.

Geroch, R. (1985). *Mathematical Physics*. Chicago: University of Chicago Press.

Geroch, R. (1996). Partial differential equations of physics. In G. Hall and J. Pulham (Eds.), *General Relativity*, Aberdeen, Scotland, pp. 19–60. Scottish Universities Summer School in Physics. Proceedings of the 46th Scottish Universities Summer School in Physics, Aberdeen, July 1995. Preprint available at arXiv:gr-qc/9602055.

Geroch, R. (2001). On hyperbolic "theories" of relativistic dissipative fluids. arXiv:gr-qc/0103112v1).

Geroch, R., L. Can-bin, and R. Wald (1982). Singular boundaries of space-times. *Journal of Mathematical Physics 23*, 432–435.

Geroch, R., E. Kronheimer, and R. Penrose (1972). Ideal points in space-time. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences 327*, 545–567.

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

Hadamard, J. (1923). *Lectures on Cauchy's Problem in Linear Partial Differential Equations*. New York: Dover Publications. A 1952 reprint of the original 1923 publication by Yale University Press.

Hall, J. (1955). The international temperature scale. In *Temperature*, Volume 2, pp. 116. New York: Reinhold.

Halmos, P. (1950). *Measure Theory*. New York: Van Nostrand and Co.

Hawking, S. (1965). Occurrence of singularities in open universes. *Physical Review Letters 15*, 689–690. doi:10.1103/PhysRevLett.15.689.

Hawking, S. (1966a). The occurrence of singularities in cosmology. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences 294*, 511–521. doi:10.1098/rspa.1966.0221.

Hawking, S. (1966b). The occurrence of singularities in cosmology. II. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences 295*, 490–493. doi:10.1098/rspa.1966.0255.

Hawking, S. (1966c). Singularities and the geometry of space-time. Adams Prize Essay (unpublished).

Hawking, S. (1966d). Singularities in the universe. *Physical Review Letters 17*, 444–445. doi:10.1103/PhysRevLett.17.444.

Hawking, S. (1967). The occurrence of singularities in cosmology. III: causality and singularities. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences 300*, 187–210. doi:10.1098/rspa.1967.0164.

Hawking, S. and G. Ellis (1973). *The Large Scale Structure of Space-Time*. Cambridge: Cambridge University Press.

Hawking, S. and R. Penrose (1996). *The Nature of Space and Time*. Isaac Newton Institute Series of Lectures. Princeton: Princeton University Press.

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

Herrera, L. and J. Martínez (1997). Dissipative fluids out of hydrostatic equilibrium. arXiv:gr-qc/9710099v1).

Herrera, L. and J. Martínez (1998). Thermal conduction before relaxation in slowly rotating fluids. arXiv:gr-qc/9804035v1.

Herrera, L. and D. Pavón (2001a). Hyperbolic theories of dissipation: Why and when do we need them? arXiv:gr-qc/0111112v1.

Herrera, L. and D. Pavón (2001b). Why hyperbolic theories of dissipation cannot be ignored: Comment on a paper by Kostädt and Liu. *Physical Review D 64*(088503). Preprint available at arXiv:gr-qc/0102026v1.

Herrera, L., A. D. Prisco, J. Martín, J. Ospino, N. Santos, and O. Troconis (2004). Spherically symmetric dissipative anisotropic fluids: A general study. arXiv:gr-qc/0403006v1.

Herrera, L., A. D. Prisco, and J. Martínez (1998). Breakdown of the linear approximation in the perturbative analysis of heat conduction in relativistic systems. arXiv:gr-qc/9803081v1.

Hiscock, W. and L. Lindblom (1985). Generic instabilities in first-order dissipative relativistic fluid theories. *Physical Review D 31*, 725.

Hocking, J. and G. Young (1988). *Topology*. New York: Dover Publications, Inc. Originally published in 1961 by Addison-Wesley Publishing Co.

Hunter, J., J. Milton, P. Thomas, and J. Cowan (1998). Resonance effect for neural spike time reliability. *JNP 80*(3), 1427–38.

Johnson, R. (1977). The bundle boundary in some special cases. *Journal of Mathematical Physics 18*, 898–902.

Johnson, R. (1979). The bundle boundary for the Schwarzschild and Friedmann solutions. *General Relativity and Gravitation 10*, 967–968.

Joshi, P. (1993). *Global Aspects in Gravitation and Cosmology*. Number 87 in International Series of Monographs on Physics. Oxford: Oxford University Press.

Jou, D., J. Casas-Vázquez, and G. Lebon (2001). *Extended Irreversible Thermodynamics* (3rd ed.). Berlin: Springer-Verlag.

Knight, B. (1972). Dynamics of encoding in a population of neurons. *JGP 59*, 734–766.

Kobayashi, S. and K. Nomizu (1963). *Foundations of Differential Geometry*. Number 15 in Interscience Tracts in Pure and Applied Mathematics. New York: John Wiley & Sons. Volume 1.

Kolmogorov, A. and S. Fomin (1970). *Introductory Real Analysis*. New York: Dover Publishing, Inc. Trans. R. Silverman. A 1975 re-print of the edition originally published in 1970 by Prentice-Hall, Inc.

Konkowski, D. and T. Helliwell (1992). Singularities in colliding plane-wave spacetimes. In G. Kunstatter (Ed.), *General Relativity and Relativistic Astrophysics*, pp. 115–119. Singapore: World Scientific.

Kostädt, K. and M. Liu (2000). On the causality and stability of the relativistic diffusion equation. *Physical Review D 62*(023003). Preprint available at arXiv:cond-mat/0010276v1.

Kundt, W. (1963). Note on the completeness of spacetimes. *Zeitschrift für Physik 172*, 488–489.

Landau, L. and E. Lifschitz (1975). *Fluid Mechanics* (Second ed.). Oxford: Pergamon Press. An expanded, revised edition of the original 1959 edition. Translated from the Russian by J. Sykes and W. Reid.

Leray, J. (1934). Essais sur la movement d'un liquide vis quenx emplissant l'espace. *Acta Mathematica 63*, 193–248.

Lewis, D. (1986). Causal explanation. In *Philosophical Papers*, Volume 2, pp. 214–241. Oxford: Oxford University Press, 1986.

Mackie, J. (1980). *The Cement of the Universe* (paperback ed.). Oxford: Clarendon Press.

Malament, D. (1977). The class of continuous timelike curves determines the topology of spacetime. *Journal of Mathematical Physics 18*(7), 1399–1404. doi:10.1063/1.523436.

Maxwell, J. C. (1877). *Matter and Motion*. New York: Dover Publications, Inc. Originally published in 1877. This edition is a 1952 unaltered republication of the 1920 Larmor edition.

Mayer, J. (1842). On the forces of inorganic nature. In R. Lindsay (Ed.), *Energy: Historical Development of the Concept*, Volume 1, pp. 277–283. Stroudsberg, PA: Dowden, Hutchinson, Ross. Originally published as "Bemerkungen über die Kräfte der unbelebten Natur", *Annalen der Chemie und Pharmakie* 42(1842):233.

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

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.

Misner, C. (1963). The flatter regions of Newman, Unti, and Tamburino's generalized Schwarzschild space. *Journal of Mathematical Physics 4*, 924–937.

Misner, C., K. Thorne, and J. Wheeler (1973). *Gravitation*. San Francisco: Freeman Press.

Müller, I. and T. Ruggeri (1993a). *Extended Thermodynamics*. Berlin: Springer-Verlag.

Müller, I. and T. Ruggeri (1993b). *Rational Extended Thermodynamics* (2nd ed.). Berlin: Springer-Verlag. This is the second edition of *Extended Thermodynamics* by the same authors.

Norton, J. (1985). What was Einstein's principle of equivalence? *Studies in History and Philosophy of Science 16*, 203–246.

Norton, J. (1993). General covariance and the foundations of general relativity: Eight decades of dispute. *Reports on Progress in Physics 56*(1), 791–858. doi:10.1088/0034-4885/56/7/001.

Pauli, W. (1921). *The Theory of Relativity*. New York: Dover Publications, Inc. A 1981 reprint of the 1958 edition from Pergamon Press, a translation by G. Field of the original "Relativitätstheorie", in *Encyklopädie der matematischen Wissenschaften*, vol. V19, B. G. Teubner, Leipzig, 1921.

Peirce, C. S. (1898). *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.

Penrose, R. (1960). A spinor approach to general relativity. *Annals of Physics 10*, 171–201.

Penrose, R. (1965). Gravitational collapse and space-time singularities. *Physical Review Letters 14*(3), 57–59. doi:10.1103/PhysRevLett.14.57.

Penrose, R. (1968). Structure of spacetime. In C. DeWitt and J. Wheeler (Eds.), *Battelle Rencontres*. New York: W. A. Benjamin.

Putnam, H. (1975). Explanation and reference. In *Mind, Language and Reality: Philosophical Papers*, Volume 2, Chapter 11, pp. 196–214. Cambridge: Cambridge University Press. Originally published in *Conceptual Change*, eds. G. Pearce and P. Maynard (Dordrecht: Kluwer Academic Publishers, 1973), pp. 199–221.

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

Quinn, T. and R. Wald (1999). Energy conservation for point particles undergoing radiation reaction. arXiv:gr-qc/9903014v1.

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

Ruelle, D. (1981). Differential dynamical systems and the problem of turbulence. *Bulletin of the American Mathematical Society 5*, 29–42.

Russell, B. (1927). *The Analysis of Matter*. New York: Dover Press. A 1954 reprint of the original 1927 edition by Kegan Paul, Trench, Trübner & Co. of London.

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

Sachs, R. (1962). Gravitational waves in general relativity: VIII. The waves in asymptotically flat space-time. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences 270*, 103–126.

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

Schmidt, B. (1971). A new definition of singular points in general relativity. *General Relativity and Gravitation 1*, 269–280.

Scott, S. and P. Szekeres (1994). The abstract boundary—a new approach to singularities of manifolds. *Journal of Geometry and Physics 13*, 223–253.

Shepley, L. and G. Ryan (1978). *Homogeneous Cosmological Models*. Princeton: Princeton University Press.

Shimony, A. (1993). Reality, causality and closing the circle. In *Search for a Naturalistic World View: Scientific Method and Epistemology*, Volume 1, Chapter 2, pp. 21–61. Cambridge: Cambridge University Press.

Sommerfeld, A. (1964). *Partial Differential Equations in Physics*, Volume VI of *Lectures on Theoretical Physics*. New York: Academic Press. Trans. E. Straus.

Spivak, M. (1979a). *A Comprehensive Introduction to Differential Geometry* (Second ed.), Volume 1. Houston: Publish or Perish Press. First edition published in 1970.

Spivak, M. (1979b). *A Comprehensive Introduction to Differential Geometry* (Second ed.), Volume 2. Houston: Publish or Perish Press. First edition published in 1970.

Steenrod, N. (1951). *The Topology of Fibre Bundles*. Number 14 in Princeton Mathematical Series. Princeton, NJ: Princeton University Press.

Stein, H. (1972). On the conceptual structure of quantum mechanics. In R. Colodny (Ed.), *Paradigms and Paradoxes: The Philosophical Challenge of the Quantum Domain*, Volume 5 of *University of Pittsburgh Series in the Philosophy of Science*, pp. 367–438. Pittsburgh, PA: University of Pittsburgh Press.

Stein, H. (1984). The Everett interpretation of quantum mechanics: Many worlds or none? *Noũs 18*, 635–652.

Stein, H. (1989, June). Yes, but...: Some skeptical remarks on realism and anti-realism. *Dialectica 43*(1-2), 47–65. doi:10.1111/j.1746-8361.1989.tb00930.x.

Stein, H. (1994). Some reflections on the structure of our knowledge in physics. In D. Prawitz, B. Skyrms, and D. Westerståhl (Eds.), *Logic, Metholodogy and Philosophy of Science*, Proceedings of the Ninth International Congress of Logic, Methodology and Philosophy of Science, pp. 633–55. New York: Elsevier Science B.V. I do not have access to the published version of Stein's paper, but rather only to a typed manuscript. All references to page numbers, therefore, do not correspond to those of the published version. The typed manuscript I have is about 17 pages long, and the published version about 22. Multiplying the page numbers I give by $\frac{17}{22}$ and adding the result to 633 (the number of the first page in the published version) should give approximately the page number in the published version.

Stein, H. (2004). The enterprise of understanding and the enterprise of knowledge—for Isaac Levi's seventieth birthday. *Synthese 140*, 135–176. I do not have access to the published version of Stein's paper, but rather only to a typed manuscript. All references to page numbers, therefore, do not correspond to those of the published version. The typed manuscript I have is 65 pages long, and the published version about 41. Multiplying the page numbers I give by $\frac{41}{65}$ and adding the result to 135 (the number of the first page in the published version) should give approximately the page number in the published version.

Stein, H. (unpub.). How does physics bear upon metaphysics; and why did Plato hold that philosophy cannot be written down? Typed manuscript not submitted for publication. Available for download at `http://strangebeautiful.com/phil-phys.html`.

Stimson, H. (1949, March). The international temperature scale of 1948. Technical Report RP 1926(209), The Journal of Research of the National Bureau of Standards.

Stimson, H. (1961, September 8). The international temperature scale of 1948, text revision of 1960. Technical Report 37, National Bureau of Standards Monograph.

Sussmann, R. (1988). On spherically symmetric shear-free perfect fluid configurations (neutral and charged). II. Equation of state and singularities. *Journal of Mathematical Physics 29*, 945–970.

Synge, J. (1957). *The Relativistic Gas*. Amsterdam: North-Holland Publishing Co.

Synge, J. (1960). *Relativity: The General Theory*. Amsterdam: North-Holland Publishing Co.

Taub, A. (1979). Remarks on the symposium on singularities. *General Relativity and Gravitation 10*, 1009.

Temam, R. (1983). *Navier-Stokes Equations and Non-Linear Analysis*. Philadelphia: Society for Industrial and Applied Mathematics.

Thorpe, J. (1977). Curvature invariants and space-time singularities. *Journal of Mathematical Physics 18*, 960–964.

Tolman, R. (1934). *Relativity, Thermodynamics and Cosmology*. New York City: Dover Publications, Inc. A 1987 facsimile of the edition published by the Oxford University Press, at Oxford, 1934, as part of the International Series of Monographs on Physics.

Trautman, A. (1962). Conservation laws in general relativity. See Witten (1962), pp. 169–198.

Trautman, A. (1970a). Fibre bundles associated with space-time. *Reports on Mathematical Physics 1*, 29–62.

Trautman, A. (1970b). Invariance of Lagrangian systems. In L. O'Raifeartaigh (Ed.), *General Relativity: Papers in Honour of J. L. Synge*, Chapter 5, pp. 85–99. Oxford: Clarendon Press.

Trautman, A. (1980). Fiber bundles, gauge fields, and gravitation. In A. Held (Ed.), *General Relativity and Gravitation*, Volume 1, Chapter 9, pp. 287–308. New York: Plenum Press. 2 Volumes.

van Fraassen, B. (1989). The Charybdis of realism: Epistemological implications of Bell's inequality. In J. Cushing and E. McMullin (Eds.), *Philosophical Consequences of Quantum Theory: Reflections on Bell's Theorem*, pp. 97–113. Notre Dame, IN: University of Notre Dame Press. Originally published in *Synthese* 52(1982):25–38. Appendix added in 1989.

Vickers, J. and J. Wilson (1998). A nonlinear theory of tensor distributions. arXiv:gr-qc/9807068v1.

Wald, R. (1984). *General Relativity*. Chicago: University of Chicago Press.

Wald, R. (1999). Gravitation, thermodynamics and quantum theory. arXiv:gr-qc/9901033.

Wiesenfeld, K. and F. Moss (1995). Stochastic resonance and the benefits of noise: From ice ages to crayfish and SQUIDs. *Nature 373*, 33–36.

Witten, L. (Ed.) (1962). *Gravitation: An Introduction to Current Research*. New York: Wiley & Sons Press.

Wloka, J. (1987). *Partial Differential Equations*. Cambridge: Cambridge University Press. Originally published in German as *Partielle Differentialgleichungen*, Stuttgart: B. G. Teubner, 1982. Trans. C. B. and M. J. Thomas.

Yang, C. (1961). *Elementary Particles: A Short History of Some Discoveries in Atomic Physics*. Princeton, NJ: Princeton University Press.

Zimdahl, W., D. Pavón, and R. Maartens (1996). Reheating and causal thermodynamics. arXiv:astro-ph/9611147v1.