

Schematizing the Observer and the Epistemic Content of Theories[†]

Erik Curiel[‡]

I dedicate this paper, with affection and admiration, to Howard Stein, mentor and friend, one of the few people I have met in my life to deserve the honorific ‘philosopher’ in the deep Platonic sense.

So I suggest that the principal difficulty is not that of how to leave the theory outside the laboratory door, but that of how to get the laboratory inside the theory. ¶ Well, how *do* we do it? For of course we do put theory and experiment in relation to one another; otherwise it would be impossible to test theories, and impossible to apply them. It would also, I should add, be impossible to *understand* a theory, as anything but a purely mathematical structure—impossible, that is, to understand a theory *as* a theory of physics—if we had no systematic way to put the theory into connection with observation (or experience).

Howard Stein

“Some Reflections on the Structure of Our Knowledge in Physics”
(italics are Stein’s)

ABSTRACT

I argue that, contrary to the standard view, one cannot understand the structure and nature of our knowledge in physics without an analysis of the way that observers (and, more generally, measuring instruments and experimental arrangements) are modeled in theory. One upshot is that standard pictures of what a scientific theory can be are grossly inadequate. In particular, standard formulations assume, with no argument

[†]I thank Jeremy Butterfield for delightfully edifying discussion and for his usual detailed, extensive and penetrating comments on a draft. I thank Adam Caulton, Neil Dewar, Robert DiSalle, Michael Friedman, Bill Harper, Sebastian Lutz, Lydia Patton, Chris Pincock, and Jim Weatherall for many enjoyable and illuminating conversations covering all the topics I treat here. I am, as always, grateful to Howard Stein for many conversations over many years on these topics, and for having written papers that continue to inspire me. This paper owes a clear and large debt especially to his paper “Some Reflections on the Structure of Our Knowledge in Physics”.

[‡]**Author’s address:** Munich Center for Mathematical Philosophy, Ludwig-Maximilians-Universität; Black Hole Initiative, Harvard University; Smithsonian Astrophysical Observatory, Radio and Geoastronomy Division; **email:** erik@strangebeautiful.com

ever given, that it is possible to make a clean separation between, on the one hand, one part of the scientific knowledge a physical theory embodies, *viz.*, that encoded in the pure mathematical formalism and, on the other, the remainder of that knowledge. The remainder includes at a minimum what is encoded in the practice of modeling particular systems, of performing experiments, of bringing the results of theory and experiment into mutually fruitful contact—in sum, real application of the theory in actual scientific practice. This assumption comes out most clearly in the picture of semantics that naturally accompanies the standard view of theories: semantics is fixed by ontology’s shining City on the Hill, and all epistemology and methodology and other practical issues and considerations are segregated to the ghetto of the theory’s pragmatics. We should not assume such a clean separation is possible without an argument, and, indeed, I offer many arguments that such a separation is not feasible. An adequate semantics for theories cannot be founded on ontology, but rather on epistemology and methodology.

Contents

1	The Complexity of the World and the Simplicity of Theory	2
2	What Is a Theory?	6
3	Why We Need to Schematize the Observer	8
4	What a Theory Is	13
5	What Measurements and the Observer Are	19
6	The Different Forms and Aspects of Scientific Knowledge	22
7	How Theory and Experiment Make Contact with and Inform Each Other	29
8	The Breakdown Scale and the Regimes of a Theory	37
9	Semantics Is Epistemology, Not Ontology	44
10	Valediction	49
	Appendix: Summary	49
	References	50

1 The Complexity of the World and the Simplicity of Theory

For essentially every physical theory we have (Navier-Stokes theory, general relativity, quantum field theory, *etc.*), we have very little detailed knowledge of the structure of generic solutions.

Usually, we know exact solutions only under conditions of perfect or near-perfect symmetry or some other unrealistic assumption (two bodies, no external perturbative influences, . . .), and then argue that we can apply such solutions to real physical systems, because the approximation is adequate under appropriately controlled circumstances. That is to say: we have no real idea *at all*, in a representational sense, “what the world would be like if the theory were true or largely true of it”, for any physical theory.

This becomes particularly clear when one stops to think, to really think hard, about the mind-blowing complexity and richness of texture of real physical systems in the world. Take this glass of water resting on the table, seemingly in equilibrium with its environment. The water does not in fact have a constant, static density, pressure, temperature, shear-stress, heat flux, fluid velocity and volume, as we would naturally ascribe to it in a basic treatment using Navier-Stokes theory, the classical theory of viscoelastic, thermoconductive fluids. There are tiny temperature gradients from the edge of the glass to the interior, and likewise tiny pressure gradients. These drive microscopic flows and eddies and vortices, which generate fluctuations in the shear-stress. There is evaporation of the water at the surface, adsorption of a thin layer of particulate matter at the surface, and thereby absorption of particulates from the air and of air molecules themselves. There is thus exchange of thermal energy at the boundary, complicated by the little air bubbles and particulate matter suspended there, and inhomogeneities in the glass at the boundary. The water itself, at a finer level of detail, consists of a hyper-dense, stereometrically complex network of Hydrogen bonds with extraordinarily rich characteristic symmetry patterns as web to the west of instances of dozens of species of ionic combinations of H and O (H_2O being not even a majority, only a plurality). The water is pervaded and perturbed by the ambient electromagnetic field, itself composed of endless linear superpositions of radiation from the antennæ of radio stations playing Berg’s “Lulu” and the BBC news, mobile phones all over the Earth, magnetic induction from small movements of power cords, the Earth’s own magnetic field, infra-red and radio-waves from Sagittarius A* (the supermassive black hole at the center of the Milky Way), the cosmic microwave background radiation, the jiggling of the water’s own molecules, the rotation of the iron core of some planet 30 light-years away, and the gamma ray bursts from distant astronomical cataclysms. The water is inundated by cosmic rays. In every atom of every molecule, continual weak and strong nuclear reactions manifest themselves. And on and on and on.

I think we tend to forget in philosophy, we overlook, lose sight of, how complex real physical systems are and what a miracle it is that our almost naively, recklessly simple-minded theories, and the childishly sketched models we construct in those theories, can still capture them with astonishing accuracy, and do so in ways, moreover, that seem to give us real understanding of the nature of the world in a broader sense. And correlatively, we forget how much distance there is between those simple models we actually do know in any given theory and the real physical systems they purport to represent, and so we forget how many of the theory’s other models would almost certainly do the representational tasks required and *prima facie* do it better (by adding ever more finely grained detail, for example), if one of them does at all—how many of those models that we have no idea how to construct in any way graspable by the human mind, or indeed even how to identify if someone gave us one gift-wrapped.

The standard view today in much if not most of philosophy of science in general, and philosophy of physics in particular, however, especially since the seminal work of Suppes (1960; 1962),

presupposes that a theory is characterized in a strong sense by its family of “models”, where a model is always meant to be something that, at a minimum, captures the essence of a solution to the theory’s equations of motion or dynamical field equations. This is true whether one hews to the semantic view of theories (Suppe 1974; van Fraassen 1980) or the Best-Systems picture (Cohen and Callender 2009) or a semantics based on possible worlds (Lewis 1970b; Butterfield 2019), or one is a neo-Carnapian (Demopoulos 2013), or a structuralist (Stegmüller 1979; da Costa and French 2005), or a neo-Kantian (Friedman 2001), or an inferentialist (Suárez 2004), or one tries to reconcile the syntactic and the semantic views by the use of category theory (Halvorson and Tsementzis 2017), or one uses category theory directly to represent the models of the theory (Weatherall 2017), or one champions a sophisticated syntactic view—declared dead many times during and after a long period of mordant vilification—on its own (Lutz 2014), or one adopts any other of the contemporary popular accounts of scientific theory.¹

In light of my remarks about the paucity of our knowledge of actual solutions to our theories’ equations, therefore, to accept the standard view is to accept one of two options:

1. we stipulate by fiat that there is a family of actual (or possible, in some sense) physical systems, though we cannot identify most of them in practice, that are the ones represented by the solutions, and those are the models;
2. we stipulate by fiat that the family of solutions themselves (*e.g.*, vector fields on phase space), though we cannot represent them in anything like closed form or even identify almost any of them if they were to bite us on the ass, are the models.

Call this ‘the standard view’s dilemma’. In both cases, we are saying that the overwhelming majority of the content of the theory is something that we do not know, are in fact nowhere near knowing, and have good reason to think we will never know in anything remotely resembling thoroughness and detail. Nonetheless, the story goes, there is a preternatural relation of subsistence among the formalism, the models, and the physical systems that univocally and unambiguously fixes them; and that in turn gives rise to a magical relation of “reference” or “representation” that zaps out from the formalism to the world and latches on to the salient physical systems in a physically significant way. We, in our supernal cognitive puissance, just *know* this.

We are talking here of a Fichtean faculty of pure intellectual intuition that grasps without cognitive or perceptual mediation the Kantian *Ding an sich*. Or more precisely, we are talking of nothing at all. A semantic relation that is not known and cannot be known by humans in our current epistemic circumstance has no possible use in actual science and has no possible bearing on a fruitful and illuminating analysis of scientific knowledge.

My rhetoric may make it sound as though only the realists are subject to my wrath, but that is not so. I intend my wrath to be directed as well at all philosophers beholden in any way to

1. I find, moreover, that in the practice of philosophers of physics who do not work directly on the semantics and structure of theories, but rather work on problems that either directly or implicitly require an account of a semantics, without hesitation, indeed usually without comment or (I suspect) thought, they employ the standard semantic view or, its close kin, a semantics based on Lewisian possible worlds. It is in this sense as well that I intend the claim that the view I sketch is “standard” today. It may well be that very few today really believe the standard view as I sketch it—but then I cannot understand the arguments they do make, *e.g.*, in trying to articulate a criterion for theory equivalence based on formalism alone, with vague gestures at how it is all to make contact with empirical content as an afterthought.

something like the standard view, whether of a realist or an anti-realist bent. What is this fairyland of formal models that are the semantic tools the instrumentalist uses to make his predictions? All those who hew to something like the standard view, no matter how they think of the character of the semantic relations between a theory and the world in detail, have the same general problem: what is the family of formal structures that “characterizes” the theory, and how can we grasp them in a comprehensive way? That is to say, which formal structures of the theory are even such as to be appropriate to serve as adequate representations of parts of the world, and how can we know this, given our woeful lack of knowledge about the details of the formal structures? For not all models in the theory as standardly formulated can designate or represent in a meaningful way all physical systems we naively think they should be able to, though we have no way of determining this in general. *What is the theory?*

We need a mechanism or device for bringing a mathematically characterizable situation in the world of experience into substantive, physically significant contact with the mathematical structures of our theories. This mechanism would be an intermediary that identifies the junctions where meaningful connections can be made between the two and embodies the possibility of the epistemic warrant we think we construct for our theories from such contact and connection. That mechanism, I shall argue, is the schematic representation of the observer. To quote Stein (1994, pp. 649–650):

In actual fact, the experimental physics is treated separately as a discipline in its own right, that is partly an art: an affair of both knowledge and manipulative and perceptual skill. But the possibility of connecting this art with the theory is closely connected with a certain possibility within the mathematical structure that is the theoretical framework: using a word I have introduced earlier, the possibility of representing experiments, and of representing the observer, “schematically.” ... I want to speak ... of “schematizing the observer within the theory”; ... the intention is ... to secure empirical content—content within experience—for an abstract structure. ¶ ... [S]o far as the fundamental theory is concerned—or rather, so far as mathematically defined structures and rigorous arguments are demanded—the “schematic” representation of observers, experiments, and observations, is, I believe, as far as we know how to go.

The standard view gets right the fact that the family of physical systems a theory appropriately and adequately treats is a central, essential component of the epistemic content of a theory.² The standard view goes wrong in thinking that it makes sense to characterize this family in the abstract, without recourse to and reliance on all the scientific knowledge we have pertaining to the theory and its applications. To identify that family of systems, we need to demarcate the theory’s regime of applicability. To do that, we need to fix error tolerances, acceptable levels of precision, *etc.*, in our experimental practice. To do that, we need real models of actual experimental techniques, representing the actual tools we have at our disposal for performing such measurements, for the regime of applicability changes over time with advances in practical knowledge—we need to schematize the observer.³ (I give more detailed and extended arguments to defend this claim in

2. I characterize my technical notion of “appropriate and adequate” in §8 below—for the time being, I use it as a term of art. I use ‘propriety’ as the nominal form of ‘appropriate’ when I intend it in this (to be characterized) technical sense.

3. Although the recent wave of philosophical work on practice in science, exemplified by such thinkers as Margaret

§3, and throughout the paper after.)

If demarcating the regime of applicability did not depend on practical knowledge of the kind that changes—the current reach of experimental technique that determines fixed levels of precision and accuracy, the availability of approximative theoretical techniques and heuristic forms of argument to make calculations tractable, and so on—then there would be a fixed number, once and for all—say $3.14159265358\dots \times 10^{-3}$ cm—at scales smaller than which the representational capacities of Navier-Stokes theory breaks down. But that is simply false. It breaks down at different scales for different types of fluids under different conditions, and the increasing acuity and resolution of our measuring instruments, in conjunction with our increasingly powerful control of experimental error, both practically and experimentally, and the increasing sophistication of the models we can construct of our measuring devices and the fine details of their interactions with different Navier-Stokes fluids under different conditions, all mean that such numbers become ever smaller. They are not fixed *sub specie aeternitatis*.⁴ (I explain the ideas of a breakdown scale and the regime of applicability of a theory more fully in §8.)

Thus, neither is the family of physical systems that a theory appropriately and adequately treats fixed *sub specie aeternitatis*. I think this is right. A scientific theory is not a Platonic form. It does not exist independently of our actual epistemic state.

2 What Is a Theory?

What is the theory of general relativity? Here is one way to make that question somewhat more precise: what is the more or less secure fund of scientific knowledge that constitutes the theoretical and empirical content of the theory?⁵ This formulation of the question is certainly not a standard one in philosophy of science. On that standard view, the emphasis is almost invariably on constructing formal models corresponding to something like solutions to the field equations or equations of motion of a theory, and perhaps asking when two such families of models (or other

Morrison and Hasok Chang, has much in common with the views I propound here, I do not explicitly engage with the literature. It is too vast, and too often the motivations and aims of the work differ too much from those of this paper, for me to do so without expanding the scope of the paper beyond what it can profitably bear. For the same reason, I do not directly engage with the literature on complexity in science, such as the marvelous work on generative entrenchment, levels of organization and causal thickets by William Wimsatt.

4. Although it is, lamentably, a theory that has received almost no philosophical attention, I will use Navier-Stokes theory as a recurring example to illustrate and substantiate my arguments throughout the paper. I do this for several reasons, among the most important to my mind being that: it provides extraordinarily good models of complex, real phenomena we are all familiar with (water— H_2O —under normal circumstances, for example); it is, with respect to physical content, a well understood theory; and it is a straightforward theory, in the sense that its physical content can be easily explicated without the use of any heavy technical machinery. (The *loci classici* for the general theoretical treatment of hydrodynamic phenomena is Lamb 1932 and Landau and Lifschitz 1975, and that for issues of stability and equilibrium in particular is Chandrasekhar 1961; see, *e.g.*, Pope 2000 and Foias et al. 2001 for a treatment of non-equilibrium fluid phenomena such as turbulence.) Nonetheless, the full richness and complexity of the theory, both mathematically and physically, lie well beyond us—we have only a poor grasp of the mathematical character of generic solutions to the Navier-Stokes equations, and many of the phenomena most characteristic of Navier-Stokes fluids, such as turbulence, remain among the most deeply puzzling and mysterious in a physical sense that we are aware of.

5. If one is permitted to speak of explicating a question, one may consider this a first step on the way to a precise explication of the question “What is a physical theory?”

formal structure) are isomorphic in a (supposedly) relevant sense, with perhaps a gesture at the possible “representational capacities” of the formalism as an afterthought.⁶

I do not like a formulation of the question that admits as an appropriate answer such a purely formal structure, along with perhaps the fixing of something like a Tarskian semantics or a sketch of representational capacities. A formulation of that kind assumes, with no argument ever given, that it is possible to make a clean separation between, on the one hand, one part of the scientific knowledge general relativity embodies, *viz.*, that encoded in the pure formalism and, on the other, the remainder of that knowledge. The remainder includes at a minimum what is encoded in the practice of modeling particular systems, of performing experiments, of bringing the results of theory and experiment into mutually fruitful contact—including all the *inexact* mathematical methods that cannot be formalized, approximative and heuristic techniques motivated by loose physical arguments and principles not part of the formalism or of any interpretive postulates, and justified only by practical success—including, in sum, real application of the theory in actual scientific practice. This assumption comes out most clearly in the picture of semantics that naturally and standardly accompanies the standard view of theories: semantics is fixed by ontology’s shining City on the Hill, and all epistemology and methodology and other practical issues and considerations are segregated to the ghetto of the theory’s pragmatics. We should not assume such a clean separation is possible without an argument. Call this the ‘separability problem’.

If one accepts my formulation of the question as fruitful and interesting, then it becomes plausible that the theory of general relativity does not—can not—in any interesting sense consist of the theory of Lorentzian 4-manifolds along with a few interpretive postulates or a family of Tarskian models or a class of possible worlds or a category-theoretic structure. Now, it may turn out that a compelling analysis attempting to answer my formulation will show that, in the end, the standard assumption of the clean separation of the relevant domains of knowledge—the “theoretical” and the “practical”, for lack of better terms—is a good and justified one. I personally doubt it.

I will argue that the interplay between theory and practice (experimentation, observation, modeling, heuristic and approximative forms of argument, ...) is far more subtle, nuanced, complex and rich than many if not most philosophers steeped in formal traditions tend to appreciate, and in particular that the two are inextricably intermingled in a strong sense. The sum total of scientific knowledge constituted by the theoretical and practical content of the theory—which for lack of a better term I’ll call the “epistemic content”—essentially includes knowledge we have gained, and more importantly could have gained only, through practical modeling, experimentation and observation, heuristic and approximative techniques of argument, and in particular the application of representations and models of such experiments and observations in our theory by use of heuristic and approximative techniques of argument. It essentially contains such knowledge, moreover, in a way that cannot at bottom be cleanly separated from any knowledge grounded in the theory’s pure formalism. This is why I think my formulation of the question based on the idea of knowledge is a fruitful sharpening and refinement, for it is exactly this sum of epistemic content that distinguishes general relativity as a *physical* theory from a merely mathematical structure. (It also, in the end,

6. Weatherall (2017) is a classic example. I do not single this paper out because I think it is a poor one; to the contrary, I think it is a beautiful paper, rich with important insights. It does, nonetheless, serve as a particularly useful foil to my views, in part because of its clarity and depth.

is what must serve as the decisive basis for any criterion by which we judge general relativity to be both a different and a better theory than, *e.g.*, Newtonian gravitational theory.) I emphasize from the start that my conception of knowledge here is not limited to anything like “true, justified belief”, but includes more importantly comprehension and understanding of the sort that often accompanies explanatory practice (on which, more below).

I will further argue that a crucial component of that epistemic content consists in knowing how to schematize the observer in theoretical representations of experiments, and cannot be grasped without models of experiments in which the observer is schematized.⁷ By “schematize the observer”, as will become clearer, I mean something like: in a model of an experiment, to provide a representation of something like a measuring apparatus, even if only of the simplest and most abstract form, that allows us to interpret the model *as* a model of an experiment or an observation. As I discuss below, even such “trivial” schematic representations as a line of sight for telescopic observations in astronomy can allow one to determine and illuminate important and substantial parts of the way that theory and experiment come into contact, in this example by way of grounding an analysis, *e.g.*, of stellar aberration (on which, more below in §3). Such schematic representation, moreover, is not possible in a physically significant and cognitively substantive sense without recourse to all the other components of what I am calling the practical part of scientific knowledge, whose explication I will give in §6 below.

Because this is a long and dense paper, and the overall argument has not leant itself to a clean, simple, linear exposition, I provide a 1-page summary of the main claims of the paper in an appendix (§10), along with my main criticisms of and problems with what I call the standard view of theories.

3 Why We Need to Schematize the Observer

Because the meaning of scientific terms and propositions must rest on the knowledge we have of the physical world, and most of all on the knowledge we have gained through controlled observation and measurement, that is, through experiment, epistemic content accrues to a scientific theory in no small measure through the successful construction of empirically appropriate and adequate representations of physical systems in the theory’s terms. At bottom, then, what secure epistemic content a scientific theory has must rest in large part on the meanings expressed in the sound articulation of experimental knowledge. This requires at a minimum that we be able, at least in principle, to construct appropriate and adequate representations of actual experiments and observations in the frameworks of our best scientific theories, that is, representations of physical systems and experimental apparatus in relation to each other as required by particular, actual experiments, not just representations of physical systems *simpliciter*, in abstraction from actual experiments.⁸ It is characteristic of the ways that such models are constructed and applied, moreover, that they require methods and warrant—forms and aspects of knowledge—that go far

7. van Fraassen (2008, ch. 12), for example, an exemplar of what I am calling the standard view, explicitly argues for the opposite conclusion, even though he attempts to construct what he argues to be a pragmatist account of representation of physical systems by formal theory.

8. Indeed, it is our incapacity to do this in a consistent way in the context of quantum theory that lies at the bottom of the Measurement Problem. This alone shows the importance of the idea.

beyond what can plausibly be claimed to be captured by the theoretical formalism of a theory alone, no matter how many representational capacities one ascribes to it or interpretive postulates one affixes to it.

What is primarily at issue here is not the relation of representation itself, which can at most tell us “what the world would be like if the theory were true or largely true of it”, but rather the understanding and comprehension we have of the world in so far as we are warranted in thinking of our theories’ models *as* representations, which goes far beyond a description of how the world would be under given conditions. Such understanding and comprehension come from arguments and investigations that cannot always be “represented” by the formalism of “the theory itself”, but rather requires for their elucidation a rich and nebulous halo of inexact mathematical techniques and heuristic physical principles applied in ways that do not lend themselves to clean formalization.

There are, as I have intimated, several reasons why I claim that schematizing the observer is required for an adequate philosophical analysis of the structure and semantics of theories. I discuss now a few of them in more detail. The first is a shallow but still important one: sometimes the nature of the observational process itself results in “distortion” of the magnitudes measured, and a proper computation of the “real” values of the magnitudes of the system’s properties requires explicit modeling of the interaction between the measuring instrument and the system itself to correct for the effect.⁹ Call this ‘the problem of seeing through a glass darkly’.¹⁰ An example is stellar aberration: when light from a star enters a telescope, the motion of the telescope transverse to the path of the light while the light traverses the telescope (because of the diurnal rotation of the Earth, the orbital motion of the Earth around the Sun, and so on) makes the star appear displaced from its actual position in the sky; in order to correct for the effect, one must compute the actual motion of the measuring device, which requires an explicit representation of it in one’s model of the observation.¹¹ Another good example from precision astrometry in determining the “real” position of an observed body is the correction required by accounting not only for the bending of light around the Sun according to general relativity, but also for the fact that the bending changes with the time of the year as the position of the Earth changes in its orbit around the Sun (Thompson, Moran, and Swenson 2017, ch. 12, §6).

The second reason is a middling deep one. The quantitative results of all measurements and observations, even the best ones, inevitably deviate from those predicted by theory, even if only by a small amount; likewise, there is an inevitable imprecision in the measured values. Such deviations largely accrue to measurements on account of systematic errors arising from the idiosyncrasy of the particular experimental apparatuses used and the ways they are actually configured and deployed during the measurement process; the imprecision is an inevitable artifact of the limited acuity of any probe.¹² In order to compute reasonable values for the expected errors and imprecision (so as, for example, to be able to say when a measured result differs by an inadmissibly large amount

9. For simplicity, construe this claim only in the context of non-quantum theories; the same reason applies in many cases in quantum mechanics, but there an appropriate analysis requires far more mathematical, physical and conceptual sophistication and subtlety than I have room for in this paper, and, more to the point, than I feel capable of giving a sound philosophical articulation of, given our (my!) current state of understanding of quantum theory.

10. I thank Jeremy Butterfield for suggesting this excellent moniker.

11. See any good book on astrometry, such as Kovalevsky and Seidelmann (2004), for a full treatment of aberration and the details of computing corrections for it.

12. I ignore error arising from unknown sources, “noise”, which raises its own fascinating set of problems, well beyond the scope of this paper.

from a theoretically predicted result), one must often take account of fine details of the measuring apparatus and the particulars of its coupling to the system under study in one's model of the experiment. Call this 'the expected error problem'. Thermometry provides an excellent example: different sorts of thermometers (bulbs of gas, pyrometers, inhomogeneous thermocouples, *et al.*) couple to systems in radically different ways, the fine details of which must be handled on a case-by-case basis in order to correct for the effects of such phenomena as convective currents in fluids. (I discuss the example of temperature at greater length at the end of this section.)

The third reason is a deeper one. In order for theory to be able to provide guidance to experiment in the design of new types of tools for probing novel sorts of phenomena and in the design of new types of tools for probing known phenomena in novel ways, and conversely for experiment to provide guidance to theory in modeling practically constructed novel ways of coupling to systems known and unknown so as to place constraints on the possible soundness of theoretical description and prediction, theory must be able to represent the fine details of the apparatus as it is actually to be used in the experiment. This particular interplay between theory and experiment, in theoretical guidance in the construction of instruments and in experimental constraint on the soundness of theory, is one of the most profound ways that theory and experiment are able to make contact with each other; without it, it is difficult to see how any epistemic content could accrue to theory in the first place.

The search by Hertz for ways to produce and detect free electromagnetic waves as predicted by Maxwell's theory provides a beautiful illustration of the delicate dialectic required between theory and experiment, especially in the construction and modeling of instruments in the attempt to produce and probe a phenomenon so poorly understood.¹³ During most of the career of the investigation, Hertz had very little idea what sorts of arrangements of what sorts of physical system would produce electromagnetic waves in the first place, and even less of what sorts of instrument could reliably detect them. His search necessarily included the construction of finely detailed models both experimental and theoretical, each guiding the other in turn, of different proposed methods of coupling the electromagnetic field to its environment and different instruments to try to realize those couplings.¹⁴ Of course, everyone at the time knew that one could derive the wave equation for the electromagnetic field from Maxwell's equations, and they knew the plane-wave solutions and the principle of linear superposition for electromagnetic fields, but I would say that, at that point in the history of Maxwell theory, the theory did not include a characterization of electromagnetic radiation with substantive epistemic content. In one sense, as Hertz famously remarked, Maxwell's theory may be his equations, but the epistemic content, including the practical knowledge, that characterizes those equations *as* a theory of *electromagnetism* does not exist in a meaningful way until we understand it well enough to use it to make testable predictions—but a prediction is not testable if we don't know in principle how to test it!¹⁵ This is not verificationism

13. See Hertz (1893, *passim*), including the preface by Helmholtz, for an absorbing account.

14. Buchwald (1994) provides a detailed and comprehensive exposition of Hertz's search for an instrumental design and a theoretical representation that jointly would do the job. There is a close affinity between my own arguments and Buchwald's conclusion that, in the end, Hertz required only a highly abstract, "schematized" representation of the instrument—the dipole oscillator—for his purposes, and yet he still did need an explicit representation of it in his theoretical work, no matter how "simple" and abstract. I take the lesson to be that Hertz needed that schematized representation in order for the theoretical and the experimental parts of his work to make substantive contact with each other.

15. This all too brief discussion of Hertz bears profitable comparison with Harry Collins' idea of the experimenter's

about meaning. It is rather the simple point that, if we do not know how to bring a theoretical structure into contact with the world by experimentation or observation, then we have little or no grounds for trying to understand that structure as representing anything in the world.¹⁶

The fourth reason is the deepest, and depends in part on the previous three. Without an understanding of the fine details of the way that theories do and do not successfully model actual experiments of the kinds of physical systems they purport to treat, we have no way to demarcate the regime in which the theory is in fact appropriate and adequate for representing those kinds of systems—its *regime of applicability*. It is only by dint of the practical forms and aspects of scientific knowledge, in other words, that we can determine a theory's regime of applicability—when, *e.g.*, what we pre-theoretically conceive of as an actual body of fluid, manifesting some dynamical behavior (say, what looks like turbulence), is appropriately and adequately treated by Navier-Stokes theory and when it is not. Knowledge of that regime, however, embodies perhaps the most important part of the epistemic content of a theory—without it, we have no warrant at all for trusting our use of the theory, much less for understanding it *as* a theory of a particular class of physical systems. Nothing in the theory's formalism can tell us this, without the input of knowledge that can be had only by experimental practices that require schematizing the observer in theoretical models.¹⁷

In order to more deeply explore the last two reasons—instrument design and demarcating a theory's regime of applicability—I will discuss in some detail the International Practical Temperature Scale of 1927, as revised in 1948, 1955 and 1960,¹⁸ as it provides an excellent, concrete example of the sorts of considerations one must take account of in attempting to define a physical quantity—that is, to render epistemic content for any representation of it—in the context of physical theory.¹⁹ The thermodynamic temperature scale between the primary fixed point 0.01 °C (the triple point of water at one standard atmosphere) and the secondary fixed point 630.5 °C (the equilibrium point between liquid and solid antimony at one standard atmosphere) is defined by interpolation, using the variation in resistance of a standard platinum wire according to the equation of Callendar

regress; see, *e.g.*, Collins (1992, ch. 4). In contradistinction to Collins, however, I consider this relation between theory and experiment to be not only unproblematic, and in particular not circular in the sense of leading to a regress, but rather necessary for the progress of science. I am thus more closely aligned in some ways with the views of Franklin (1986), many of whose criticisms of Collins I endorse.

16. See also the philosophically rich essays by Maxwell (1871; 1876) himself on this necessary sort of interplay between theory and experiment.

17. Although it does not bear on the question of how the schematic representation of the observer informs and is required by work in physics itself, I feel it is important to note that deep philosophical investigation of the foundations of physical theories also at times relies on it. The work of Geroch (1977), Malament (1977), and Manchak (2009) on predictability in and observational indistinguishability of relativistic spacetimes, and the work of Malament (2002, 2003) on defining relative rotation in relativistic spacetimes, are exemplary in this regard.

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

19. See, *e.g.*, Benedict (1969) for detailed exposition of all the technical matters I discuss here. It is not the most up-to-date reference with regard to the international agreement on defining the standard methods for the determination of temperature (for which see, *e.g.*, Haynes 2014), but I have found no better guide to the nuts and bolts of thermometry when it comes to developing in a physically illuminating way how the mathematical models and the experimental techniques bear on each other. Although I do not endorse all his conclusions, in particular those with a more Kuhnian tang, I recommend that the interested reader see Chang (2008) for an extensive and edifying discussion of many of the practical issues that arise in thermometry, and how they bear on philosophical problems closely related to those I am addressing.

(1887),

$$\tau = 100 \left(\frac{\rho_\tau - \rho_0}{\rho_{100} - \rho_0} \right) + \delta \left(\frac{\tau}{100} - 1 \right) \left(\frac{\tau}{100} \right)$$

where:

- τ is the temperature being interpolated from knowledge of the measured electrical resistance of the platinum;
- ρ_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 $\pm 0.001^\circ\text{C}$) by barometric pressure variations from 28.50 inches to 31.00 inches of mercury, and the resistance of the wire is independent of the static water pressure up to a depth of immersion of 6 inches at sea-level;
- ρ_{100} is the resistance of platinum as measured with the thermometer immersed in saturated steam in 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);
- ρ_τ is the resistance of platinum at temperature τ (the temperature being measured), *i.e.*, ρ_τ is itself the quantity being measured that allows the calculation of the system's temperature;
- and δ is a constant characteristic of the particular type and features of the platinum thermometer employed, defined at the primary fixed point 444.6°C (the equilibrium point between liquid and solid sulphur at one standard atmosphere).

Below 0.01°C and above 630.5°C , the Callendar equation quickly diverges from the thermodynamic scale. From 0.01°C down to the primary fixed point -182.97°C (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 630.5°C up to the primary fixed point 1063.0°C (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 1063.0°C , 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. 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 under the influence of fluctuations in the temperature of the system itself and its environment, under the given conditions, when the device comes into proper contact with the system under study.

I take it to be clear from this exposition that one cannot even define quantities such as temperature without explicit schematic representation of the observer, much less have understanding of how to employ their representations in scientific reasoning. Without such a definition, however, any theoretical proposition referring to temperature can have no epistemic content.

4 What a Theory Is

What, then, is the theory of black holes? It cannot be the subset of Lorentzian 4-manifolds with non-trivial event horizon.²⁰ That cannot tell us that the Milky Way has at its center a black hole of about 4 million Solar masses, about 88 million kilometers across, referred to by astronomers as ‘Sagittarius A*’ (usually abbreviated ‘SgrA*’)—even putting aside the fact that it is likely that the thing has nothing associated with it remotely resembling an event horizon in the standard, formal sense. Even if we throw in interpretive postulates or a Tarskian-like semantics, it cannot tell us how we know this, why we think we are justified in believing it, the ways we may be justified in using it as evidential warrant for further knowledge claims, and so on—all of which counts as part of the epistemic content of the theory.

To this observation, an advocate of the standard view may reply that a particular prediction of the theory starting from particular initial data is not “part of the theory” in any interesting sense—we must, she will say, sharply distinguish between, on the one hand, the formal structures of the theory that encode its form of nomic possibility and, on the other, the possibly realized individual physical systems characterized by contingent initial or boundary conditions. Nonsense, I say. First, an *ad hoc* argument: on the standard view, SgrA* is one of the Tarskian models, possible worlds, or what have you, that characterizes the theory. Thus scientific knowledge that we have about SgrA* forms part of the theory even on the standard view. That leads to the deeper response, to point once again to the question whether one can cogently articulate the character and content of that knowledge in isolation from the other forms or aspects of scientific knowledge, the more practical ones, we have about SgrA*: one cannot identify the putative Tarskian-like model of the theory, or find the correct application of the interpretive postulates required to identify a solution to the field equations as representing SgrA*, in isolation from the practical knowledge used to characterize it as a physical system amenable to representation by general relativity. Standard interpretive postulates such as “freely falling particles traverse timelike geodesics” are manifestly inadequate for the job.

The problem for the standard view is even more severe than these remarks suggest. Criteria for what exactly we are to count as part of SgrA* and what not to count for the purpose of tallying the theory’s models are never specified by advocates of the standard view. Should we include the accretion disk? The mini-cluster of rapidly orbiting stars immediately surrounding it? The frozen magnetic field surrounding the horizon and ionized plasma jets shooting along it? Does adding or subtracting each of those yield a new model of the same target system or just a model of a different target system? Does including inhomogeneities in the accretion disk or the magnetic field or the plasma count as refining an existing model or as creating (or identifying) a completely new model? What about accounting for the tidal deformation of the orbiting stars? And so on. And in each case, once we have decided, what type of model are we to use to represent the phenomena in question? One using magnetohydrodynamics to model the accretion disk, or a statistical mechanics of charged plasma? Both are encompassed in the framework of general relativity.

In our current state of knowledge—given the actual epistemic content of general relativity as

20. I put aside here the fascinating question of what physicists mean in the first place by the idea of a black hole, the manifest fact that different fields of physics, and often even different physicists in the same sub-field, have different, often mutually inconsistent, definitions of ‘black hole’. See Curiel (2019a) for discussion.

we know it—there is no single, unambiguously correct answer to those questions. Different ways of modeling the phenomena yield different families of solutions to the Einstein field equation in general relativity, none of which are related to each other in any simple way: none of the families “reduce” to the others as one lets some parameter smoothly go to zero, none of the families are subsets of the others, none of the families non-trivially intersect with the others, and so on. For example, depending on one’s purposes, one can treat the accretion disk as an uncharged dust or fluid or as a charged plasma, and, if the last, one can either treat the plasma as a source or treat it as not a source (“test matter”) for the ambient magnetic field. This will be justified, *e.g.*, by a solid practical understanding of the aspects, features and components of the system one’s observational instruments will couple to in the kinds of measurements one plans to make. For some wavelengths of electromagnetic radiation one aims to detect, one can safely ignore the inhomogeneities and anisotropies in the magnetic field arising from treating the plasma as a source, and for others not, depending, *inter alia*, on whether those inhomogeneities and anisotropies are comparable in scale to the wavelengths one plans to study. Whether one treats it as charged is irrelevant, for example, if one is making measurements in the gamma-ray wavelength spectrum, as electromagnetic radiation of that type is insensitive to the differences. Similarly, when one treats the constituents of the accretion disk as uncharged, then whether one treats it as a dust or as a fluid is irrelevant if one is making observations in the x-ray band, as the system’s generation of x-rays is insensitive to the differences.

In all these cases, it may seem as though the one type of model is just a simplification of the other, as (say) one allows some quantity or coupling to go to zero. This is not so. One may think that the models based on dust are just the fluid models in which the pressure and viscosity are taken to zero. The procedure just described, however—“let pressure and viscosity go to zero”—is ambiguous, as there are many ways to construct such a limiting family in general relativity that each yields a different kind of spacetime as the limit (Geroch 1969). In any event, the “obvious” ways of doing it give the wrong answer in some cases: it is a theorem by Ellis (1967) that if a pressureless dust is shear-free, then either its twist or its expansion must be zero. Many fluid models of accretion disks, however, are both shear-free and have non-vanishing twist and expansion, but there are no dust models that correspond to such behavior. (Recall that the Einstein field equation is highly non-linear, so the behavior of the limits of such procedures as “take this parameter to 0” cannot always be simply “read off” the mathematics in a naive way.) There is no systematic relation between the dust and the fluid models of the accretion disk.²¹

In sum, there is no single, canonical model that represents SgrA*, nor even a single, canonical, naturally related family of models. We rather have several different, not unambiguously related families of models, each of which captures some aspect, feature or component of the system not captured by the others, and each of which is appropriate and adequate for the different inves-

21. Indeed, even just from a purely mathematical perspective, the situation is more complex than I have sketched, and more dire for the standard view. What, *e.g.*, is the level of differentiability we demand or require of solutions? Depending on the answer to the questions, we may or may not be able to represent such phenomena as shock waves. It is still, in fact, an open theoretical and mathematical question whether generic shock waves in the standard sense (Landau and Lifschitz 1975) can be cogently represented in general relativity outside the regime of exact spherical symmetry (Reintjes and Temple 2017). So, are representations of shock waves in non-symmetrical systems part of the theory of general relativity itself? Please do not say there is an answer to that question in the Platonic heaven of the space of formal solutions to the Einstein field equation. That would not answer the question about the physics.

tigative purposes one may have—each gives what we believe to be excellent answers to kinds of questions that the others answer terribly. The propriety and adequacy of the model for many given purposes, moreover, is determined in large part by the kind of knowledge that can be had only by comprehension of the practical part of the theory—what sorts of instruments with what sorts of sensitivities are appropriate and adequate for probing and studying those different aspects, features and components of the system, in which states and under what sorts of conditions, and the kinds of arguments (almost never articulable strictly in the theoretical formalism) that legitimately show this.

There is also the fact that we do not think general relativity is at bottom correct as a theory of spatiotemporal structure and its relation to matter; at some point, presumably, quantum effects must be taken into account, and we have little to no idea how to do that. In that sense, *no* model general relativity produces can possibly be “the right one” for SgrA*. Physicists almost universally ignore this fact in their work, and with good reason: even were we to have a satisfactory theory of quantum gravity, it is almost certainly the case that we would be unable to use it to construct models of SgrA* to answer the questions we have about it.

Indeed, a crucial part of the epistemic content of the theory is the knowledge of the physical regime in which it is appropriate and adequate for representing a given class of phenomena at all. As Geroch (2001, pp. 6–7) concisely puts it, in discussing attempts to formulate a relativistic version of Navier-Stokes theory,

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.

The physical quantities modeled by a theory’s equations have a regime in which they are simultaneously well-defined, satisfy the theory’s equations and have values stable with respect to fluctuations and effects that are not representable in the theory (such as quantum effects in classical general relativity, or molecular-scale fluctuations in a theory of fluids). That regime itself forms part of the core of the theory, in so far as it determines what the scope and depth of the theory’s empirical content can be. As I’ll discuss below in §§8–9, moreover, until one has specified a regime of applicability for the theory, such questions as I’ve raised about SgrA* cannot be answered, but to specify such a regime requires of necessity knowledge that can in part be had only by the construction of models of experiments that include schematic representation of the observer. One cannot determine the regime of applicability by examination of the formal structures on the one hand and a cleanly separated family of interpretive postulates or Tarskian-like models on the other. The interpretive postulates on their own cannot do the job, for they require the systems to which they apply to fall already within the regime of applicability; they do not determine it.

That we feel secure in claiming that SgrA* is a black hole of that mass and size—that we believe this to be *scientific* knowledge—depends in large part on the confidence we have in the experiments and observations that delivered the data that allow us to identify and characterize SgrA* as a black hole, and moreover as *that* black hole. That confidence in turn depends on the confidence we have in the methods we use to model those experiments and observations and to manipulate and impose structure on the raw data collected so as to put it into a form that can make fruitful, substantive, physically significant contact with the abstract, formal structures

of pure theory. And much of the knowledge that constitutes and underlies the experiments and observations—their results, the confidence we have in those results, the way we use those results as evidence, the confidence we have in doing that—both derives in large part from and consists in large part of the knowledge that comes from constructing a model of the experiments and observations themselves, which inescapably includes schematizing the observer.²² This also necessarily involves epistemic content from other theories—those, *e.g.*, we use to model the measuring instruments and the environment in which the experiment or observation is taking place—and the knowledge we have of how the epistemic content of those other theories interacts with that of the theory at issue. That itself forms a critical part of the epistemic content of the theory we are considering, how it admits input from other theories.

A theory in my sense thus includes such things as accounts of experimental devices appropriate for studying the relevant kinds of physical system, good practices for employing them, sound techniques for the collection of raw data and for the statistical and other analysis and organization of the same, reliable methods of approximative and heuristic reasoning for constructing models and solving equations, and guidelines for determining whether a system of the given kind is in a state and experiencing interactions with its environment such as to be jointly amenable to appropriate and adequate representation by the theory (*viz.*, whether or not the system falls into the theory’s regime of applicability), and so on.

Contrarily, it is characteristic of an appropriately unified kind of physical system, one treated by a theory in my sense, that there exist a set of scales at each of which all quantities the theory attributes to the kind of system simultaneously lose definition. In other words, every theory, in so far as it treats an appropriately unified kind of physical system, not only has a regime of applicability, but it has a *single, unified* one, bounded on all sides by scales characterized by the values of different combinations of its physical quantities. For classical fluids, for example, the definitions of their pressure, fluid flow, viscosity, and all the rest break down at spatial and temporal scales a few orders of magnitude greater than those of the mean free-path of the fluid’s constituent particles. There is no *a priori* reason why the definitions of all the different physical quantities represented by the theory should fail at the same characteristic scale, even though, in fact, those of all known examples do, not only for theories of classical fluids but for all physical theories we have. This seems, indeed, to be one of the markers of a physical theory, the existence of a set of characteristic scales for its physical quantities, at each of which all the theory’s physical quantities simultaneously lose definition.²³ (I will give a more thorough and detailed explication of the ideas of the breakdown scales of a theory and its regime of applicability in §8 below.)

In sum, a theory must include all the practices and principles that allow one to meaningfully bring formal theory and practical experiment into contact with each other so that the former may be used to interpret the latter, and the latter may be used to constrain the former, *i.e.*, so that

22. See Genzel et al. (1997) and Ghez et al. (2000) for accounts of the detailed infrared studies that nailed once and for all our confidence that SgrA* is indeed a black hole, and in particular the way that models of the experimental apparatus and the observations themselves were crucial to the arguments. Collmar et al. (1998) and Eckart et al. (2017) provide illuminating discussions, almost all by the physicists directly involved in the work themselves, of the fascinating methodological and epistemological problems associated with trying to ascertain that what we observe astronomically when we point our telescopes at SgrA* is really a black hole.

23. This is a fact that deserves philosophical investigation. I discuss it in more detail in Curiel (2017b), but with no pretense of even touching on all the interesting questions this fact raises.

data may be structured in such a way as to be directly comparable with, and even identified with, *e.g.*, formal solutions to the equations of motion. This will include, *inter alia*:

- mathematical structures, relations and formulæ over and above the abstract equations of motion or field equations (*e.g.*, in Newtonian mechanics that $\vec{v} =_{\text{af}} \dot{\vec{x}}$, that mass is additive, that spacetime has a flat affine structure, and so on);
- standards of good argumentation (accepted approximative techniques for solving equations, sound heuristics for informal arguments, and so on);
- families of accepted experimental and observational practices for systems of different types;
- rules for connecting experimental outcomes with formal propositions (principles of representation, guidelines for the construction of concrete data models from raw observations, rules for reckoning expected experimental precision and error, and so on);
- rules of evidential warrant (what can be evidence, how to apply it, reckoning of error tolerance, and so on);
- and guidelines for judging the legitimacy of proposed modifications, extensions, and restrictions of all these.

Most of this will be difficult if not impossible to articulate and record in an exhaustive and precise way, so as to lend itself to use in formal philosophical investigations. We must trust that all such collateral principles and practices are there, and can be, now and again, each more or less precisely articulated as the occasion demands. The same holds true, however, for formal reconstructions of all forms of reasoning in science (*e.g.*, the ‘auxiliary hypotheses’ of the hypothetico-deductive method, which always hide an ugly mob of philosophical sins.)

On this view, the entirety of a theory is a dynamic entity, evolving over time as new theoretical and experimental techniques and practices are developed and accepted. I believe this is the right way to think about these matters for many if not most purposes in those parts of philosophy of science studying scientific theories. The contemporary practice of treating theories as static, fixed entities, especially in work of a more technical and formal character, can lead to serious philosophical error. An adequate semantics of a theory, for instance, should reflect and accommodate its dynamic nature.

Before moving on, I must emphasize that I am not opposed to the use of formal machinery in philosophy of science in general or in the study of theories in particular, far from it. I am opposed only to a certain conception of how formalism can be fruitfully conceived of as representing or capturing (part of) the structure of scientific knowledge. In the end, the standard view is a philosophical idealization. Unlike idealizations in physics, however, its use is not susceptible to justification or defense in the way that idealizations in physics often are, by the use of a quantitative measure of the expected error introduced by the idealization, and the subsequent judgement whether that error falls within one’s error tolerance. The kinds of error the standard view introduces are not controllable in that fashion. One must argue for the goodness of the idealization on a case by case basis, in a way that respects the peculiar demands of the given philosophical context. There are certainly cases in which it is manifestly a good idealization. A good example are the arguments by Malament (2002, 2003) about the ambiguity of the idea of relative rotation in general relativity, and some of the unintuitive features that different criteria for whether a body

is rotating relative to another or not can have. Malament's claims, based on arguments relying only on the formalism in conjunction with a few interpretive postulates, are indisputably correct, and tell us something of great conceptual and physical interest about the theory. It can also be of great interest, and teach us much about a theory, to explore its space of formal possibilities, even when we know that no corresponding physical systems can exist in the actual world. Again, however, conclusions about what such explorations can teach us must be made with account taken of the fact that one ignores much if not most of the epistemic content of a theory while exploring.

I conclude this section with a lengthy quote giving a beautifully illustrative description of how a practicing physicist views theories in his own thought and work (Geroch 1981, ch. 7, pp. 182–183), with obvious points of contact and sympathetic relations with the view I have been advocating, and from which I have taken inspiration:

One might very well be left with the impression that the theory itself is rather hollow: What are the postulates of the theory? Where are the demonstrations that all else follows from these postulates? Where is the theory proven? On what grounds, if any, should one believe the theory? I can only answer these questions with my own opinions. It seems to me that “theories of physics” have, in the main, gotten a terrible press. The view has somehow come to be rampant that such theories are precise, highly logical, ultimately “proved.” In my opinion, at least, this is simply not the case—not the case for general relativity and not the case for any other theory in physics. First, theories, in my view, consist of an enormous number of ideas, arguments, hunches, vague feelings, value judgments, and so on, all arranged in a maze. These various ingredients are connected in a complicated way. It is this entire body of material that is “the theory.” One’s mental picture of the theory is this nebulous mass taken as a whole. In presenting the theory, however, one can hardly attempt to present a “nebulous mass taken as a whole.” One is thus forced to rearrange it so that it is linear, consisting of one point after another, each connected in some more or less direct way with its predecessor. What is supposed to happen is that one who learns the theory, presented in this linear way, then proceeds to form his own “nebulous mass taken as a whole.” The points are all rearranged, numerous new connections between those points are introduced, hunches and vague feelings come into play, and so on. In one’s own approach to the theory, one normally makes no attempt to isolate a few of these points to be called “postulates.” One makes no attempt to derive the rest of the theory from postulates. (What, indeed, could it mean to “derive” something about the physical world?) One makes no attempt to “prove” the theory, or any part of it. (I don’t even know what a “proof” could mean in this context. I wouldn’t recognize a “proof of a physical theory” if I saw one.)

We get misled in philosophy by a picture—we fixate, I think, on how we are *taught* the physics, but we forget how we *learned* it, and most of all how we know that we are warranted in believing it—believing it, that is, to be appropriately and adequately about some part of the physical world.

5 What Measurements and the Observer Are

The relevant notion of measurement and observation at play here is a “non-primitive” one, whose explication in the context of a particular theory, or as an adjunct to a particular theory, is always at least in part mediated by the theory.²⁴ I mean something like: “measurement as coupling between two individual systems, usually of different types”, such that:

1. one can distinguish the target system (the one being measured) from the measuring system, in the sense that there are some quantities of the target system and some quantities of the measuring system that dynamically evolve independently of each other;
2. there is at least one fixed quantity (the *measured quantity*) borne by the target system whose value will be determined by the measurement;
3. there is a set of quantities (possibly a singleton, but not empty, the *measuring quantities*) borne by the measuring system that jointly couple to the quantity being measured, in the sense that the measured quantity and possibly some of its derivatives are algebraic functions of the measuring quantities;²⁵
4. the minimal temporal interval required for a measurement to take place is small compared to some relevant, characteristic time-scale associated with the dynamics of the state of the measured system during the time of measurement.

Some simple examples will clarify what is meant. For a gas-bulb thermometer, the gas’s volume (the measuring quantity) is an algebraic function of the temperature of the system being measured; the equilibration time, *i.e.*, the time it takes for the volume of the thermometer to achieve its final state directly proportional to the temperature of the measured system, is small compared to the characteristic time-scale in which the measured system will change its temperature in response to any external influences that may be present. When the acceleration of a test-particle of known mass and charge is used to measure the vectorial value of an ambient, static electric field, the particle’s acceleration is an algebraic function of the value of the electric field; in this case, the acceleration is an “instantaneous” response to the intensity and orientation of the electric field. Components of the Riemann tensor are algebraic functions of the differential accelerations of a collection of nearby freely falling particles, in accord with the equation of geodesic deviation, and those differential accelerations serve as measuring quantities for those components.

It is a deep and difficult question, to ask how one “reads off” the results of a measurement from the values of the measuring quantities. Naively, it seems as though the process must either continue in an infinite regress or terminate in a measuring quantity whose values are amenable to direct inspection by the ordinary human sensory apparatus. But such a conclusion comes from

24. It is, therefore, exactly the kind of measurement that does not pertain to the Measurement Problem in quantum mechanics.

25. *Cf.* Geroch (1996, pp. 3–4, arXiv version):

Roughly speaking, two fields are part of the same physical system if their derivative-terms cannot be separated into individual equations; and one field is a background for another if the former appears algebraically in the derivative-terms of the latter. The kinematical (algebraic) interactions are the more familiar couplings between physical systems.

too quickly accepting the raw empiricist's naiveté—I think rather that the distinction drawn in discussions of empiricism and realism, between “theoretical” terms and “observable” terms, is a red herring. Our sensory organs are instruments, and we know—through theoretically grounded experimental investigations—that they are often unreliable ones, and we know, moreover, in quantifiable ways, experimentally determined under precisely specifiable conditions, how they are unreliable: the ways and the magnitudes of the errors. We know the same thing about other instruments of measurement we use in experiments, often in greater, more precise and more quantitatively exact detail. It is only knowledge gained through such theoretically controlled means, *i.e.*, gained using instruments that we have some measure of a theoretical grip on the inaccuracies of, that can contribute in the strongest form to the epistemic content of a theory. Thus, the relevant distinction for the analysis of scientific knowledge and the nature of scientific theories is between purely theoretical propositions and experimental propositions based on measurements using well understood instruments, whether those be sensory organs or electron microscopes. So far as *scientific* knowledge is concerned, the epistemological status of the different kinds of instruments is the same, so long as their various forms of inexactitudes and systematic errors are amenable to theoretical and experimental characterization and testing.²⁶

The results of perception are *not* more epistemically fundamental, primitive or secure *with respect to* the constitution of an evidentiary basis for a scientific claim. *That* is a true dogma of empiricism, that perception is the most epistemically secure source of knowledge. Both the results of perception and those of experimentation must themselves be secured—their reliability, veracity, accuracy, precision, scope, fallibility, *etc.*—by scientific investigation, each every bit as much as the other.

To my mind, the theoretical in science, as supported and substantiated by controlled experimentation, is on a more sound epistemic footing than the perceptual. To be clear, I mean that the theoretical both has more epistemic warrant and plays a more fundamental role in the construction of epistemic warrant than the perceptual *with regard to their respective roles in scientific enterprises*, precisely because of the substantive contact of the theoretical with the experimental, in the relation between which we have many powerful methods of measuring, tracking, manipulating, controlling for, and justifying essentially all the kinds of approximation, idealization and error that such investigations must suffer, none of which we have for perception. And, in any event, the perceptual is almost entirely irrelevant for the philosophical comprehension and understanding of the nature of science and of scientific knowledge, as is shown by the fact that, outside physiology,

26. My view has some affinity with that of Helmholtz (1870, p. 12, from the English translation given in the bibliography reference):

When we measure, we only effect with the best and most reliable aids known to us the same thing that we ordinarily ascertain through observation by visual estimation, judgment by touch, or pacing off. In these cases, our own body with its organs is the measuring instrument that we carry about in space. Our compasses are now the hand, now the limbs; or the eye, turning towards all directions, is our theodolite, with which we measure arc-lengths or surface-angles in the visual field. ¶ Every comparison of magnitudes of spatial relations, whether by estimation or by measurement, thus proceeds from a presupposition about the physical behavior of certain natural bodies, whether our own body or the measuring instruments employed; a presupposition which, moreover, may possess the highest degree of probability, and may stand in the best agreement with all other physical relationships known to us, but which in any case reaches beyond the domain of pure spatial intuition.

no science has the theoretical capacity to include sense organs in its models. This, I think, has been one of the main and most pernicious sources of error in so much of Twentieth Century philosophy of science, inherited from Kant as worked through the mangle of the Logical Empiricists: that the most relevant epistemic distinction in the analysis of science is the theoretical versus the observable. It's not. The most important and fundamental distinction is between the theoretical and the experimental, and the most important correlative question is, how does it happen that these two dramatically different representations of the physical world not only come into fruitful contact with each other, but indeed need and rely on each other in such intimate ways in order to be able to come into contact with the physical world in the first place. (My one-sentence sketch is, of course, a caricature of both Kant and the Logical Empiricists, but it's one that, astonishingly, continues to guide much contemporary philosophical work.)

Science itself shows us how unreliable our perceptual faculties are, with great precision and rigor. If we accept those findings—and I think we must, if we are to accept anything in science at all, as those findings are grounded in and derive from the same sorts of experimentation, reasoning, *etc.*, as all our most deeply held scientific results—then it follows *a fortiori* that perception and the deliverances thereof are less epistemically secure than experimental knowledge as represented in our theoretical formalism, for we rely on experiment to characterize the epistemic reach, precision and accuracy of perception.

In any event, how a real human comes to know the result of a measurement is an independent question from that of what *constitutes* a measurement and what constitutes an epistemically warranted judgement of a sound and accurate *outcome* of an experiment. That we see a pointer on an instrument indicating the value measured as a result of an experiment to be 3 is a fact of little to no relevant epistemic import. How and that we have good reason to believe that the result we expect is 3 and that the reading '3' is the result of the correct functioning of the experiment as a whole is what matters, and that comes entirely from our reasons for believing that our experiment is latching on to the germane parts or aspects of the world in the right way and that our theoretical representations of the experiment appropriately and adequately captures that latching on to.²⁷

For all these reasons, I want to speak of experimentation, measurement and observation as practices and processes that can be characterized independent of the direct experience of humans, though certainly not independent of the current epistemic state and capacity of humans. By “schematize the observer”, therefore, I mean something like: in a model of an experiment, to provide a representation of something like a measuring apparatus, even if only of the simplest and most abstract form, that allows us to interpret the model *as* a model of an experiment or an observation. Such schematic representation, moreover, is not possible in a physically significant and cognitively substantive sense without recourse to all the other components of what I will call the practical part of scientific knowledge, as spelled out in the next section.

27. Compare the remarks of Popper (1959, part II, ch. 5) to the effect that an observation sentence is not an autobiographical report, but a conventional statement of socially accepted scientific fact.

6 The Different Forms and Aspects of Scientific Knowledge

In theory, there's no difference between theory and practice.

In practice, there is.

– Yogi Berra

An adequate semantics should be able to support the articulation of, to represent the meaning of, all knowledge claims a theory embodies, for all the forms and aspects of knowledge thus embodied. As the discussion of SgrA* in §4 highlights, my formulation of the question “what is a theory?” brings into focus what I consider another shortcoming of the standard view (that there is a clean separation between, on the one hand, one part of the scientific knowledge a theory embodies, *viz.*, that encoded in the pure formalism and, on the other, the remainder of that knowledge)—that, according to it, all scientific knowledge is essentially of one undifferentiated kind: approximately true (or at least epistemically justified) propositions about states of the world, and perhaps about the global or nomic structure of the world independent of the state the world is actually in. This comes out most clearly in the picture of a theory's semantics that naturally accompanies the standard view of theories, that it is fixed by ontology.

This is not an accurate picture of scientific knowledge. Scientific knowledge has many forms and aspects that cannot be well captured with this simple picture. The structure of the knowledge that characterizes a theory as a *physical* theory of the types of physical systems it purports to treat must include *all* forms and aspects of knowledge the theory has endowed on us, including that gathered and expressible only in the context of experimental practice. A semantics that captures only knowledge—epistemic content—that it is impossible in principle for humans to know, such as that embodied in the “family of all models or possible worlds” of the theory, is useless as a semantics on which to ground an analysis of the nature of scientific knowledge and its epistemic warrant. If the semantics does not capture all the epistemic content of a theory, then what *does* it do? And why should we care about it?

A useful classification of what I have been calling the theoretical and the practical forms of scientific knowledge is suggested by the following observation of Stein (1994, p. 637, emphases his):

It is hard, but possible, to learn theory by self-study from books; it is surely much harder to learn experimental techniques without a teacher to help one acquire skills; but what I suspect to be impossible is to learn the principles of experiment without *actual* experience with the relevant *instruments*.

How one is able to learn different parts of the epistemic content of a theory may seem at first glance irrelevant to an analysis of the character of the knowledge to be learned. To the contrary, I think Stein has put his finger on something very deep about epistemic content: how it is possible to come to learn something fundamentally shapes and constrains the kinds of epistemic warrant we can have for it, and how we may bring it to bear in possible evidential relationships to other pieces of potential knowledge. The nature of the object of knowledge determines how one may learn about it, and so that constraint itself can teach us something about its nature. Thus, if one wants to understand the structure and character of our knowledge, in all its forms, one of the

deepest and most powerful clues we have is what we do and can learn, and how we do and can learn it, for how knowledge can be learned tells us how it can (and cannot) be represented—and thus whether it can be accommodated by formalism in conjunction with “coordinating principles” or “representational capacities”. One of the most important distinctions here is that between a subject matter’s being *teachable* and its being *learnable*.²⁸

Theoretical knowledge of the sort Stein describes in the first part of the quote has the peculiar and remarkable property of being representable by a well arranged sequence of propositions (a well written book, say). If it did not have that property, we could not learn it from such a sequence. Epistemic warrant for knowledge acquired in that way has its strengths and weaknesses. It relies on a particular form of trust in the scientific community,²⁹ and on a belief that one’s own idiolect shares enough in common with that of the author’s (or, at least, that one can learn enough of how the author uses language) to be certain that one has understood the epistemic content of what the author is trying to convey. It also, however, admits of being made precise and clear in a way that conduces to verification and defense.

Practical knowledge cannot be effectively represented in this way. An anecdote will perhaps show this better than an argument. During the 2017 run of the Event Horizon Telescope (<https://eventhorizontelescope.org/>), when a worldwide team of astronomers were co-ordinating the use of several observatories around the world to simultaneously record observations of SgrA* (for later structuring and collation) in the attempt to record the first direct image of a black hole, the head of the organization, Shep Doeleman, masterminded the proceedings in a control center two doors down from my office at the Black Hole Initiative. Naturally, I spent some time watching them, as this was scientific history in the making. One evening, while the team was deciding whether or not to make observations that night, I witnessed the following exchange between Shep and an instrumentalist at ALMA, the Chilean observatory (5000 meters above sea level), who was having trouble with one of the specially designed instruments. A quartz crystal oscillator was malfunctioning. Shep said, “Oh, yeah, I’ve see this before. Here’s what you do. Take’em down to 1000 meters [above sea level], and leave’em for a few hours; they’ll settle down, and then when you take them back up, they’ll work.” Chilean instrumentalist: “Why does that work?” Shep: “I have no idea, but it does.” This is a part of the practical knowledge associated with observational relativistic astronomy that I doubt anyone will ever find written in any book. Moreover, even if written, it is the sort of thing one could not *learn* from just having read the book. Of course, one could learn that particular trick—but to learn how to *play* with the instruments so as to be able acquire the feel for them necessary to figure out such tricks for solving problems—that is something that cannot be learned from a book, as Stein emphasized so beautifully in the passage I quoted.³⁰

28. It is instructive in this regard to compare Stein’s remarks about how one can and cannot teach and learn different forms of knowledge in physics with Plato’s discussion, in *Protagoras*, of how one can and cannot teach and learn virtue: Plato concludes (as I read the dialogue) that virtue is learnable but not teachable (at least not based solely on the output and intake of discursive language use); one must practice it in order truly to understand and master it.

29. Of course, not only the learning of theoretical knowledge requires such trust: “When scientists present the results of an experiment they take responsibility for those results by attaching to them the most precious coin of the scientific realm: the individual scientist’s pledge to speak the truth” (Lockman 2005, p. 2).

30. Before making the anecdote public, I emailed Doelman to ask his permission, asking him to read as well the remarks I made surrounding it. He wrote me back a very interesting elaboration on the anecdote, explaining the process by which he acquired the knowledge about the oscillator, which illuminates the theme of this section nicely:

I think this lesson generalizes: how it is possible to learn different pieces or areas of scientific knowledge is one very important criterion for determining the kind of knowledge one is dealing with. This practical form of knowledge, inextricable from the epistemic content of the theory one considers, can not be formalized or even just well captured in book form; most theoretical forms of knowledge can. The latter is, to speak crudely, propositional in a way the former is not. That does not, however, *ipso facto*, imply that it is more important for a proper understanding of the character and content of the sum of knowledge that a theory embodies, that characterizes it as a *physical* theory of *that* sort of family of physical systems. The proponent of the standard view may well reply that it's an amusing anecdote, but tells us nothing of the epistemic content of general relativity. I disagree. It is knowledge of this sort, and the trust we have in experimentalists and instrumentalists of the calibre of Doelman and his collaborators, that gives us the epistemic warrant we have for believing that we understand and comprehend the physical systems that general relativity provides us models of—that justifies our belief that general relativity is a good theory of those systems, and that we know what the mathematical formalism represents, and know in particular which parts of the formalism represent which parts of the world. More importantly, this shows how the two types of knowledge are not cleanly separable—it is only practical knowledge of the kind had by the experimentalists commingled with the formal theoretical apparatus that allows us to understand the formalism as part of a physical theory.³¹

[The anecdote you're relating concerns] a quartz crystal oscillator we had running at Chile for a while that then went on the fritz. After we brought it back to the SAO it seemed to return to normal, and we attributed that to the high elevation and low pressure (up at 16,500 ft). That's not to say that the standard operating procedure should be to send such crystals to a low elevation 'spa' for rejuvenation every so often. It was more just me speculating about what was going on. However the general point you are making is very valid. As with any profession, we develop intuition about the instruments and systems we work with, just as any craftsman understands the feel of good tools working good materials. The lesson in the anecdote you quote is that the quartz oscillator in question was brought up to high altitude at the ALMA site as a reference to check the performance of the hydrogen maser, which is critical for VLBI. When the crystal-maser comparison went South, there was some concern that it might have been the maser that was the problem, but after working with these atomic clocks for so long, I was sanguine. There was a greater probability that it was the crystal (with a power supply not designed specifically for long-term operation at high altitude) that was at fault. When I brought a second crystal to Chile to check it out, that was indeed the case.

Doelman then recommended that I read Lockman (2005), which examines why going to the telescope and being there with the instruments in person can have great benefits, both epistemic and practical, over remote observing. I did read it, and now I recommend that you do as well. Here is a gem from it (p. 3):

Someone who actively plays with the equipment, tries out various combinations of things, and constantly iterates on technique, not only gets data and a sense of its correctness, but also develops skills which can make the next data better. The way to become a skilled observer is to participate in the observations as completely as possible, to seek active control or understanding of every phase of the process; to try to recognize the difference between the basic limitations of an instrument and those limitations which are rooted in style or tradition.

31. I do not want to leave the reader with the idea that I think all theoretical knowledge is articulable in propositions and learnable from books, nor that all practical knowledge can be had only by experience. The situation is far more complex and subtle than that, and a complete discussion would take us too far afield. To get an idea of how the simplistic picture I have sketched falls short, consider all the inexact techniques theoreticians learn and master in order to solve their problems, and how, moreover, they can in general learn to do so only by practice, by trial and error, not from text books—for example, the “feel” for when a given perturbative or approximative technique is the

To make this argument a little more precise, I again follow Stein (1994, p. 635, italics his), now in distinguishing among different aspects of scientific knowledge:

I am construing the word knowledge in a wide and ambiguous sense. The reflections I am proposing have as their object (a) our knowledge in physics as an *achieved result*: knowledge as *the knowledge we have of X*; (b) our knowledge as susceptible of *justification* or *defense*—that is, as involving a structure of “evidence” for its asserted contents; and (c) knowledge—science—as (to appropriate a word of Isaac Levi’s) an *enterprise*: an activity aimed at increasing our knowledge in sense (a), by means appropriate to the constraints of (b).

I would add to this list a fourth, closely related to, but distinct from, Stein’s (b): knowledge as ground for epistemic warrant for other knowledge claims and for furthering the enterprise of science. Call this ‘(d)’. The standard view, I claim, captures at most one of these aspects, knowledge as achieved result, and that at best only in part. The standard view’s picture of the knowledge content of a theory, entirely captured by indicative propositions about the state and nature of the world, does not allow one even to ask how such claims are justified, how their justification depends on the knowledge we have in virtue of other theories, the limitations on the scope of such claims, and so on, for the practical knowledge I have been discussing informs and provides much of the body of the theory’s knowledge in all these aspects, and the standard view does not allow this knowledge to be represented. Thus, it cannot explain the epistemic warrant we have for trusting and using the theory, and thinking that the understanding and comprehension it seems to give us of the world *actually is about the world*.³²

The standard view, however, is deficient, I claim, even with regard to knowledge of aspect (a) (knowledge as achieved result), in so far as it precludes fundamental understanding and comprehension we gain about parts of the world that can be had only by heuristic, informal arguments that cannot be formalized. A good example is the way that one of the fixed relations among quantities in Navier-Stokes theory, *viz.*, that heat flux is always independent of the pressure gradient (Landau and Lifschitz 1975, ch. v, §49), partly encodes the fact that Navier-Stokes theory is valid only for fluids quite close to hydrodynamical and thermodynamical equilibrium of a certain kind. (This relation is not one of the equations of motion, the so-called Navier-Stokes equations; it forms part of what Curiel 2017a calls the kinematical constraints of a theory, which I describe in §8 below.) The only arguments I know of for this are, strictly speaking, incorrect in parts, incoherent as a whole, and even inconsistent in several ways (all with respect to the formalism of the theory).³³ Here is a sample of some of the propositions such proofs rely on:

- that a “particle” of fluid could “switch flow lines” only if the line it traversed crossed another (whereas in fact this can happen also if two lines merely osculate, or, in technical terms, at

right one, how it should be applied, when it is legitimate, when illuminating, where it breaks down, *etc.*

32. *N.b.*: nothing I am saying commits me to any form of realism about theories—I can believe that our best theories tell us a lot about the world, and even that those theories provide us with deep understanding of some aspects of the world, without also believing that those theories are exactly or even approximately true of the world in any deep or strong representational sense.

33. I actually do not know of any published arguments for the proposition. It is a fact widely known among physicists practicing in the relevant areas, part of the common lore of the field, passed down from teacher to student in lectures and such.

- a point at which a Jacobi field of the flow vanishes);
- that flow lines can “cross” in a sense, even in the absence of turbulence, by converging to a single, isolated point, a “singularity” (as happens, more or less, when water runs down a narrow drain);
- that, strictly speaking, however, flow lines can never cross, whether in turbulent flow or smooth, as flow lines lose their definition the moment they converge.

The propositions wear their individual incorrectness and incoherence, and their mutual inconsistency, on their sleeves. Nonetheless, they serve their purpose adequately in practice, allowing the formulation of an argument that sketches a picture of part of the epistemic content of a theory of a genus of physical phenomena (Navier-Stokes fluids), a picture that in essence conveys the important physical features of the phenomena at issue. In particular, such arguments convey the kind of understanding and comprehension Navier-Stokes theory yields of the systems it treats. The nature of the arguments, moreover, are such that they manifestly could not be made or had sense given to them strictly within the confines of the formal apparatus of the theory and any epistemic content that could accrue to it from a clean semantics based on ontology alone. In any event, “close to hydrodynamical and thermodynamical equilibrium” is not a concept that lends itself to precise formalization once and for all, independent of context, but it is of the most fundamental importance in coming to understanding the epistemic content of Navier-Stokes theory that it is essentially an equilibrium theory. No model in the theory represents this fact. No formal, rigorous proposition, can capture this idea once and for all.

Both theoreticians and experimentalists employ that sort of loose, heuristic argument as a matter of course in their workaday activity, as one of the most important tools in their respective tool-boxes. To my mind, the most striking and salient feature of such arguments lies in the fact that, strictly speaking, one cannot prove them either correct or incorrect, in so far as, strictly speaking, several of the terms and relations used in such arguments lose their definition at various points. (The proposition “ A bears the property ϕ ” can be neither true nor false if neither ‘ A ’ nor ‘ ϕ ’ nor ‘bears the property’ has an unambiguous definition.) One cannot analyze the meaning of the propositions of such arguments by analyzing their possible respective truth-conditions in isolation from those of the others, by judging relations of reference to pure ontological models for example. And yet those propositions do have meaning, albeit not precisely determined, and can be more or less appropriate and adequate. One must base the meaning of the propositions on the import of the argument as a whole, and in particular on what we know to be the sort of “mild” incoherencies and inaccuracies we can meaningfully and usefully allow ourselves in contexts like this, as determined and verified by the experimental probing and practical fixing of the theory’s regime of applicability. We are back in the neighborhood of needing to know how to schematize the observer in order to grasp and articulate large and important parts of the theory’s epistemic content.³⁴

34. None of this has to do with meaning holism in anything like Quine’s sense, as expressed, *e.g.*, in the famous (and glib) aphorism, “our statements about the external world face the tribunal of sense experience not individually but only as a corporate body” (Quine 1980, p. 41). The meaning of the—strictly speaking—incoherent and inconsistent propositions in the kinds of arguments I discuss indeed cannot be fixed without reliance on a body of other propositions. Here, however, essentially all required propositions concern either the propriety and adequacy of such descriptions in experimental models of fluids, or the way such propositions can be treated as approximations, as it

I have argued that the standard view gives too little as an answer to the question of the scope of a theory's epistemic content. In another sense, however, the standard view gives too much—it says that every question one can pose to the theory has a determinate answer. In a purely formal sense, this may be correct. But if one wants to know whether the answer given is *relevant to the real world*—whether the answer counts as scientific knowledge in the senses, for instance, of Stein's (b) and (c), and my (d), and indeed even for the factual kind in Stein's (a)—then one has to know, among other things, the scope of the theory's regime of applicability, *i.e.*, the spatial, temporal and energetic scales (*e.g.*) in which it is a good theory and those where it breaks down. For that, one must have recourse to knowledge that one can come by only through experimentation and the application of knowledge embodied by other theories, and thus, more importantly, knowledge that is not and cannot be encoded in the formalism of the theory itself.

In so far as knowledge of the regime essentially represents knowledge in Stein's aspect (b), moreover, as being susceptible to justification by involvement in a network of evidentiary relations, it cannot be given as a physical interpretive postulate in any standard sense (*e.g.*, a Reichenbachian coordinating principle, a Carnapian reduction sentence, or contemporary gestures at representational capacities), *i.e.*, one that encodes knowledge cleanly separable from any knowledge represented purely by the theory's formalism.³⁵ Our knowledge of where and how Navier-Stokes theory, *e.g.*, breaks down is in large part based on and justified by experimental results as represented in and put into meaningful contact with the mathematical formalism. A physical interpretive postulate, therefore, that encoded such knowledge would have to make inevitable reference to such facts that are known only by experimentation, and thus require knowledge of how the observer and measuring devices are represented in the theory. Whether a given question, therefore, that seems in the abstract to be one the theory can address is in fact one the theory can give a physically relevant answer to—whether it is about a system that the theory appropriately and adequately treats—can be determined only by recourse to knowledge that goes well beyond that which can be encoded only in the formalism or that can be expressed by propositions making no reference to detailed experimental knowledge. It is knowledge, in other words, that cannot be cleanly split into a formal part and a practical part. The standard picture cannot apply.³⁶

Much of this discussion has implicitly relied on a picture of knowledge that goes far beyond what can be captured in propositions expressing true, justified beliefs. The most important component of the epistemic content of a theory, on my view, is more akin to what the (notoriously problematic) concepts of “comprehension” and “understanding” try to get at. I will not attempt to give a

were, of rigorous and precise theoretical propositions. There is no implication that *all* the propositions of theory face the tribunal of experience as a corporate body, much less that such propositions depend *at all* on all our statements about the external world. Indeed, it is difficult even to make sense of Quinean holism on my view. As the reach of experimental prowess extends itself, and ever more physical systems in ever more states fall with the regime of applicability of a theory, the meaning of such terms as “pressure” (in Navier-Stokes theory) does not change. If one is committed to a role for a relation of designation in one's semantics, then one will certainly conclude that the reference of the term changes. I deny, however, that meaning is exhausted by, much less grounded in, referential relations (as I argue in §9 below, and at much greater length in Curiel 2017b).

35. This, I think, is what is right about the claims of Friedman (2001) to the effect that, in so far as the relativized *a priori* provides the grounds for endowing theoretical formalism with empirical content, it itself is not susceptible to justification by experimental means.

36. I shall give in §§8–9 a more detailed discussion of a few particularly damning ways in which the standard view allows one to ask too many questions and formulate too many propositions.

systematic account of these ideas, but I can say a few things to try to convey my sense of them. “Understanding” is something like “the capacity to operate successfully in the scientific enterprise, in all its aspects and parts”: to use our representations and instruments as the basis for the fruitful continuation of the enterprise, as part of evidential warrant in testing, as basis for predictions and characterizations, as inspiration for potentially fruitful new investigations, as the grounds for conceptual clarification and innovation in foundational work; and perhaps most of all to grasp how our representations and our instruments relate to, inform, and substantively contribute to the constitution of each other, and to grasp that in such a way that it grounds the work successfully continuing the enterprise.

Consider, as an example, the two formulations of the Second Law of thermodynamics given by, respectively, the Kelvin Postulate and the Clausius Postulate.³⁷

Postulate 6.1 (Lord Kelvin) *A transformation whose only final result is to transform into work heat extracted from a source that is at the same temperature throughout is impossible.*

Postulate 6.2 (Clausius) *A transformation whose only final result is to transfer heat from a body at a given temperature to a body of a higher temperature is impossible.*

One knows the Second Law in an important sense when one learns each of these two postulates as facts, and comes to see that and how they are equivalent. One *understands* the Second Law when one masters the facts, not just their objective, individual content, but their inter-relations, and, more, how to fruitfully deploy them—that is what it is to *master* in this context. Such mastery cannot be captured in any linguistic expression. This stands in opposition to mere knowledge of the Second Law, which may consist only of the possession of a linguistic expression—the ability to recall and state it, perhaps to use it in some form of deductive reasoning that does not extend epistemic content in any substantive way, but to be still unable to use it fruitfully in a truly ampliative context: the possibility of creating or discovering new knowledge. That is understanding in the deepest and richest sense.

Before moving on, I want to address a question that my discussion has raised. One may well object that the kind of knowledge shown by, say, astronomers in good manipulation of their instruments may be part of the epistemic content of *some* theory, but surely not that of general relativity. That may well be right in some sense, although I suspect that one can never impose a clean, sharp demarcation, once and for all, between a given theory and auxiliary theories needed by experimentalists in the study and use of it. What, then, can we say, if anything at all, about the criteria for when a particular piece of knowledge has been delivered by a particular theory in the relevant sense, so as to count as part of *its* epistemic content, rather than, say, as part of the content of an auxiliary theory used to model an experimental apparatus? I don’t know. At a minimum, it should be a piece of knowledge that we could not, even in principle, have without using the theory at issue, at least as determined by the context of our actual current state of achieved knowledge. In principle, one should always be able to perform experiments of different kinds, with instruments modeled by different theories, so as to determine whether or not the theory at issue makes an ineliminable contribution to the epistemic content of the knowledge, and so to attempt to isolate that contribution. In the end, as Geroch so piquantly put it in the long quote at the end

37. My treatment of this example was inspired by the discussion of Maxwell (1888, ch. VIII, p. 153).

of §4 (page 4), a theory is a nebulous mass. I doubt very much that there can be a fixed, final, determinate answer to the question whether this or that bit of knowledge properly belongs to its epistemic content, an answer relevant and appropriate for all philosophical and scientific purposes.

In the end, all the problems I discussed in this section arise from the same deficiency in the standard view, to wit, its inability to allow schematization of the observer in a way relevant to the articulation of the epistemic content of a theory.

7 How Theory and Experiment Make Contact with and Inform Each Other

That we must have the capacity in our theories to construct explicit models of complete experimental situations including instruments and the actual methods of their deployment in order to represent actual observations, however, immediately raises a serious problem, one that Stein (1992, p. 290) trenchantly poses:

... we have no language at all in which there are well-defined logical relations between a theoretical part that incorporates fundamental physics and any observational part at all—no framework for physics that includes observational terms, whether theory-laden or not... I cannot think of any case in which one can honestly deduce what might honestly be called an observation. What can be done, rather, is to represent ... “schematically,” within the mathematical structure of a theoretically characterized situation, the position of a “schematic observer,” and infer something about the observations such an observer would have.

In other words, we do not have a formal account of the epistemic content of the theories of physics even minimally adequate for any account of their actual empirical application; this is not to say that such applications in real scientific practice have no foundation or are unjustified, only that we have no adequate comprehension of the process. Forget how we get the theory into or out of the laboratory—how do we get the laboratory into the theory? This, I think, is the fundamental issue one must address in trying to give an account of the epistemic content of scientific theories and how its components bear on each other, and in particular in answering the question whether they are mutually extricable in the way demanded by the standard story told by philosophers about the character of scientific theories, and relied upon by them in their work. (See again the long passage from Stein 1994 that I quote on page 5.)

In order to address this problem, in this section I explain with a little more precision (but not much) what I mean by ‘theory’, with the goal of drawing a—schematic!—picture of how experimental results are placed into physically significant contact with formal theoretical structures, and what the nature of that contact is. I neither promise nor threaten you with anything near a precise, complete account of any of this. For my purposes here, it suffices to provide only a rough characterization sketched in broad strokes.

A theory, as implied by the discussion up to this point, is a system that allows one, among other things, to formulate propositions and affirm them in principled ways based on evidence gathered according to good principles, to apply them in turn as evidence for other propositions, and to use

them as the inspiration and basis for new investigations.³⁸ A theory in this sense is a system that allows for the unified representation and modeling of a particular kind of physical system so as to render that kind amenable to investigation by scientific reasoning and practices of all forms. A particular kind of physical system is one such that all individuals falling under the kind bear the same physical quantities whose properties are characterized by and whose behavior is governed by the same set of equations of motion and collateral mathematical relations. A theory in my sense also includes such things as accounts of experimental devices appropriate for the probing of the relevant kind of physical system, good practices for employing them, sound techniques for the collection of raw data and statistical and other analysis and organization of the same, reliable methods of approximative and heuristic reasoning for constructing models and solving equations, and guidelines for determining whether a system of the given kind is in a state and experiencing interactions with its environment such as to be jointly amenable to appropriate and adequate representation by the theory (*viz.*, whether or not the system falls into the theory’s regime of applicability), and so on.

Before diving into the details of how I think that all these components hang together so to make it possible for the theoretical and the experimental parts of a theory to make fruitful contact with each other, my analysis and arguments will need further distinctions with regard to “level” of theoretical representation, so it will be useful to characterize them by laying down terminology. I will not attempt to give formal, precise definitions. These ideas are for my purposes adequately clarified by examples.

Theories are often formulated in the context of a framework. Newtonian mechanics, the heart of which is embodied in the definitions and the three laws Newton lays out in his *Principia*, is perhaps the exemplar *nonpareil* of a framework. Newton’s Second Law in the abstract is not the equation of motion of any particular kind of physical system. It is rather the template that any equation of motion for any particular kind of system treated by the framework must instantiate. Newton’s theory of gravity is a theory formulated in the framework of Newtonian mechanics. It treats that kind of physical system characterized, *inter alia*, by the possession of inertial mass, of gravitational mass (quantitatively equal to the inertial mass), of spatial position, and of a velocity expressible as the temporal derivative of spatial position, such that its dynamical evolution is governed by Newton’s gravitational force law. Navier-Stokes theory is another theory formulated in the framework of Newtonian mechanics. It treats that kind of physical system characterized, *inter alia*, by the possession of inertial mass, shear viscosity, bulk viscosity, thermal conductivity, fluid velocity, heat flux, pressure, and shear-stress, all satisfying among themselves fixed relations of constraint (*e.g.*, that heat flux is always independent of the pressure gradient), and whose

38. My conception of a theory is in many ways similar to, and indeed in part inspired by, Carnap’s (1956) conception of a linguistic framework, particularly in the way that a theory in my sense serves to define a fixed sense of physical possibility relevant to the kinds of system the theory treats, just as a Carnapian framework does for the types of entity whose existence is analytic in the framework. Carnap’s conception is too broad and vague, however, to do the work I require of it. Stein (1992) provides an insightful and illuminating, albeit brief, discussion of the differences between a Carnapian framework in the original sense and a theory in the sense of a formal structure in theoretical physics related to, but more restricted in scope than, the type I am sketching here. Lakatos (1970) has some affinity with the gist of this view, in particular his notion of the “hard core” of “research programs”, though, again, the differences in detail outweigh the similarities. There is perhaps more affinity with the “research traditions” of Laudan (1977), in so far as different ones can share and swap important methodological and theoretical principles, as can happen with theories in my sense. A discussion of these comparisons is beyond the scope of this paper.

dynamical evolution is governed by the Navier-Stokes equations.

The template in a framework for equations of motion and other mathematical relations I call *abstract*. Canonical examples are Newton's Second Law, the Schrödinger equation in non-relativistic quantum mechanics, and so on. Structure and entities at the highest level of a theory formulated in a given framework I will call *generic*. In particular, generic structure has no definite values for those quantities that appear as constants in the theory's equations of motion and other mathematical relations. The symbol ' k ' appearing in the generic equation of motion of an elastic spring modeled as a simple harmonic oscillator, $\ddot{x} = -\frac{k}{m}x$, denotes Hooke's constant (the coefficient of proportionality of a force applied to the spring and the resulting displacement from its equilibrium position), but possesses no fixed value, and the same for the mass m . It is important to keep in mind, however, that all these formal representations of physical quantities at the generic level in a theory do have determinate physical dimensions, for Hooke's constant, *e.g.*, $\frac{\text{m}}{\text{t}^2}$. Otherwise, one could not say that this is the generic equation of motion for, say, a spring rather than a pendulum or oscillating string or electric circuit or any other type of system whose dynamical evolution is governed by the equation of motion for a simple harmonic oscillator.

One can, in the same way, write down generic solutions to the generic equations of motion; these are formal representations of the dynamically possible evolutions such systems can manifest. One generic solution to the equation of motion of an elastic spring is $x(t) = \cos(\sqrt{\frac{k}{m}}t)$, where one may think of ' k ' and ' m ' as dummy variables, not determinate real numbers. Generic structure defines a *genus* of physical system, all those types of physical system the theory appropriately and adequately treats.

One obtains *specific* structure by fixing the values of all such constants in generic structure, say $m = 1$ and $k = 5$ (in some system of units) for the elastic spring. This defines a *species* of physical system of that genus, all springs with those values for mass and Hooke's constant.³⁹ One now has a determinate space of states for systems of that species, and a determinate family of dynamically possible evolutions, *viz.*, the solutions to the specific equations of motion, represented by a distinguished family of paths on the space of states. A path in this family is an *individual model* (or *individual solution*) of the specific equations of motion; one common way of fixing an individual model is by fixing definite initial conditions for the specific equations of motion. An individual model, as the name suggests, represents a unique physical system of the species, that whose dynamical quantities satisfy the initial conditions.⁴⁰ In the case of the spring, that may mean the unique one whose position and momentum at a given time have the values given by the initial conditions.⁴¹

39. One can as well consider mixed systems, with, say, a fixed value for mass but indeterminate value for Hooke's constant. These raise interesting questions, but they are beside the point here, so I ignore them.

40. Everything I say so far applies only to deterministic theories that represent the dynamical behavior of its target systems with something like partial-differential equations. The discussion is easily modified to accommodate, *e.g.*, stochastic theories, and so on.

41. I have deliberately taken my terminology from biological taxonomy, inspired by the remark of Peirce (1878, p. 143):

Now, the naturalists are the great builders of conceptions; there is no other branch of science where so much of this work is done as in theirs; and we must, to great measure, take them for our teachers in this important part of logic.

There is much of insightful relevance in the lead-up to this remark, about how one individuates and characterizes genera and species of physical systems in my sense, which it would be illuminating to discuss, but it would take us

Finally, I call a *concrete model* a collection of experimentally or observationally gathered results structured and interpreted in such a way as to allow identification with an individual model.⁴² (I will sometimes also refer to a concrete model as *structured data*.) If I continually measure the position and momentum of a spring with known values for m and k , oscillating in one dimension for an interval of time equal to its period of oscillation, and graph the results “in a natural way”, *e.g.*, as a curve parametrized by time on a Cartesian plane whose x -axis represents position and y -axis represents momentum, then I will produce a curve that can be, “in a natural way”, identified with exactly one dynamically possible evolution on the space of states used to represent that species of spring. Kepler’s organization and structuring of the planetary ephemerides into parametrized ellipses satisfying the Area Law and the Harmonic Law is a marvelous example from the actual history of physics of this procedure.⁴³

It is this identification of individual and concrete models that embodies the substantive contact between theory and experiment required to comprehend the full epistemic content of a theory.⁴⁴ There are many fascinating and deep problems associated with characterizing the ways that experimental and observational results are collected and appropriately transformed into structured data, but I put them aside.⁴⁵ Here, I will give only a (schematic!) outline of how I envisage the procedure of bringing theory and experiment into contact, based on the concepts I just articulated. I will give a more finely grained analysis of the procedure in the next section.

too far afield. I am tempted to describe structure at the level of a framework as *phylar*, and to call the family of all physical systems treated by a framework a *phylum*—and so all physical systems would fall under the kingdom of physics—but I suspect it would just be distracting to the reader.

42. This idea bears obvious and interesting comparison with the distinction between data and phenomena as drawn by Bogen and Woodward (1988). In so far as I understand their distinction, their idea of data more or less corresponds with my idea of “experimentally or observationally gathered results”, but their notion of phenomena, in so far as it seems to try to capture something like general patterns in the world, does not neatly square with my conception of a concrete model, which is the result of structuring the results of a single experiment (or family of related experiments). A complete discussion is beyond the scope of this paper.

43. I am making a serious over-simplification here. Not every concrete model contains data for a sufficient number of dynamic quantities to enable the identification of the concrete model with a single, unique individual model. A set of data may contain, *e.g.*, data only for the pressure of a Navier-Stokes fluid, which allows one to identify the concrete model only with a family of individual models, all those with those values for pressure, but having any values for all the other dynamic quantities Navier-Stokes attributes to fluids. I also ignore the fact that, in light of the inevitability of expected error and experimental imprecision, one can *always* identify a given concrete model with a family of individual models, *viz.*, all those whose values collectively lie within the error and imprecision intervals associated with the data. I will not discuss these issues, as they raise too many problems orthogonal to my purposes here, though the problems are of great interest in their own right. I assume from hereon that a concrete model is such as to allow identification with a single, unique individual model. In particular, I assume it contains values for enough dynamic quantities to determine individual states—at least as many as the “dimension” of the state space—and that determinate, individual values have been fixed for all data, not just intervals defined by the range of expected error and imprecisions.

44. This classification of levels of theory and data and the concomitant idea of identifying individual and concrete models have obvious similarities with the proposals of Suppes (1960; 1962), and indeed much of my own thought has been inspired by reactions to Suppes’ extensive work on these questions (not just those two papers).

45. As is so often the case, Suppes (1960, p. 297) sums the matter up with elegance and concision:

The maddeningly diverse and complex experience which constitutes an experiment is not the entity which is directly compared with a model of a theory. Drastic assumptions of all sorts are made in reducing the experimental experience, as I shall term it, to a simple entity ready for comparison with a model of the theory.

The identification—I think essentially always—comes down in the end to the brute comparison, one by one, of the values of theoretically calculated quantities on the one side with the values of the analogues of those quantities in the structured data on the other, with the choice being made each time whether the values are close enough, given one’s understanding of the errors and imprecisions in measurement, the introduction of skew in statistical reconstruction, one’s pragmatically chosen error tolerance, and other such factors. An excellent example is provided by the comparison of the observed to the calculated values for the planetary orbital periods and the semi-major axes of the orbits in Book III of the *Principia*, in Newton’s claim that the orbits of the planets and of Jupiter’s and Saturn’s satellites all respectively satisfy Kepler’s Harmonic Law. His explicit discussion of the errors involved in constructing the concrete models, primarily due to limitations in the observational instruments, and the errors involved in constructing the individual models, primarily due to idealizations, and why they can all safely be ignored when comparing the data to the theoretically calculated results, exemplifies perfectly what I am trying to get at. (See especially Newton’s discussions of Phenomenon I, II, and IV, and Harper 2011 for a comprehensive and detailed discussion of Newton’s analysis of the data.)

Newton’s discussion also exemplifies perhaps the most important feature of such identification, at least for my purposes, *viz.*, the fundamentally pragmatic character of the whole affair. One does not evaluate the truth value of the proposition “all the values are in close enough agreement”. One rather chooses to accept (or not) the theoretical structure as an appropriate and adequate representation of the structured data, and one does this for any of a number of reasons, one of which may be that the differences in values one focuses on fall within one’s pragmatically chosen error tolerance. Even when that does not happen, there may be other principled grounds on which one can reasonably choose to accept the identification of theoretical structure with concrete model, *e.g.*, to serve as the ground of further, more refined investigation. This is what Newton did in his studies of the moon’s librations and other irregularities in its motion—he knew his initial, simple theoretical models (2-body interaction ignoring the Sun, circular orbit) were inadequate (in the sense of predictive accuracy) by any reasonable measure, but he relied on them anyway in attempting to construct better ones, taking into account ever more finely grained details of the system (*Principia*, Book I, Proposition LXVI, and Book III, Propositions XVII, XXII, and XXV–XXXV; see also Smith 1999a, 1999b)—and he had *damn* good reason to do so.

With regard to the exact relation between theory and experiment in making such identifications, it is worth noting that one can “Keplerize” the planetary ephemerides without Newton’s Laws or any of the other most characteristic theoretical concepts of his system, such as inertial mass and force. One needs only the pre-theoretical ideas of “ellipse”, “area swept out in given time”, “distance to sun”, “orbital period”, and so on. One constructs the structured data using these pre-theoretical concepts, and then *shows* that it can be *interpreted* using the concepts of the Newtonian framework in such a way as to support identification of the concrete model with a particular individual model. And indeed that is what Newton did.

To try to reassure the reader that I am not cherry-picking particularly apt examples from the Greatest Hits in the History of Science, I will also briefly discuss the recent detection of gravitational waves by LIGO (Abbott, B. *et al.* (LIGO Scientific Collaboration and Virgo Collaboration) 2016). This example also shows in a particularly clear light exactly where theory enters into the design of the experiment and the construction of structured data, and where it does not. I give only a

wildly simplified, schematic description of the way observations are transformed into structured data and compared with theory in the experiment.⁴⁶

1. Design the instrument and their relative configuration so as to be susceptible kinematically to coupling with the kinds of physical phenomena one is interested in, according to the models of one's theories—in this case, stress-strains caused by “distortion” of spacetime as gravitational waves pass, in the persona of the relative acceleration of contiguous portions of continuous bodies (the instruments).
2. Arrange the instruments in the desired configuration, fire them up, and begin collecting raw data.
3. In the first pass through raw data, discriminate between “noise” and “signal”, throwing out “bad measurements”, winnowing signal pollution from other sources, ignoring mistakes and instrument malfunctions, and so on. (This will have to be done at several points later in the process as well.)
4. Match and collate different numerical “data streams” of raw stress-strain (in this case, time-stamps with measured stress-strain coming from instruments in geographically separated locations).
5. Perform numerical interpolation, extrapolation, and manipulation (best-fit statistical analysis, *etc.*) on the raw data so as to produce a continuous curve of aggregated stress-strain plotted against time; it is now structured data.
6. Using the theory of general relativity, construct models of the production of gravitational waves in different phases of the process (in-spiral of a binary black hole system, coalescence, ring-down of the resultant black hole to equilibrium); calculate the waveforms as they would appear when passing through the Earth.
7. Based on those waveforms, construct continuous graphs representing the responses of the instruments to the stress-strain resulting from them.
8. Compare the structured-data graph with the theoretically constructed graph.
9. Decide, based on calculation of expected errors, statistical analysis of the rate of expected false positives, and so on, whether the structured-data graph matches the theoretical graph to within one's error tolerance.
10. If the match is good enough, use the physical principles of general relativity to interpret the measured stress-strain in the instruments as the result of passing gravitational waves produced in the coalescence of a binary black hole system.
11. Identify the structured data with an individual model (or family of models) of a coalescing binary black hole system.

46. I completely ignore many subtleties, such as the fact that radically different kinds of theoretical models and calculational methods are used to model different parts of the single process of in-spiral, coalescence and ring-down, and also the fact that the stress-strain is determined by the use of Michelson interferometers, and so optical theories are required as well as mechanical theories in the design and performance of the experiment and the structuring of the data. There are many fascinating philosophical problems of methodology and epistemology here. See, for instance, the project being run by Lydia Patton on the problems associated with parameter estimation in LIGO (<https://www.researchgate.net/project/LIGO-and-parameter-estimation>).

One of the most interesting points for our purposes is where general relativity itself enters: only in steps 6, 7, 10 and 11. Step 1 may appear already to depend on general relativity, but the appearance is illusory, arising from the description I used. The instrument needs only to be sensitive to stress-strain induced by “gravitational forces”, irrespective of how one interprets, conceives or represents such a thing. In particular, one needs to know only how to measure the relative acceleration of contiguous parts of continuous bodies (induced stress-strain in the instrument), and to ensure that the stress-strain does not arise from any source other than “gravitational” (electromagnetic forces, *e.g.*, or mechanical disturbances arising from vibrations in the Earth’s surface caused by passing semi trucks).⁴⁷ None of that depends on the assumption (or lack thereof) that the Einstein field equation holds for the matter fields one is working with. The theory’s dynamics (the Einstein field equation) plays a role only in steps 6 and 7, manifestly and directly so.⁴⁸ The way that general relativity comes into play, moreover, does not make the experiment “theory-laden” in any interesting sense, because it is only the pre-theoretical concepts of stress-strain, relative acceleration, and so on, that ground the construction of the instruments and the structuring and analysis of the data that will be used to test the predictions.⁴⁹ The Einstein field equation comes into play only in producing predictions to compare to observed results, and then in interpreting the results using its conceptual framework by the identification of a structured data model with an individual model (or family of models) in the theory.⁵⁰

Stein (1994, p. 639, emphasizes his) speaks of “*a mathematical structure [of a theory] discernible in the world of phenomena, of observations, of experience.*” I endorse this idea by the requirement that structured data be identifiable with formal theoretical structures. I mean, however, something quite weaker than what many contemporary philosophers would assume on reading Stein’s phrase in the absence of any context. I mean only that the theory facilitates the successful identification of individual and concrete models in such a way as to make it possible for its entire formal and conceptual apparatus to come into play and form the basis for further successful reasoning about the systems at issue. In particular, I make no claims that structures in the individual models and in the concrete models, much less “in the world itself”, are “isomorphic” or “homomorphic” or “similar” in any way above and beyond the fact that they are relevantly identifiable with each other.⁵¹ (I do not even know what it means to speak of an isomorphism or homomorphism or

47. *Cognoscenti* should note that the relative acceleration between neighboring curves at issue here is not an instance of geodesic deviation, as the instruments (and their parts) are not following geodesics: they are bolted to the floor of the laboratory, which is not in free fall.

48. It may also come into play in step 9, in the determination of expected errors, but that discussion is far beyond the scope of this paper.

49. These concepts are pre-theoretic with respect to general relativity, obviously not with respect to the theories (*e.g.*, of the continuum mechanics of solid bodies) that are used to construct detailed models of the instruments themselves and their reactions to imposed forces.

50. At this point, the experiment has not definitively ruled out the possibility that other theories of space, time and gravity besides general relativity might also produce predictions appropriately and adequately matching the structured data. This fact highlights the important point that different theories—*e.g.*, general relativity and many alternative theories of gravity—can share a substantial amount of epistemic content, *viz.*, that encoded in the structured data formed from observations that can be modeled by and can be used to test theoretical propositions in the different theories. It should thus be clear that I reject the “incommensurability” of theories in any deep or interesting sense.

51. See the discussion in the final section of Curiel (2014) for examples of how one can identify the same concrete model with individual models in different theories (the Lagrangian and Hamiltonian formulations, respectively, of

similarity between a mathematical structure and “the world itself”.⁵²)

Indeed, the sense in which there may be any structure in the concrete model isomorphic with structure in the individual model is etiolate at best. In the case of LIGO, the extrapolated continuous curve of stress-strain is isomorphic with the analogous curve in the individual model general relativity yields, but that is only the minimal consistency condition required in order to justify the identification of the two. The additional structures in general relativity one needs in order not only to justify the identification but further to give empirical content to the individual model—additional structure such as the relation of the relative acceleration of neighboring non-geodetic curves to the Riemann tensor, *e.g.*—have no correlate, isomorphic or otherwise, in the concrete model.

In any event, the identification is, strictly speaking, not a relation between “different structures” at all, in the sense of being representable by a proposition whose truth value one can evaluate based only on consideration of theory and data alone—it is, as I already have stated, a pragmatic choice: does one accept that the two can be identified with each other or not? It is a pragmatic choice in the same way as use of a framework in the sense of Carnap (1956) is decided by pragmatic choice. This is why pragmatics and epistemological concerns, not ontology, lie at the ground of semantics, as I will argue in §9 below. Nonetheless, I will continue to speak of this identification as a “representation relation” for convenience, though I emphasize again that by “representation” I do not mean a relation whose content can be exhausted by reference to the theory and data alone.

One can agree with this analysis of how theory and experiment make contact with each other while remaining agnostic about all issues pertaining to realism and anti-realism. Indeed, I want to emphasize that none of this has anything to do with any issue pertaining to structural realism in particular, nor to debates about realism and anti-realism in general. There is no claim made or needed that the structure manifest in the structured data is “really” part of the furniture of the world, in some deep metaphysical sense. That it is manifest in the structured data suffices for the soundness of the representational relations at issue. Those relations are agnostic about realism, as any good epistemic relation should be.

Where in all this, however, does the schematic representation of the observer enter? In the case of Newton’s investigations, they appear explicitly in his defense of the way he identifies his theory’s models with the Keplerized orbits, by explaining the way, *e.g.*, one expects variation in the observed size of Jupiter’s apparent diameter, and thus in the distance its satellites appear

the theory of an oscillating spring) without any physically significant morphism between any of the models at issue.

52. Jeremy Butterfield countered this claim in conversation by adverting to work on the correspondence theory of truth, such as Davidson (1969), in which there seems to be a straightforward sense of isomorphism between a formal semantic structure and parts of the world. I can accept the cogency of such work but still deny that the isomorphism at issue consists of anything more than the kind of bare, thin agreement between formal semantic structures and parts of the world captured by the identification of concrete models and individual models in my sense. For the claim that there is a cogent sense of isomorphism between a mathematical structure and the world to be true in a substantive and relevantly interesting sense, I would want to see an example like the following: there is an isomorphism between the (standard) topology of \mathbb{R}^4 , as used in flat FLRW cosmological models in general relativity, and the topology of the real world. I do not know, however, what “the topology of the real world” means. I can certainly envisage ways to operationalize its meaning, but different operationalizations, each with manifest physical content, will yield different topologies. One obvious way will yield the standard manifold topology (at least locally). Another, however, based on time functions, will yield a different topology (Geroch 1970). So which one is “the” topology of “the world itself”?

to be from its center, when using telescopes of different sizes and resolving powers (Book III, Phenomenon 1). In the case of LIGO, they appear in several places, perhaps most notably for our purposes in the calculation of the response of the instrument to applied stress-strain from passing gravitational waves. It is true that often, if not indeed almost always, explicit representations of the observer and the instruments will not appear in the concrete models—they get washed out in the structuring of the data. They are still there implicitly in the background, however, and the eventual identification of individual and concrete model can happen only because of them.

At a deeper level of analysis, schematic representation of the observer is what grounds the circumscription of the regime of applicability of a theory, because different kinds of instruments, naively modeled and deployed, give different outcomes for “measuring the same thing” near the boundaries of the regime. Circumscribing the regime, however, grounds all other aspects of what one may want to call the semantics of a theory in particular and its epistemic content more generally, so it is to this problem that I now turn.

8 The Breakdown Scale and the Regimes of a Theory

This shaking keeps me steady. I should know.
What falls away is always. And is near.
I wake to sleep, and take my waking slow.
I learn by going where I have to go.

– Theodore Roethke
“The Waking”

As I discussed in §4, a necessary condition for a collection of (possible) physical systems to hang together in the right way, to share a common structure, to form a unified kind, in the sense of being adequately treated by a single theory, is that they all share a family of common breakdown scales: “places” where all the kinematically and dynamically relevant structures of the theory break down all at once, in the sense that the theory becomes inadequate for an appropriate treatment of any of the systems in the collection.⁵³ These places are not characterized by anything intrinsic to the formalism of the theory, but are rather almost always determined by knowledge that one can gather only from investigations grounded in that part of the epistemic content of the theory not captured by the formalism by itself.

Although we perhaps naively tend to think of scales determined by spatial, temporal and energetic quantities when considering how and where theories break down in their capacity to provide sound representations of phenomena, any quantity in any theory can provide such a measure.

53.

Und wir: Zuschauer, immer, überall,
dem allen zugewandt und nie hinaus!
Uns überfüllts. Wir ordnens. Es zerfällt.
Wir ordnens wieder und zerfallen selbst.

– Rainer Maria Rilke
“Eighth Duino Elegy”

Velocity provides a breakdown scale for Newtonian mechanics, and acceleration and scalar curvature provide different breakdown scales for various theories of gravity, such as general relativity. No breakdown scale, moreover, can be a single number holding for all systems the theory treats. Navier-Stokes theory, for instance, becomes inadequate for different fluids at different energies and spatial and temporal scales. Often it is not a bound on a single quantity, such as a value of energy, a value of spatial length, *etc.*: classical Maxwell theory, *e.g.*, breaks down when the ratio of the field’s amplitude to its frequency approaches \hbar . Nor is it ever the case that there is a single characteristic scale for each theory. Navier-Stokes theory breaks down when various measures of flow complexity indicate the fluid is approaching turbulence, when the fluid is too viscous, when one tries to use the theory to specify behavior at time scales comparable to equilibration time after a sharp disturbance, when temperatures become large enough that heat loss due to emission of black-body radiation becomes non-negligible, when the ambient electromagnetic field in the environment becomes strong enough that the constituent molecules of the fluid begin to denature and ionize, and so on. A breakdown scale, then, is something like the following: a measure of or function of or relation among quantities, such that, when the joint state of the system and its environment imply that the values of some of the system’s quantities do not satisfy the measure, function or relation, then the theory can no longer provide good models of the system.⁵⁴ (I provide a more precise definition below.)

Breakdown scales can essentially never be determined by analysis of the formalism and theoretical machinery of the theory alone, without input from knowledge acquired by experimentation in particular and empirical investigation in general. As such, they change with the increasing scope and depth of our experimental reach.⁵⁵ Every theory has many breakdown scales, putting constraints on the values of all the quantities the theory treats. For simplicity, however, I will speak as though all such scales are spatial or temporal.

What does it mean to say that the theory cannot provide good models of systems outside its breakdown scales? One of the most important markers of this is that the quantities the theory attributes to the system lose unambiguous definition. For a Navier-Stokes fluid, for example, different sorts of thermometers that allow spatial discrimination on scales only a couple of orders of magnitude greater than the mean free path of the fluid’s molecules will record markedly different “temperatures” depending on characteristics of the joint system that one can safely ignore at larger scales—the transparency of each thermometric system to the fluid’s particles, for instance, and the

54. I gloss over a subtlety here: sometimes approximations used to construct models of particular phenomena, such as surface waves in fluid dynamics, have characteristic breakdown scales different from those of the material in which the phenomena manifest (Lamb 1932, ch. ix). The theory can provide appropriate and adequate models of the systems in the relevant states, only not in the way the approximations definitive of the phenomena at issue require.

55. This is not to say that one can never find relations among quantities that determine some part of the breakdown regime such that the relations are insensitive to the state of experimental prowess. Geroch (2001, pp. 2–4) points out, for example, that any relativistic theory of dissipative fluids—such as Navier-Stokes theory posed in the relativistic context—must fail on every combined time-distance scale τ - d such that $d^2 < \eta\tau$ and $d > c\tau$, where η is a value of a typical Navier-Stokes dissipation coefficient and c is the speed of light. Instead of a characteristic break-down scale, this requirement defines a characteristic break-down *area* in the τ - d -plane. Note that the complement of this region in the plane, that is, the region in which the system remains valid (at least so far as these conditions are concerned), includes arbitrarily small d -values and arbitrarily small τ -values (though not both at the same time). This all follows strictly from considerations of relativistic constraints on causal propagation alone. Such relations, however, never exhaust the boundary of the complete breakdown scale, and in general form only a small part it.

exact distribution and pattern of collisions of the particles with the thermometer's molecules. As the fluid approaches turbulence, the values of all the its quantities begin to vary rapidly in time and eventually cannot be measured by any conventional means. In practice, all that can be measured are averages (usually time averages), leading to the application of the idea of ensemble averages and the development of statistical generalizations of the Navier-Stokes equations to treat the fluid flow (Foias et al. 2001). As I discussed in §4 above, one of the most striking features of breakdown scales is that all the quantities a theory attributes to the systems it treats lose definition at the same time when the system passes beyond one of them. This makes the idea unambiguous.

In order to draw out the full import of the existence of breakdown scales for the characterization of the epistemic content of a theory, a few more definitions and distinctions must be made.⁵⁶ The most basic distinction is between what I call the kinematical and the dynamical parts of a theory. The kinematical parts, roughly speaking, are all those features of a theory's treatment of a system that remain the same irrespective of the state the system is in and irrespective of the kinds of interaction it has with its environment. The kinematic quantities are, therefore, those which are assumed to be constant. In relativity theory, this includes the speed of light *in vacuo*. In Newtonian particle mechanics, the masses of the particles are assumed to be fixed once and for all. In Navier-Stokes theory, the fluid's shear and bulk viscosities and its coefficient of thermoconductivity are assumed to be constant. The dynamic quantities are those that are allowed to vary as the system evolves, such as position, velocity, momentum, shear-stress, and so on.

All theories, moreover, attribute not only constant quantities to its systems, but also fixed mathematical relations among both its kinematic and dynamic quantities whose form is always the same, what I call its kinematical constraints. In Newtonian mechanics, the velocity is always the time-derivative of position. (This is not so trivial a requirement as it may sound, as there are theories in which it fails, *e.g.*, in some formulations of relativistic fluid mechanics; see Landau and Lifschitz 1975 and Earman 1978). Navier-Stokes theory requires that the shear-stress tensor be symmetric, and that heat flux be independent of the pressure gradient. The attribution of definite values for the kinematic quantities to a species of system also is a kinematical constraint, *e.g.*, that water under "normal" conditions have the value (approximately) 1000 $\mu\text{Pa}\cdot\text{s}$ (micro Pascal-seconds) for its shear viscosity.

Every theory postulates as well dynamical relations among its quantities, the more familiar equations of motion or field equations. These have the characteristic property that their specific form depends on the interactions the system has with its environment. (\mathbf{F} always equals $m\mathbf{a}$ in Newtonian mechanics, but the functional form of \mathbf{F} may be anything one wants, depending on what forces the system at issues is subject to.)

Now, to be able to bring theory and experiment into physically significant contact by having the capacity to identify individual and concrete models, one must demonstrate that the structures intrinsic to the theory are appropriate and adequate for representing and reasoning about the genus of physical systems the theory purports to treat. I am now in a position to give (somewhat) precise definitions for these terms. The theory is *appropriate* (or has *propriety in representation*) for treating a given system if all the quantities the theory attributes to a system are well defined and they jointly satisfy all the theory's kinematical constraints, given the state of the system in conjunction with the interactions it has with its environment. Thus, a theory can be used with

56. See Curiel (2017a) for a more comprehensive and detailed discussion of all the issues I discuss here.

propriety to treat a type of physical system it putatively represents if and only if the system's environment and its own state jointly permit the determination, within the fineness and ranges allowed by their nature, of the system's quantities over the variously relevant scales appropriate for the representation of the relations among the quantities manifested in the phenomena at issue.

Propriety is a property accruing to the representation of individual states, not to the course of a dynamical evolution: one may not be able to identify an entire concrete model with an entire individual model, but, if one is to apply the theory in a meaningful way at all to the treatment of a system, at a minimum one must be able to identify those substructures of each that refer only to single states. Since the kinematical constraints are the same for all states, it makes sense to ask whether they are satisfied so long as the quantities themselves are well defined. Indeed, I now impose a stronger criterion for the presence of a breakdown scale, in the form of a definition.

Definition 8.1 *A breakdown scale of a theory is a bound on or function of or relation among its quantities such that the kinematical constraints are satisfied to the required degree of accuracy given the experimental techniques used for probing them if and only if the joint state of the system and its environment imply that the bound or function or relation is satisfied.*

If the kinematical constraints are not satisfied, one has no reason to think that the system at issue is one the theory can treat at all.

It is a brute fact about the actual physical theories we use in real science that they become predictively inaccurate in regimes well separated from all its breakdown scales, *i.e.*, well before the theory's kinematical constraints are no longer satisfied. The theory is *adequate* if its dynamical relations are satisfied to some prescribed degree of accuracy, given the experimental techniques used for measuring it. Thus, a theory is adequate if and only if entire concrete models can be identified with entire individual models. Clearly, if the theory does not have propriety in representation, it cannot be adequate. It is only now, when we have reached the point when it makes sense to inquire after a theory's adequacy, that the issue of the accuracy of the theory's *predictions*—their "truth"—becomes meaningful.

This discussion suggests the following.

Definition 8.2 *The regime of propriety of a theory is the family of physical systems whose states, in conjunction with their interactions with their environments, are bounded on all sides, in all ways, by all the breakdown scales of the theory.*

In this regime, the theory's representational resources are appropriate for modeling the kinds of system at issue, even if they are not predictively accurate. In particular, all quantities are well defined, and all kinematical constraints are satisfied. A system is in the theory's regime of propriety if and only if one can identify those parts of an individual model representing a single state with those parts of a concrete model representing the same.

Definition 8.3 *The regime of applicability of a theory is a subset of the regime of propriety, in which the theory's dynamical equations are predictively accurate in their modeling of the behavior of the family of systems at issue.*

A system is in the regime of applicability if and only if one can identify an entire individual model with an entire concrete model.

A theory may appropriately treat a family of phenomena even when it does not model the dynamical behavior of all members of the family to any prescribed degree of accuracy, *i.e.*, even when the equations of motion are not satisfied in any reasonable sense. A theory, that is to say, can and does tell us much about the character and nature of physical systems for which it does not give accurate representations so long as they are in its regime of propriety—systems, in other words, it cannot soundly represent in totality, cannot be true of, and so systems that, according to all the standard contemporary accounts of theory structure and semantics, the theory should have nothing to say about at all. I shall now discuss four roles that propriety, even in the absence of adequacy, plays in informing and contributing to the epistemic content of a theory.

Theories do not predict kinematical constraints; they demand them. I take a prediction to be something that a theory, while appropriately modeling a system, can still get wrong. Newtonian mechanics, then, does not predict that the velocity of a body equal the temporal rate of change of its position; rather it requires it as a precondition for its own applicability. It can't "get it wrong". If the kinematical constraints demanded by a theory do not hold for a family of phenomena, that theory cannot treat it, for the system is of a type beyond the theory's scope. By contrast, if the equations of motion are not satisfied, that may tell one only that one has not taken all ambient forces on the system (couplings with its environment) into account; it need not imply that one is dealing with an entirely different type of system. Even in principle, one can never entirely rule out the possibility that the equations of motion are inaccurate only because there is a force one does not know how to account for, not because the system is not appropriately treated by those equations of motion. This can never happen with a kinematical constraint. It is either satisfied, to the appropriate and required level of accuracy given the measuring techniques available and the state of the system and its environment, or it is not. Thus, the propriety of a theory constitutes its necessary preconditions of applicability.⁵⁷

Indeed, satisfaction of kinematical constraints is required in general for the equations of motion of a theory to be well posed or even just consistent. The initial-value formulation of the Navier-Stokes equations, for example, is well set (in the sense of Hadamard) only if the shear-stress tensor is symmetric, and it is thermodynamically consistent only if the heat flux is independent of the pressure gradient (Landau and Lifschitz 1975, ch. v, §49). One cannot even formulate Newton's Second Law if velocity is not the first temporal derivative of position. More generally, in a sense one can make precise (Curiel 2014), if the kinematical constraints of Lagrangian mechanics are not satisfied ($\mathbf{v} = \dot{\mathbf{q}}$), then one cannot formulate the Euler-Lagrange equation; and similarly, if the kinematical constraints of Hamiltonian mechanics are not satisfied (the \mathbf{p} s and \mathbf{q} s do not

57. In this sense, the kinematical constraints play a role analogous to that proposed by Reichenbach (1965) and Friedman (2001) for what they call the relativized *a priori* of a theory. Kinematical constraints, however, differ in this respect: they are part of the theory itself, not supra-theoretical principles, as are the relativized *a priori* of the neo-Kantians. Also, satisfaction of the kinematical constraints can, indeed must, be *experimentally* verified in order for one to ascertain that the theory represents the system with propriety. To the contrary, satisfaction of the relativized *a priori*, in the neo-Kantian sense, does not admit of empirical verification, but rather grounds the possibility of empirical investigation in a stronger sense.

satisfy the canonical Poisson-bracket relations), then one cannot formulate Hamilton’s equation.⁵⁸ Thus satisfaction of the kinematical constraints is required as a precondition for the appropriate application of a theory in modeling a kind of system.

It is essential that all the kinematical constraints be satisfied for the theory to be able to represent any aspects of the system at all, for it is those kinematical constraints, not the equations of motion, that characterize the genus of system at issue. Many systems have, *e.g.*, shear-stress (Navier-Stokes fluids, elastic continua, Maxwell fields, charged plasmas, *etc.*); what makes a shear-stress the shear-stress of a Navier-Stokes fluid as opposed to that of a Maxwell field is its satisfaction of the Navier-Stokes kinematical constraints, *e.g.*, that the shear-stress be, in a sense one can make precise, transverse to the fluid flow. By contrast, you can’t throw a rock without hitting a system whose equations of motion are that of a simple harmonic oscillator: if I know only the equations of motion, I cannot in general tell you what kind of system I am dealing with.⁵⁹

Finally, it is the kinematical constraints, not the equations of motion, that guide the experimentalist in the design of instruments for probing and measuring the quantities the theory attributes to the systems it treats. (Recall the general discussion of this issue in §3, and the analysis of this issue in the particular case of LIGO in §7.) An instrument that is to measure Newtonian velocity, for instance, must be sensitive to differences in spatial location at ever smaller measured temporal intervals, even if only indirectly, in accord with the kinematical constraint $\dot{\mathbf{x}} = \mathbf{v}$. Such an instrument, if well designed, does not care about how the system accelerates, *i.e.*, about its dynamics. Similarly, an instrument that would measure shear-stress of a Navier-Stokes fluid must conform to the equality of pressure and reversed sense of shear across imaginary surfaces in fluid that is represented by the symmetry of the shear-stress tensor. Again, the instrument need not care at all about the dynamics of the fluid to measure the shear-stress, and indeed must be insensitive to how the shear-stress changes over time, if it is to measure the instantaneous value of the quantity. In this way, kinematical constraints provide the foundation for the operationalization of the meaning of theoretical terms.

There is one more subtlety in the idea of the regimes of a theory that must be addressed. It is in fact never the case that an individual model of a given physical system represents the system appropriately and adequately *in toto*. There is always a sense in which some “parts” of the model represent nothing physical at all. Consider again individual solutions to the Navier-Stokes equations. Because Navier-Stokes is a continuum theory, the solutions to its equations allow one to make, indeed they *necessarily encode*, what seem to be physical predictions at arbitrarily small spatial and temporal scales. We know, however, that the theory is not appropriate for that task, and we know we can safely ignore those “predictions” of the theory in assessing its adequacy and soundness, for real fluids are not continua. General relativity is scaleless—one cannot distinguish a model of a Schwarzschild black hole of radius $10^{-10^{39876}}$ cm from one of radius 44 million km.

58. (q_i, p_j) satisfy the canonical Poisson-bracket relations if

$$\begin{aligned} \{q_i, q_j\} &= 0 \\ \{q_i, p_j\} &= \delta_{ij} \\ \{p_i, p_j\} &= 0 \end{aligned} \tag{8.1}$$

where δ_{ij} is the Kronecker delta symbol, which equals 1 for $i = j$ and 0 otherwise.

59. Thus, kinematical constraints are constitutive of the systems the theory treats, in a way analogous to the relativized *a priori* of the neo-Kantians.

And yet we also know that the theory is appropriate only for treating the latter, not the former. Indeed, even in the case of a Schwarzschild black hole of radius 44 million km, which does find appropriate and adequate representation in the theory, not all “parts” of the individual model representation. One finds again the same circumstance as in Navier-Stokes theory: we have good reason not to trust “predictions” general relativity makes in such models about the geometry of spacetime at spatiotemporal scales approaching the Planck scale. Propositions about geometry at such scales formulated in the model have no empirical content. That in no way detracts from the sound epistemic content the model accrues by the propositions it allows one to formulate about geometry at much larger scales.

One comes to recognize and understand these limitations in the representational capacities of theories only by knowledge of the breakdown scales, and correlatively of the regime of propriety. Thus, the knowledge of how the mathematics is to be applied, how it represents, and what it represents, is not separable from the practical knowledge that grounds the determinations of the scales and regimes. This is not a matter of “segregating part of the mathematical structure as representational fluff (‘gauge’)”—for there just is the solution itself. The fact that one cannot use it to model phenomena happening on small spatiotemporal scales does not mean that “part” of the solution—“that part representing small stuff”—is gauge, for there is no such part cleanly separable from the rest. It is not like the case of the Faraday tensor and one of its 4-dimensional gauge potentials, two different mathematical entities with a fixed relation between them. And it is not that there is a simple rule, *e.g.*, “ignore everything in cubes whose edge is smaller than the mean free-path of the fluid’s constituents”, because such rules will in general be incorrect. Sometimes what happens on such scales is physically relevant (say, a discontinuity in a boundary condition—such a discontinuity obtains across arbitrarily small scales). One simply has to use one’s judgment, on a case by case basis, to determine how the formalism encodes potential knowledge, how it successfully represents parts of the world, how it can be used as part of sound, legitimate reasoning.⁶⁰

To hark back to the discussion, in §6, of the different forms and aspects of scientific knowledge, I argued there that that the standard view of theories gives one too little in answer to the question of a theory’s epistemic content. Here, by contrast, is a case in which the standard view tries to answer too many questions, tries to embody too much knowledge in Stein’s sense (a)—simple, factual knowledge—in so far as it allows one to formulate propositions that can have no empirical content *even about systems falling in the theory’s regime of applicability*.

This characterization of the regime of propriety and the regime of applicability makes explicit the ways that schematizing the observer is required for a theoretical structures to make substantive contact with experimentation, and so for a theory as a whole to have non-trivial epistemic content.

60. This problem underscores the shallowness of much of the debate about realism in science, in particular the almost ubiquitous yet deeply problematic implicit assumption that “realism” is a holistic attitude—a mathematical structure *in toto* “really represents” or it does not. That, however, is too coarse a conception. Some aspects of a mathematical structure may represent, and others not, *e.g.*, the integrated phase differences in a quantum wave-function (Berry phase) may be “real” whereas the phase itself may not, or those “parts” of a Navier-Stokes model representing behavior on scales greater than 10^{-4} cm may “really designate” whereas those parts of representing smaller scale behavior may not. Halvorson (2018) provides a splendid discussion of these matters.

9 Semantics Is Epistemology, Not Ontology

The conclusions of my discussion of breakdown scales and the regimes of a theory have a deep consequence for how meaning and truth may cogently be related in an adequate semantics for theories. According to the kind of semantics that naturally accompanies the standard view of theories, all of a theory's propositions must be (at least approximately) true of a system in order for the system to represent it. By my arguments, however, the meanings of a theory's theoretical terms must be fixed by the epistemic content associated with the regime of propriety, not the regime of applicability, *i.e.*, in part by situations in which *not* all of a theory's propositions about the system are true. The meanings of the theory's propositions, therefore, cannot be fixed solely by their truth-conditions in any standard sense—the seductive intuition grounding essentially all contemporary thought on the semantics of scientific theories, as Carnap (1942, ch. B, §7, p. 22) concisely expresses it: “. . . to understand a sentence, to know what is asserted by it, is the same as to know under what conditions it would be true.” One cannot even begin to investigate what the truth conditions of some of the sentences in the theory may be, however—in particular, its dynamical predictions—until one already knows enough about the meaning of its terms to ascertain the truth of some of its other sentences, *viz.*, the kinematical constraints, that guarantee that the system at issue falls within the regime of propriety. To verify that the kinematical constraints are appropriately satisfied, however, is necessarily grounded in the possibility of the schematic representation of the observer in our formal models.

One can think of propriety in representation, in part, therefore, as what a theory must have for it to have the capacity to produce propositions whose truth-value can be cogently investigated—not a fixing of truth conditions, but rather the securing of the possibility to investigate whether or how truth-conditions for a given proposition can be determined in the first place. A theory does not possess even the capacity to be accurate or inaccurate in its treatment of a family of phenomena if it does not represent the phenomena with propriety. It follows that one can not even entertain questions about the truth of many sorts of propositions—most of all those depending on the identification of individual models with concrete models—until one has determined that the theory has the resources, both practical and formal, to represent the system at issue with propriety, *i.e.*, until meaning already has accrued to the formal structures of the theory. The fact that the regime of propriety is strictly larger than the regime of applicability, therefore, shows that the fundamental idea of semantics should rather be: to understand a sentence, to know what is asserted by it, is the same as to know under what conditions its constituent terms can cogently be assigned meaning and so allow one to investigate the possibility of ascertaining the truth conditions of the sentence.

Indeed, I believe that all the problems I have discussed with the standard view—in particular, its lack of accounting for all kinds of knowledge, and its incapacity to handle the breakdown of models—boil down to its implicit reliance on something like an essentially truth-conditional semantics, in conjunction with designation as the fundamental semantic relation: semantics is fixed by ontology. This is what the assumed clean separation of formal from practical spheres of knowledge amounts to in current philosophical work, and, conversely, that clean separation is implied by such a semantics.⁶¹

61. My reliance on the formally characterized distinction between kinematical and dynamical structures (§8) is

More precisely, a view about the structure and semantics of physical theory based ultimately on ontology, grounded in the assumption of a clean split between the parts of the epistemic content of a theory captured wholly by its formalism and all other parts, is inadequate for (at least) two reasons. First, it does not allow us, within the scope of the theory itself, to understand why models of systems in the regime of propriety but not in the regime of applicability are not accurate even though all the quantities the theory attributes to the system are well defined and the values of those quantities jointly satisfy all kinematical constraints the theory requires. Second, we miss something fundamental about the meaning of various theoretical terms by rejecting such models out of hand merely on the grounds of their inaccuracy. It is surely part of the semantics of the term ‘hydrostatic pressure’, *e.g.*, that its definition as a physical quantity treated by classical fluid mechanics breaks down when the fluid approaches turbulence closely enough; because, however, the theory’s equations of motion stop being accurate long before, in a precise sense, the quantity loses definition in the theory and long before the kinematical constraints of the theory stop being satisfied, any account of the structure of theories and their semantics that rejects the inaccurate models in which the term still is well defined will not be able to account for that part of the term’s meaning. Thus, an adequate account of physical theory must be grounded on notions derived from relations in some sense prior to the theory’s representations of the dynamical behavior of the physical systems it treats, relations that govern the propriety of the theory’s representational resources for modeling the system at issue. These are the theory’s kinematical constraints.

One may think that this discussion about how, where and when theories break down more properly belongs to pragmatics (in the sense of semiotic theory) than to semantics. That is not so. A system of formal semantics that would ground itself in the family of possible physical systems for which it provides sound models cannot even get started until that family is demarcated. But that is exactly to require an investigation of the boundary of the theory’s regime of propriety, which is thus logically and conceptually prior to any such system of semantics.⁶²

According to a semantics that requires predictive accuracy, such as a Tarskian one and in general almost every one conforming to the standard view, the idea of the regime of propriety is meaningless. Such a semantics cannot explain or even accommodate this fact about our theories, for *predictively inaccurate models cannot be Tarskian models or possible worlds or objects of the category of solutions to the equations of motion of the theory*. Semantics founded on the standard

not inconsistent with my contention that a good semantics allows no clean separation of epistemic content into that part captured by the formalism and that part that is not. It may still be the case—and in fact is the case—that different parts of the formalism play different roles in the analysis of the epistemic content, indeed in characterizing the inextricable integrity of epistemic content itself as not being cleanly separable into formal and non-formal parts.

62. Jeremy Butterfield has suggested to me that an intensional semantics of the sort, *e.g.*, that Lewis (1970a) propounds—which manifestly seems to lend itself to the articulation of a semantics for theories in conformity with the standard view—can handle these pragmatic issues as part of the formal semantics itself, in so far as such pragmatic issues are encoded in the intensions. I do not think that is right. Such a semantics would evaluate a proposition in Navier-Stokes theory purporting to describe the behavior of a fluid at spatial scales of 10^{-100} cm as false. Such propositions are not false, however. They are meaningless. None of the quantities Navier-Stokes theory attributes to systems are well defined at such scales, so no “proposition” purportedly referring to them can have meaning. It is folly to require that every well formed formula constructible in the terms of the formalism of a physical theory must have a determinate truth value, much more a meaning. If one could re-tool such a scheme of intensional semantics so that its domain is not “all possible worlds” (or “all possible models” in some other appropriate sense), however, but rather a pragmatically characterized regime of propriety, then it might in fact be a useful tool for trying to define a semantics grounded on propriety in representation. I am not sure how such a thing could be done.

view, therefore, does not exhaust the representational capacity of the theory, and the theory gains non-trivial semantic content from everything it can significantly represent, whether in all accuracy or not. The set of possible worlds picked out by satisfaction of the equations of motion is not a rich enough family of worlds to express or encode all the information the theory can give us about the possibilities of the actual physical world. The theory tells us more about physical quantities like pressure in the actual world than there would be to learn about it in a world the theory would be true of, in the standard sense.

Thus, the regime of propriety must be included as part of the theory's semantics—at least so far as a real semantics of a real physical theory goes, not just a formal semantics of a formal theory: if I don't know the family of actual and physically (not mathematically) possible systems the theory applies to, I don't, by lights of the standard view itself, know the semantics; but if I don't know the regime of propriety, I don't know that family; and nothing in the formalism of the theory itself can tell me the regime—I cannot fix the “ontology” by reasoning grounded on a clear separation of the theory's epistemic content into one part captured by the formalism and “interpretive postulates” and another part corresponding only to practice.

A semantics grounded on pure ontology assumes a kind of Fichtean direct intellectual grasp of the world by our representational systems: the posited relation between our symbolic systems and the “objects in the world” (irrespective of whether one is a realist about such things or not) is unmediated by the actual state of our knowledge and by the practices and techniques we have for probing, experimenting on, and more generally investigating the world and evaluating the results of those investigations.⁶³

Have we learned nothing from Kant? What sense is there in trying to articulate “truth conditions” that are forever beyond our cognitive grasp? What was wanted was an account of “meaning” analytically connected to truth, completely divorced from human concerns. But this is self-defeating, for such a conception *eo ipso* completely divorces, unbridgeably separates, semantics from the *fundamental sources* of scientific knowledge—experimental knowledge—which in the end must ground the empirical content and significance of our theoretical representations. Relations of “direct designation” serve—can serve—no philosophical or foundational purpose. They can tell us nothing about meaning. They are vacuous chicanery, nonsensical will-o-the-wisps leading us to a morass of philosophical quicksand into which we hopelessly sink, suffocating on our own confusion.⁶⁴

63. These criticisms do not apply to the use of any Tarskian-like semantics in logic and mathematics, where one cannot cleanly and unambiguously separate our symbolic systems from the objects they purport to represent; or, at least, the kind of access we may have there to such objects is mediated only by the symbolic systems, not by experimental practice.

64. Stein (1989, p. 50, emphasizes his) puts the point forcefully, when he imagines Kant posing the following dilemma to Leibniz's ghost:

How can you *know* that things are as you say they are? If the claimed “reference” of the theory is something beyond its correctness and adequacy in representing phenomena – if, that is, for a given theory, which (we may suppose) does represent phenomena correctly and adequately, there are still two possibilities: (a) that it is (moreover) *true*, and (b) that it is (nonetheless) *false* – then how in the world could we ever tell what the actual case is?

Stein talks here of representing phenomena “correctly and adequately”. His distinction, as I understand it, is related to my own “appropriately and adequately” in interesting ways. For Stein, “correctly” is something like my “appropriately and accurately”, and “adequately” is something like covering a large enough extent to think one has

These criticisms assume a link between semantics and knowledge in all its human forms, as explicated and discussed in §6. I think this must be right, that there must be such a link between semantics and real human knowledge. Any semantics *à la* the standard view, however, pertains, at best, only to the first kind, knowledge as achieved state. A semantics completely divorced from our actual state of knowledge (as achieved state, as provider of evidentiary relations for epistemic warrant, as guide to future investigation), and from the ways we have of improving our epistemic state, is useless.

We must expel from philosophy the myth of a “human-free” semantics. *All* aspects of scientific knowledge, not just the “ontological”, should be reflected to at least some degree in a theory’s semantic content, and most of all that knowledge ultimately grounding the semantic content—and that is indubitably, inextricably, inexorably *experimental* in large part. This view may entail a blurring of the lines between the traditional conception of semantics and pragmatics (in the sense of semiotic), but that, I think, is all to the good, since the traditional notions are (*pace* Carnap, Suppes, *et al.*) appropriate for mathematics, not the empirical sciences. In any event, I am not convinced that one needs to lose a sharp distinction between semantics in a formal sense and pragmatics in order to ground the kind of view I advocate—one needs only to characterize semantics in a way that is not wholly “ontological”, based on a primitive relation of designation.⁶⁵

Much of what semantics based on the standard view by its nature must leave to the pragmatics does not properly belong there. That addition of momentum in Newtonian mechanics means something different in different contexts—in this instance, a representation of a physical interaction (say, particle collision), in that one, part of the calculation of a factitious quantity such as the center of momentum—does not turn on the fact that this or that individual investigator employs it differently. The two uses differ in intrinsic physical significance, the sort of thing, if anything, a semantics ought to capture. A semantics fixed entirely in terms of the formalism, however, cannot distinguish the two cases. They are formally identical. Vector addition is vector addition no matter where and when it happens. Divorcing semantics from real human epistemology means that we lose the capacity to analyze in a principled way how epistemic warrant accrues to scientific propositions, and, more importantly, to investigate how different amounts of warrant accrue to different propositions, and how much warrant a given proposition should have. This all must be based on our understanding of the connection of theory with experiment.

If I were developing a semantics that at bottom grounded semantic content on something like truth as characterized by predictive accuracy, I could at this point call upon Stalnaker’s admirably elegant and succinct summation of the appropriate relation between pragmatics and semantics in such systems (1981, pp. 44–45, emphases mine):

Now that we have found an answer to the question, “How do we decide whether or
latched onto *general* features of the world, not just something peculiar to the idiosyncrasy of *this* kind of system studied by *this* kind of experiment. This kind of comprehension in coverage is a fundamental part of the epistemic content of a good theory, but it would take us too far afield to discuss it here.

65. I want to emphasize that I am not opposed to referential relations and the idea of designation itself in a semantics for theories. I am opposed only to the idea that such relations be primitive. If one wants to incorporate a notion of reference into one’s semantics, I see no reason one cannot do so, so long as it is a relation that (1) is defined with respect to a prior characterization of meaning, such as one based on satisfaction of kinematical constraints, and (2) has a target domain that is delimited by the kinds of pragmatic considerations I have discussed. (See footnote 62.)

not we believe a ... statement?" the problem is to make the transition from belief conditions to truth conditions; that is, to find a set of truth conditions for statements ... *which explains why we use the method we do use to evaluate them.*

In other words, the relation of pragmatics and semantics in a system that grounds meaning in truth conditions should reflect the relation between epistemology and ontology: one's formal truth conditions (semantics by way of ontology, as truth conditions ought not depend on the context of the individual knower) should relate in the appropriate way to belief conditions (pragmatics by way of epistemology, as belief conditions depend on the context of the individual believer). This must hold at a minimum if our beliefs are even to have a shot at tracking truth, when meaning is defined by its relation to truth in such a way.

In a system of semantics that would capture *all* forms and aspects of scientific knowledge, by contrast, in which meaning cannot be grounded on truth conditions alone, such as accuracy in prediction, but is rather grounded on the way a theory provides appropriate structures for the meaningful representation of a given (type of) system, before one can even ask about the truth values of all the propositions in the theory—in such a system of semantics, the relation between pragmatics and semantics should remain silent about the relation between epistemology and ontology; in the best of cases, it should be compatible with many possible such relations.⁶⁶

As I have argued, in order to know how to investigate whether or not a theory provides an accurate representation of a system—whether that system falls in the theory's regime of applicability—one must be able to verify first whether or not the system satisfies the theory's kinematical constraints, *i.e.*, whether or not it falls in the theory's regime of propriety. Our problem, therefore, the analogue to Stalnaker's, is how to move from an understanding of how to verify whether or not a theory's kinematical constraints are satisfied to an understanding of how to use the resources of the theory to represent a system once we have verified the kinematical constraints are satisfied. As Stalnaker says, moreover, our account of this process should "[explain] why we use the method we do use to [verify] them." Because, as I argued, the identification of (parts) of a concrete model of a system with (parts) of the formal structures of a theory is a pragmatic choice, the problem is to find justifications for those choices in the success and fruitfulness of the subsequent use we make of theory based on them. Pragmatics, on this picture, still has to do with the way that the contexts of *individuals*, in their utterances, express meaning, and that context includes, *inter alia*, the individuals' beliefs; but what individuals believe and express, in any given context, does not bear on the objectivity of the knowledge embodied in the epistemic content theory, of which the semantics ought to ground the articulation and representation.

Semantics, therefore, must ground analysis of epistemology, and be grounded in turn by our grasp of it. NOT:

$$\begin{aligned} \text{semantics} &\approx \text{ontology} \\ \text{pragmatics} &\approx \text{use} \end{aligned}$$

66. In any event, I would argue that ontology is a much more subtle affair than can be expressed by a semantics of the kind naturally accompanying the standard view, a broader and vaguer enterprise than formal truth-conditional semantics can accommodate; and there is much more to content and representation than is captured by standard views of semantics. For my reasons for saying this, see my discussions of physicality and the different ways that one may talk about existence in physics in Curiel (2016, 2019b).

RATHER:

semantics \approx epistemology, methodology
pragmatics \approx acceptance, choice

The Slogan:

Meaning comes before truth.

10 Valediction

The miracle of science is that theory and experiment are consonant with each other; the necessity of science is that they are inextricably so—not, however, as equals. Theory plays Boswell to the subtle and tragic clown of experiment’s Johnson.

Appendix: Summary

I sum up here the main claims of the paper.

1. A theory is characterized by its epistemic content, the sum total of all forms of knowledge it embodies, in all their aspects and relations to each other, as determined by our actual current state of knowledge about the world as a whole. (§2, §4, §6)
2. The total family of physical systems the theory appropriately and adequately treats, a proper subset of all those “possibly representable” by all its formal models, forms an essential part of that epistemic content. (§8)
3. There are (at least) two forms of knowledge (§6):
 - a. theoretical (what can be learned from books);
 - b. practical (what can be fully understood only by doing).
4. Those forms have (at least) four aspects (§6):
 - a. as achieved state;
 - b. as susceptible of justification, and so involving a structure of evidential relations;
 - c. as ground for epistemic warrant, and so involving a structure of evidential relations;
 - d. as an enterprise, an activity aimed at increasing the first aspect as constrained by the second and third.
5. Much of all of those forms and aspects is grounded on and embodied in our knowledge of how to schematically represent the observer in models of actual experiments and observations (§5-8). Perhaps most importantly, it is needed to:
 - a. lay down adequate definitions for physical quantities (§3);
 - b. determine the theory’s breakdown scales (§8);

- c. demarcate the theory's regime of applicability, the total family of systems the theory appropriately and adequately treats, as delimited by the breakdown scales (§8).
6. A good semantics for theory should respect the fact that the epistemic content of a theory is not cleanly separable into formal and practical parts; in particular, semantics should be based on and reflect epistemology, not ontology. (§9)

I sum up here my main problems with the standard philosophical view of theories.

1. It assumes a clean separation between the theoretical and the practical forms and aspects of scientific knowledge.
2. It assumes that only one form (theoretical), in one of its aspects (achieved state as represented purely by formalism), suffices for characterizing a theory.
3. Thus, it cannot demarcate, much less identify, the family of physical systems the theory appropriately and adequately treats.
4. Thus, it cannot explain the epistemic warrant we have for trusting and using the theory, and for believing that the understanding and comprehension it seems to give us of the world is actually about the world.

References

- Abbott, B. *et al.* (LIGO Scientific Collaboration and Virgo Collaboration). 2016. "Observation of Gravitational Waves from a Binary Black Hole Merger". *Physical Review Letters* 116 (6): 061102. Preprint: [arXiv:1602.03837 \[gr-qc\]](https://arxiv.org/abs/1602.03837), doi:[10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102).
- Benedict, R. 1969. *Fundamentals of Temperature, Pressure and Flow Measurements*. New York: John Wiley & Sons, Inc.
- Bogen, J., and J. Woodward. 1988. "Saving the Phenomena". *The Philosophical Review* xcvi (3): 303–352. doi:[10.2307/2185445](https://doi.org/10.2307/2185445).
- Buchwald, J. 1994. *The Creation of Scientific Effects: Heinrich Hertz and Electric Waves*. Chicago: University of Chicago Press.
- Burgess, G. 1928. *The International Temperature Scale*. Technical report RP 22(635). The Journal of Research of the National Bureau of Standards.
- Butterfield, J. 2019. "On Dualities and Equivalences between Physical Theories". In *Philosophy beyond Spacetime*, edited by N. Huggett, B. Le Bihan, and C. Wüthrich. Forthcoming. Oxford: Oxford University Press.
- Callendar, H. 1887. "On the Practical Measurement of Temperature". *Proceedings of the Royal Society of London* 178:160.
- Carnap, R. 1942. *Introduction to Semantics*. Volume I. Studies in Semantics. Cambridge, MA: Harvard University Press.

- Carnap, R. 1956. “Empiricism, Semantics and Ontology”. In *Meaning and Necessity: A Study in Semantics and Modal Logic*, Second, 205–221. First edition published 1947. Midway Reprint Edition printed 1988. Chicago: The University of Chicago Press. An earlier version was published in *Revue Internationale de Philosophie* 4(1950):20–40.
- Chandrasekhar, S. 1961. *Hydrodynamic and Hydromagnetic Stability*. Oxford: Oxford University Press.
- Chang, H. 2008. *Inventing Temperature: Measurement and Scientific Progress*. Oxford: Oxford University Press. doi:10.1093/0195171276.001.0001.
- Cohen, J., and C. Callender. 2009. “A Better Best System Account of Lawhood”. *Philosophical Studies* 145 (1): 1–34. doi:10.1007/s11098-009-9389-3.
- Collins, H. 1992. *Changing Order: Replication and Induction in Scientific Practice*. Chicago: University of Chicago Press.
- Collmar, W., N. Straumann, S. Chakrabarti, G. 't Hooft, E. Seidel, and W. Israel. 1998. “Panel Discussion: The Definitive Proofs of the Existence of Black Holes”. Chapter 22 in *Black Holes: Theory and Observation*, edited by F. Hehl, C. Kiefer, and R. Metzler, 514:481–489. Lecture Notes in Physics. Berlin: Springer-Verlag. Proceedings of the 179th W. E. Heraeus Seminar, Bad Honnef, Germany, 18–22 Aug 1997., doi:10.1007/978-3-540-49535-2_22.
- Curiel, E. 2014. “Classical Mechanics Is Lagrangian; It Is Not Hamiltonian”. *British Journal for the Philosophy of Science* 65 (2): 269–321. doi:10.1093/bjps/axs034.
- . 2016. “On the Existence of Spacetime Structure”. *British Journal for the Philosophy of Science*: axw014. Published online. Preprint: arXiv:1503.03413 [physics.hist-ph]. A manuscript containing technical appendices working out details of some of the constructions and arguments, and containing further discussion of the possible observability of different kinds of spacetime structure, is available at <http://strangebeautiful.com/papers/curiel-exist-st-struct-tech-apdx.pdf>., doi:10.1093/bjps/axw014.
- . 2017a. “Kinematics, Dynamics, and the Structure of Physical Theory”. Unpublished manuscript. Draft available at <<http://strangebeautiful.com/papers/curiel-kin-dyns-struc-theory.pdf>>.
- . 2017b. “On the Propriety of Physical Theories as a Basis for Their Semantics”. Unpublished manuscript. Draft available at <<http://strangebeautiful.com/papers/curiel-propriety-adequacy-basis-sems.pdf>>.
- . 2019a. “The Many Definitions of a Black Hole”. *Nature Astronomy* 3:27–34. Free read-only SharedIt link: <<https://rdcu.be/bfNpM>>. Preprint: arXiv:1808.01507 [physics.hist-ph]., doi:10.1038/s41550-018-0602-1.

- Curiel, E. 2019b. “What Does It Mean to Ask Whether There Is Something Rather Than Nothing?” Unpublished manuscript, prepared for presentation at the symposium “Understanding Our Place in the Universe: Beyond the Legacy of Stephen Hawking” organized by the John Templeton Foundation, Mishkenot Sha’ananim Conference Center, Jerusalem, 17–19 March 2019. Available online: <<http://strangebeautiful.com/papers/curiel-what-means-something-not-nothing.pdf>>.
- Da Costa, N., and S. French. 2005. *Science and Partial Truth: A Unitary Approach to Models and Scientific Reasoning*. Oxford: Oxford University Press.
- Davidson, D. 1969. “True to the Facts”. *Journal of Philosophy* 66 (21): 748–764. doi:[10.2307/2023778](https://doi.org/10.2307/2023778).
- Demopoulos, W. 2013. *Logicism and Its Philosophical Legacy*. Cambridge: Cambridge University Press.
- Earman, J. 1978. “Combining Statistical-Thermodynamics and Relativity Theory: Methodological and Foundations Problems”. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association 1978* 2:157–85. doi:[10.1086/psaprocbienmeetp.1978.2.192467](https://doi.org/10.1086/psaprocbienmeetp.1978.2.192467).
- Earman, J., C. Glymour, and J. Stachel, editors. 1977. *Foundations of Space-Time Theories*. Volume VII. Minnesota Studies in Philosophy of Science. Minneapolis: University of Minnesota Press. Freely available online: <<http://mcps.umn.edu/philosophy/vol8.html>>.
- Eckart, A., A. Hüttemann, C. Kiefer, S. Britzen, M. Zajaček, C. Lämmerzahl, M. Stöckler, M. Valencia-S., V. Karas, and M. García-Marín. 2017. “The Milky Way’s Supermassive Black Hole: How Good a Case Is It?” doi:[10.1007/s10701-017-0079-2](https://doi.org/10.1007/s10701-017-0079-2). Preprint: [arXiv:1703.09118](https://arxiv.org/abs/1703.09118) [[astro-ph.HE](https://arxiv.org/abs/1703.09118)]. *Foundations of Physics* 47 (5): 553–624.
- Ellis, G. 1967. “Dynamics of Pressure-Free Matter in General Relativity”. *Journal of Mathematical Physics* 8 (5): 1171–1194. doi:[10.1063/1.1705331](https://doi.org/10.1063/1.1705331).
- Foias, C., O. Manley, R. Rosa, and R. Temam. 2001. *Navier-Stokes Equations and Turbulence*. Volume 83. Encyclopedia of Mathematics and Its Applications. Cambridge: Cambridge University Press. doi:[10.1017/CB09780511546754](https://doi.org/10.1017/CB09780511546754).
- van Fraassen, B. 1980. *The Scientific Image*. Oxford: Oxford University Press.
- . 2008. *Scientific Representation: Paradoxes of Perspective*. Oxford: Oxford University Press.
- Franklin, A. 1986. *The Neglect of Experiment*. Cambridge: Cambridge University Press. doi:[10.1017/CB09780511624896](https://doi.org/10.1017/CB09780511624896).
- Friedman, M. 2001. *The Dynamics of Reason*. Delivered as the 1999 Kant Lectures at Stanford University. Stanford, CA: CSLI Publications.
- Genzel, R., A. Eckart, T. Ott, and F. Eisenhauer. 1997. “On the Nature of the Dark Mass in the Centre of the Milky Way”. *Monthly Notices of the Royal Astronomical Society* 291 (1): 219–234. doi:[10.1093/mnras/291.1.219](https://doi.org/10.1093/mnras/291.1.219).

- Geroch, R. 1969. “Limits of Spacetimes”. *Communications in Mathematical Physics* 13 (3): 180–193. Open access at <http://projecteuclid.org/euclid.cmp/1103841574>, doi:doi:10.1007/BF01645486.
- . 1970. “Domain of Dependence”. *Journal of Mathematical Physics* 11 (2): 437–449. doi:10.1063/1.1665157.
- . 1977. “Prediction in General Relativity”. In Earman, Glymour, and Stachel 1977, 81–93.
- . 1981. *General Relativity from A to B*. Chicago: University of Chicago Press.
- . 1996. “Partial Differential Equations of Physics”. In *General Relativity*, edited by G. Hall and J. Pulham, 19–60. Aberdeen: Scottish Universities Summer School in Physics. Proceedings of the 46th Scottish Universities Summer School in Physics, Aberdeen, July 1995. Preprint: [arXiv:gr-qc/9602055](https://arxiv.org/abs/gr-qc/9602055).
- . 2001. “On Hyperbolic “Theories” of Relativistic Dissipative Fluids”. [arXiv:gr-qc/0103112v1](https://arxiv.org/abs/gr-qc/0103112v1)).
- Ghez, A., M. Morris, E. Becklin, A. Tanner, and T. Kremenek. 2000. “The Accelerations of Stars Orbiting the Milky Way’s Central Black Hole”. *Nature* 407:349–351. doi:10.1038/35030032.
- Hall, J. 1955. “The International Temperature Scale”. In *Temperature*, 2:116. New York: Reinhold.
- Halvorson, H. 2018. “To Be a Realist about Quantum Theory”. Unpublished manuscript. URL: <http://philsci-archive.pitt.edu/14310/>.
- Halvorson, H., and D. Tsementzis. 2017. “Categories of Scientific Theories”. In Landry 2017, chapter 17.
- Harper, W. 2011. *Isaac Newton’s Scientific Method: Turning Data into Evidence about Gravity and Cosmology*. Oxford: Oxford University Press.
- Haynes, W., editor. 2014. *CRC Handbook of Chemistry and Physics*. 95th edition. Boca Raton, FL: CRC Press.
- Helmholtz, H. 1870. “Über den Ursprung und die Bedeutung der geometrischen Axiome”. In *Populäre Wissenschaftliche Vorträge von H. Helmholtz*, 3:23–51. Braunschweig. A lecture delivered in the Docentverein of Heidelberg, 1870. English translation by Howard Stein (unpublished), “On the Origin and Significance of the Geometrical Axioms”, available at <http://strangebeautiful.com/other-minds.html#helmholtz..>
- Hertz, H. 1893. *Electric Waves, Being Researches on the Propagation of Electric Action with Finite Velocity through Space*. Trans. D. Jones. Originally published as *Untersuchungen über die Ausbreitung der elektrischen Kraft*, Leipzig: J. A. Barth, 1892. This is a 1962 reprint of the English edition which first appeared in 1893, published by Macmillan and Sons. New York: Dover Press.
- Kovalevsky, J., and P. Seidelmann. 2004. *Fundamentals of Astrometry*. Cambridge: Cambridge University Press.

- Lakatos, I. 1970. “Falsification and the Methodology of Scientific Research Programmes”. In *Criticism and the Growth of Knowledge*, edited by I. Lakatos and A. Musgrave, 91–196. Cambridge: Cambridge University Press. Reprinted in his *Methodology of Scientific Research Programmes: Philosophical Papers*, vol. 1, Cambridge: Cambridge University Press, 1980, eds. J. Worrall and G. Currie, ch. 1, pp. 8–101.
- Lamb, H. 1932. *Hydrodynamics*. Sixth. New York: Dover Publications. The 1945 Dover reprint of the 1932 edition published by the Cambridge University Press.
- Landau, L., and E. Lifschitz. 1975. *Fluid Mechanics*. Second. Oxford: Pergamon Press. An expanded, revised edition of the original 1959 edition. Translated from the Russian by J. Sykes and W. Reid.
- Landry, E., editor. 2017. *Categories for the Working Philosopher*. Oxford: Oxford University Press. doi:[10.1093/oso/9780198748991.001.0001](https://doi.org/10.1093/oso/9780198748991.001.0001).
- Laudan, L. 1977. *Progress and Its Problems: Towards a Theory of Scientific Growth*. Berkeley: University of California Press.
- Lewis, D. 1970a. “General Semantics”. *Synthese* 22 (1–2): 18–67. doi:[10.1007/BF00413598](https://doi.org/10.1007/BF00413598).
- . 1970b. “How to Define Theoretical Terms”. *The Journal of Philosophy* 67 (13): 427–446. Stable URL: <<http://www.jstor.org/stable/2023861>>.
- Lockman, F. 2005. “Can Remote Observing be Good Observing? Reflections on Procrustes and Antaeus”. [arXiv:astro-ph/0507140](https://arxiv.org/abs/astro-ph/0507140). A slightly edited version of a paper published in 1993 in *Observing at a Distance*, D. T. Emerson and R. G. Clowes (eds.), World Scientific, Singapore, pp. 325–340, with a new Afterword.
- Lutz, S. 2014. “What’s Right with a Syntactic Approach to Theories and Models?” *Erkenntnis* 79 (8 Supplement): 1475–1492. doi:[10.1007/s10670-013-9578-5](https://doi.org/10.1007/s10670-013-9578-5).
- Malament, D. 1977. “Observationally Indistinguishable Spacetimes: Comments on Glymour’s Paper”. In Earman, Glymour, and Stachel 1977, 61–80.
- . 2002. “A No-Go Theorem About Rotation in Relativity Theory”. In *Reading Natural Philosophy: Essays in the History and Philosophy of Science and Mathematics*, edited by D. Malament. Chicago: Open Court Press. Essays presented to Howard Stein in honor of his 70th birthday, delivered at a *Festschrift* in Stein’s honor at the University of Chicago, May, 1999.
- . 2003. “On Relative Orbital Rotation in General Relativity”. In *Revisiting the Foundations of Relativistic Physics: Festschrift for John Stachel*, edited by A. Ashtekar. Dordrecht: Kluwer.
- Manchak, J. 2009. “Can We Know the Global Structure of Spacetime?” *Studies in History and Philosophy of Modern Physics* 40 (1): 53–56. doi:[10.1016/j.shpsb.2008.07.004](https://doi.org/10.1016/j.shpsb.2008.07.004).
- Maxwell, J. C. 1871. “Introductory Lecture on Experimental Physics”. In Maxwell 1965, volume II, 240–255.
- . 1876. “General Considerations concerning Scientific Apparatus”. In Maxwell 1965, volume II, 505–522.

- Maxwell, J. C. 1888. *Theory of Heat*. The Dover edition of 2001 republishes in unabridged form the ninth edition of 1888 published by Longmans, Green and Co., London, and also includes the corrections and notes of Lord Rayleigh incorporated into the edition of 1891. Mineola, NY: Dover Publications, Inc.
- . 1965. *The Scientific Papers of J. C. Maxwell*. New York: Dover Publications, Inc. W. Niven (Ed.). Two volumes, published as one.
- Newton, I. 1999. *The Principia: Mathematical Principles of Natural Philosophy*. Berkeley, CA: University of California Press. Translation of *Philosophiæ Naturalis Principia Mathematica* by I. Cohen and A. Whitman, based on the Third Edition text of 1726. (The first edition of *Principia* was published in 1687, the second in 1713.)
- Peirce, C. S. 1878. “The Doctrine of Chances”. Chapter 9 in *The Essential Peirce: Selected Philosophical Writings*, volume 1 (1867–1893), 142–154. Edited by N. Houser and C. Kloesel. Bloomington, IN: Indiana University Press. Originally published in *Popular Science Monthly* 12(March, 1878):604–615.
- Pope, S. 2000. *Turbulent Flows*. Cambridge: Cambridge University Press. doi:[10.1017/CB09780511840531](https://doi.org/10.1017/CB09780511840531).
- Popper, K. 1959. *The Logic of Scientific Discovery*. London: Hutchinson & Co.
- Quine, W. 1980. “Two Dogmas of Empiricism”. Chapter II in *From a Logical Point of View: Nine Logico-Philosophical Essays*, 2nd, revised, 20–46. Cambridge, MA: Harvard University Press. Emended version of the article originally published in the *Philosophical Review*, 1951, 60(1, January):20–43.
- Reichenbach, H. 1965. *The Theory of Relativity and A Priori Knowledge*. Berkeley, CA: University of California Press.
- Reintjes, M., and B. Temple. 2017. “Shock Wave Interactions in General Relativity: The Geometry behind Metric Smoothing and the Existence of Locally Inertial Frames”. Preprint: [arXiv:1610.02390 \[gr-qc\]](https://arxiv.org/abs/1610.02390).
- Smith, G. 1999a. “Newton and the Problem of the Moon’s Motion”. In Newton 1999, chapter 8.15.
- . 1999b. “The Motion of the Lunar Apsis”. In Newton 1999, chapter 8.16.
- Stalnaker, R. 1981. “A Theory of Conditionals”. In *Ifs: Conditionals, Belief, Decision, Chance, and Time*, edited by W. Harper, R. Stalnaker, and G. Pearce, 41–56. The University of Western Ontario Series in Philosophy of Science 15. Dordrecht: D. Reidel Publishing Co. First published in *Studies in Logical Theory, American Philosophical Quarterly Monograph Series*, No. 2, Oxford: Blackwell, 1968.
- Stegmüller, W. 1979. *The Structuralist View of Theories: A Possible Analogue of the Bourbaki Programme in Physical Science*. Berlin: Springer-Verlag. doi:[10.1007/978-3-642-95360-6](https://doi.org/10.1007/978-3-642-95360-6).
- Stein, H. 1989. “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](https://doi.org/10.1111/j.1746-8361.1989.tb00930.x).

- Stein, H. 1992. “Was Carnap Entirely Wrong, After All?” *Synthese* 93:275–295. doi:[10.1007/BF00869429](https://doi.org/10.1007/BF00869429).
- . 1994. “Some Reflections on the Structure of Our Knowledge in Physics”. In *Logic, Methodology and Philosophy of Science*, edited by D. Prawitz, B. Skyrms, and D. Westerståhl, 633–655. Proceedings of the Ninth International Congress of Logic, Methodology and Philosophy of Science. New York: Elsevier Science B.V.
- Stimson, H. 1949. *The International Temperature Scale of 1948*. Technical report RP 1926(209). The Journal of Research of the National Bureau of Standards.
- . 1961. *The International Temperature Scale of 1948, Text Revision of 1960*. Technical report 37. National Bureau of Standards Monograph.
- Suárez, M. 2004. “An Inferential Conception of Scientific Representation”. *Philosophy of Science* 71 (5): 767–779. doi:[10.1086/421415](https://doi.org/10.1086/421415).
- Suppe, F. 1974. “The Search for Philosophic Understanding of Scientific Theories”. In *The Structure of Scientific Theories*, edited by F. Suppe, 3–254. Chicago: University of Illinois Press.
- Suppes, P. 1960. “A Comparison of the Meaning and Uses of Models in Mathematics and the Empirical Sciences”. *Synthese* 12 (2–3): 287–301. doi:[10.1007/BF00485107](https://doi.org/10.1007/BF00485107).
- . 1962. “Models of Data”. In *Logic, Methodology and Philosophy of Science*, edited by E. Nagel, P. Suppes, and A. Tarski, 252–261. Palo Alto, CA: Stanford University Press.
- Thompson, A., J. Moran, and G. Swenson. 2017. *Interferometry and Synthesis in Radio Astronomy*. Third. Astronomy and Astrophysics Library. Cham, Switzerland: Springer International Publishing AG. The book is available as open-access on Springer’s website: <<https://link.springer.com/book/10.1007/978-3-319-44431-4>>, doi:[10.1007/978-3-319-44431-4](https://doi.org/10.1007/978-3-319-44431-4).
- Weatherall, J. 2017. “Categories and the Foundations of Classical Space-Time Theories”. In Landry 2017, chapter 13.