Against 'The System'[1]

Nancy Cartwright
LSE & UCSD

I. Introduction

Throughout the history of science people have dreamt of one grand consistent scheme that can encompass all possible scientific knowledge. I defend instead the view of my hero, Otto Neurath, one of the founders of the Vienna Circle: "'The system' is the one big scientific lie."[2] I do not argue, however, that there is no possibility for one consistent theory that rules ubiquitously. Rather I argue that

- a 'pluralistic' universe that does not lend itself to description within one consistent scheme is perfectly possible;
- there is a good deal of good evidence in its favour; and
- the evidence on the other side is not so good as it first looks.

My overall conclusion is the old-fashioned positivistic advice: Do not let metaphysical issues like these intrude into scientific practice. Where this is not possible, hedge your bets and hedge them heavily. And if you have to bet, I'd bet against the system.

When I discuss the weight of evidence, my worries are not about what we should *believe* about the metaphysics of our universe. Like Bas van Fraassen I take it that individuals are entitled to leaps of faith in arriving at their beliefs.[3] I am worried instead about what we do. We can maintain freely adopted beliefs about the consistency and unity of our universe; but we should not allow them to play a role that must be reserved for evidence alone. That I take it is what empiricism teaches; and I endorse these empiricist teachings unreservedly.

What I say in my book, The Dappled World,[4] is this:

> The yearning for 'the system' is a powerful one; and the faith that our world must be rational, well ordered through and through, can play a role where only evidence should matter. Our decisions are affected. After the evidence is in, theories that purport … to be able in principle to explain everything of a certain kind often gain additional credibility just for that reason itself. They get an extra dollop of support beyond anything they have earned by their

[1] Research for this paper was supported by the U.S. National Science Foundation and by the Latsis Foundation. Research assistance was provided by Anna Alexandrova under a grant from the University of California at San Diego. I would like to thank all of these for their help.
[2] P. 116, Neurath, O. "Einheit der Wissenschaft als Aufgabe", 1935, *Erkenntnis*, 5, pp.16-22, trans. "The Unity of Science as Task" in Neurath, O. Philosophical Papers 1913-46 ed. and trans. R.S. Cohen and M. Neurath, Dordrecht: Reidel, 1983. See also Cartwright, N., Cat, J., Fleck, L. & T. Uebel Otto Neurath : philosophy between science and politics, New York : Cambridge University Press, 1996.
[3] See his discussion in *Laws and Symmetry* (Oxford : Clarendon, 1989) pp.171-172 and *The Empirical Stance* (New Haven, CT : Yale University Press, 2002) p.77 and pp.81-90.
[4] Cambridge: Cambridge University Press, 1999.

> empirical successes or by the empirically warranted promise of their research programme for solving the problem at hand. (p.17)

Though my conclusions are guarded – we do not have good enough evidence for the possibility of one consistent scientific system to bet on it – you will notice my tone of enthusiasm. That is because I am personally enamoured of the dappled world. I delight in the thought of it. Like Gerard Manley Hopkins, I love "landscape plotted and pieced – fold, fallow and plough"; I am fascinated by "all trades, their gear and tackle and trim." Nevertheless I do not attempt to offer this aesthetic ideal as a philosophical argument. In this paper I propose to review some of the central arguments. My focus will be on physics since it is our closest approach to the ideal system, unified, precise and effective. I shall discuss two distinct aspects of the systems.

The first is the system's scope. Can no system treat everything? I shall argue that there is a trade-off. Our physics can be exact, its applications grounded in principle and its predictions precise. But the way we secure these virtues narrows the scope. Modern physics does not even treat much of what is in its own domain.

My second topic is the consistency between our physics and the world it is meant to describe. In this part of the paper I will discuss the Kantian-inspired views of Michael Friedman, which he describes in this volume, about the principles in physics that make it possible for the system to fit the world. I think he locates these principles in the wrong place, high in the system, in a sense before the theory even begins, whereas I argue they need to be located right on the edge of practice. Again we see a trade-off. The demands of the rational system – grounded in principles, exact, unambiguous and precise – oppose the looseness that is required to make physics fit the world.

II. The All-encompassing System and the Dappled World

II.1 Are Pluralism and Precision at Odds?

The first step in defence of a dappled world is to establish that pluralism is possible. That's a task that seems relatively easy from various social constructivist and relativist standpoints. In this part of my paper aim to defend pluralism from a realist point of view. I do so because that seems to be taking on the hardest case. So I suppose for the sake of argument that a great many of the claims of modern physics theory are true – or at least as true as other claims we find unproblematic. I believe that there is a tendency to suppose (often almost without realizing it) that without a universal rule of order it is impossible to have very exact and precise scientific regularities of the kind that, wearing my realist hat, I want to endorse; i.e., it seems that pockets of precision are impossible.

My basic strategy for showing that pockets of precision are indeed possible in a happenstance world is to look at science itself: I offer an image of the world copied from the methods we use in our most successful studies of it. To account for order – say the regular motions of the planets – we build a model. The model describes a set of components with fixed capacities – here masses with the capacity to attract and be attracted – in a fixed arrangement, operating without external interference. The central ideas are the dual notions of *capacity* and *interference*.

The principles governing the operation of the parts are, I maintain, claims about capacities; for example, the capacity of strength GM_1m_2/r^2 that a mass of size M_1 has to attract another mass of size m_2. They are not themselves further claims to regular associations between occurrent properties. In line with conventional philosophical views about dispositions and powers, I agree that capacity ascriptions support both factual and counterfactual claims about regularities between occurrent properties: e.g., if nothing were to interfere the acceleration of m_2 towards M_1 would be GM_1/r^2. But more than this is true. For certain very nice kinds of capacities, we have learned laws for how they combine, like the law of vector addition for forces or the far more complicated rules for how to compute currents in variously configured electric circuits from the separate capacities of the resistors, impedances and capacitors.

All this, of course, only so long as nothing interferes, where *interference* is a positive descriptive concept. It is not used just to assert that the regularity holds unless it does not. But again, in our experience more than that is true. Often the capacity successfully exerts itself even when interference is rife. Ordinary language has good modal forms for expressing this (though Humeans may revile them). My breakfast cereal box says "Shredded Wheat *may* help keep your heart healthy". And this is in accord with a wealth of day-to-day experience. My daughter and I have dropped the tiny metal back from an earring into a crack between the floor boards. I say to Emily, "Go fetch that red and silver magnet from my desk." "OK," she replies, "the magnet may just lift it."

Perhaps the behaviours resulting from the action of capacities, like the capacity of the magnet to attract metal objects, can all be cashed out in terms of long, complicated regularities. Perhaps not. The point is that we can – and do – perfectly intelligibly use and test these kinds of claims without engaging ourselves with these totally unknown regularities. We do not have to make this big leap of inference, nor do I see where we could get the epistemic warrant for doing so. We do not have to believe that for everything that happens there is some description under which it falls that will bring it under a master regularity. Some things may just occur by hap. Again, this is completely in accord with our everyday experience. The therapy can help, but it may not; and there may be no projectible probabilities to rely on.

I said that capacity and interference are dual notions. I think that is important for understanding how we test capacity claims. But later on, in the context of formal theory, I shall describe a weaker notion of interference not connected with the concept of capacity. Roughly, for a given theory an interference is anything relevant to the behaviour under discussion that cannot be correctly described using the concepts of the theory.

My argument here is exactly like Hume's in his *Dialogues Concerning Natural Religion*. The well-ordered world in which all events occur in accord with some general regularity is perfectly possible given everything we know; but it is not the most natural conclusion to reach from the evidence around us.

II.2 The Anti-Theory Theory of Theory

Larry Sklar, in commenting on *The Dappled World*,[5] offers a clear positive argument for why, despite the appearances around us, we should think that the laws of many of our sciences (fundamental physics in particular) apply everywhere: our theories tell us that they do. I disagree; I don't think our theories do tell us that. Not only do I not think that our best physics theories tell us that they apply everywhere; I wonder if the claim makes sense. That's because it relies on what I think is a misguided idea of theory.

In the heyday of Popper and the Logical Positivists philosophers came to demand of science that it be *exact*: the claims of science must be explicit, unambiguous and precise. This provided us with a weapon to fight the evils of Hegel, religion, Freud, Marx for many, and, hopelessly, Nazis and the restrictions on freedom in Eastern Europe.

The problem is that we are in danger of believing our own propaganda. Let us focus on physics because it is the hard case for my position:

1. Physics is clearly a successful science.
2. Successful science, we maintain, is exact science.
3. <u>So</u> physics is an exact science.

I realise that this direction of implication is not usual. We are often told that physics is the model from which we derive these requirements for exact science. When, for example, economics strives to be an exact science, we are told that economics is trying – rightly or wrongly – to cast itself in the mould of physics. So the more standard way of ordering my three propositions is this: physics is an exact science; physics is clearly a successful science. Therefore (it looks as if) successful science will be exact science. The semantic view of theories contributes to this image. Theory is a collection of models and the theory is correct when the models can be matched to the world. Again, theory is both completely articulated and unambiguous.

I think this is a mistake. So far as I can see, physics is not now, and never was, an exact science in the sense laid out. I don't just draw this conclusion from my own studies of physics, but from those of other philosophers, historians and sociologists. I'll remind you of the conclusions of just a few of the most well-known.

Thomas Kuhn, Larry Laudan and Margaret Morrison[6]*:* Scientist learn a set of techniques for problem solving, for prediction and for application. Claims play a role here – both claims of high theory and more concrete claims about specific situations and specific materials. But they do not function as the kind of explicit, unambiguous propositions we have been taught about, propositions that can be used jointly to derive predictions and conclusions. They function, rather, in very delicate and complicated

[5] "Dappled Theories in a Uniform World" *Philosophy of Science* 2003, v.70, no.2, pp.424-441.
[6] T.S.Kuhn <u>The Structure of Scientific Revolutions</u>, Chicago: University of Chicago Press, 1962; L.Laudan "A Problem-Solving Approach to Scientific Progress", in Ian Hacking (ed.), <u>Scientific Revolutions,</u> Oxford: Oxford University Press, 1987, pp. 144-155; M. Morrison "Models as Autonomous Agents" in M. Morgan and M. Morrison (eds.) <u>Models as Mediators</u>, Cambridge University Press, pp.38-65.

ways, as guides for the constructions of models of the target situation, models from which we can produce predictions.

Harry Collins[7]*:* Much knowledge that is necessary to produce a predictions or to build a physics device – like a measuring instrument or a piece of physics-based technology – is implicit. Even when we do our best to write it all out, we often do not succeed.

Peter Galison[8]*:* Even when we do have a body of explicit claims, they will often be understood very differently by different groups, groups with different backgrounds, different agenda and different skills, using different techniques. Galison takes the problem of univocal understanding to be so serious that he introduces the metaphor of the "trading zone" to explain how agreement is reached between different groups, despite very different understandings of what has been agreed upon: the islanders and the European traders exchange one good for another, though neither understands either good in the same way.

I conclude from considerations like these that there is no such thing as 'the theory' that we can consult to see what the world is like according to our bets scientific knowledge. There are myriads of theories all under the same name – say, "quantum field theory"; or, by standards of exact science, there are none. I do not want to deny the blatant fact that there are some 'standard' axiomatizations on offer, especially in high theory; nor that specific theory groups, or different theory groups talking closely together, can have a univocal object they focus on. I do want to deny that any of these formulations are ever 'the' theory in a given field. For

1. They neither constitute nor imply much of the knowledge, even highly theoretical knowledge, that we have in a given domain. This is an old theme of mine. But I have recently returned to it, in work with Towfic Shomar and Mauricio Suarez.[9] We argue that the propositions prized by theory groups, and many much further down the ladder of abstraction, function as guides for the constructions of models, not as true general propositions exemplified in the resulting models. Mathias Frisch's work on classical electromagnetic theory supports this view, for Frisch shows, among other things, how many of the important models are inconsistent with the central equations.[10] (I shall return to this topic in section III.)
2. Theoretical claims will generally be differently understood by different groups. What is especially important, if Galison's cases turn out to be typical, they may be differently understood by theoretical groups that wish to refine them, experimental groups that aim to test them and applied groups that use them to build new technologies.
3. They may be metaphysical add-ons or over-generalizations that do not play a role in generating predictions. This has been my general point of attack over the years.

---

[7] Changing Order. Replication and Induction in Scientific Practice, London: Sage, 1985.
[8] Image and Logic: a Material Culture of Microphysics, Chicago: Chicago University Press, 1997.
[9] Cartwright, N. Shomar, T, and Suarez, M. 'The Tool Box of Science' in Theories and Models in Scientific Processes. W Herfel, et al eds., Amsterdam: Rodopi, 1995, pp. 137-149.
[10] Cartwright, N. and M. Suarez "In defence of Theories as Tools", unpublished manuscript, 2004.

II.3 Warrant: The Ghost-train of Theory versus the Hard-knocks of Practice

How then do we figure out what the world is like? I agree with Sklar that we should do it "only by consulting our best available science"[11]. But I do not believe that there is a convenient place called 'theory' where that is encoded. I also presuppose a very strong empiricism: it is empirical success that determines what our best available science is. So to figure out what we are warranted in believing we need to find *just those claims that are genuinely used in deriving the predictions and applications that constitute our empirical successes*.[12] This means that it will be much easier to figure out what to believe about a helium-neon laser than it is to find out what to believe about coherent radiation and that is easier to figure out than what we should believe about quantum radiation and so forth.

Figuring out what our empirical successes entitles us to believe then turns out to be a hard job. But why not? As I see it, it is best to think about entitlement to believe at the sharp end, where what we believe makes a difference to what we do and thus to what happens to us. Here, I think, along with a misguided image of theory, we also operate under a misguided image of how theory interacts with confirmation and use. We think of theory as a *ghost train* that can take us from a set of appropriately strong and varied confirming instances a long way, to new predictions and results never before dreamed of. Because the instances confirm the theory, we are entitled to believe in it, and because we are entitled to believe in it, we are entitled to believe in its novel consequences.

But that is not what we do at the sharp end. Nor is it what any rational person would do had they had any alternatives. When it comes to using theory in ways that will affect our lives, we insist on exceedingly small steps. We want the claims we rely on to have proven successful in situations as much like the target situation as possible. Even then, whenever we can, we try to build a prototype, to try out the new predictions before we rely on them.

My point is one about confirmation, where confirmation is really going to have some bite. What we judge we are entitled to reply on in practice is a good guide to what we should take ourselves to have confirmed. And I don't find that that is 'the theory' in the sense of standard axiomatizations or the object that is studied by people called 'theoreticians' or the object that is studied by my comrades in the foundations of physics. That object may be exciting; it may present us with an impressive image of a Platonic world of mathematical objects and their relations; and it may serve as powerful a tool to aid in and guide the construction of precise models that allow us to, say, build lasers and predict precisely what they will do. But it does not have the right relations to these models to inherit upwards the confirmation that successful prediction confers on the models.

So on my view if we want to figure out what is confirmed at the most abstract and general level possible, we are starting out on a long and difficult task that requires a

[11] Sklar 2003, p.424.

[12] E. Winsberg, M. Frisch, K. M. Darling, and A. Fine ("Nancy Cartwright's The Dappled World: A Study of the Boundaries of Science." 2000 *Journal of Philosophy* 97, pp. 403-408) were, I believe, first to point out how much my arguments that our best bet is that the world is dappled depend on a strict view of warrant.

detailed look at *exactly what is put to use across the panoply of our empirical successes.*

The strategy I urge is familiar from recent debates in scientific realism over structuralism. Structuralism is a doctrine of Poincare[13], developed by my LSE colleagues John Worrall and Eli Zahar[14], now taken on by Stephen French, James Ladyman[15] and others. The doctrine is meant to combat the pessimistic metainduction: all great physics theories of the past have been radically mistaken; so our best bet is that so too are our current ones.

Structuralists argue that often our theories have not been as mistaken as we think. For many specific cases, if we look carefully we shall see that the content and overall world vision may have changed, but the form of the equations has not: structure has been preserved. Classical field versus particle theories of light is the canonical example. The content of the theory has changed across scientific revolutions, but the form of the equations for the propagation of light have not. So the empirical success based on this structure give reason for us to believe in *it*, reason not undermined by the dramatic change in content.

Second generation structuralists – notably Stathis Psillos[16] – are even more optimistic. If we trace through the details of how myriads of successful predictions are produced both pre-revolution and post-revolution, we see that lots of theoretical content stays the same as well – just not the content of what I have here called 'the theory'. For instance, both wave and particle theories assume that there are light rays that behave in very much the same way. This though is a crude example. Most of the content we will buy in this way will be highly detailed theoretical content, claims most of us have never heard of since we do not have much call to look at the nitty-gritty of real prediction.

So, if the new structuralists are right, there is a very great deal of modern physics that we can be realist about; a great deal that our successes give us sound epistemic ground accepting for use in decisions that matter. But to find out what it is requires hard excavation work. Personally I do not see why one would want to undertake it. Is light composed of fields or of particles or perhaps of something we have not yet conceived? Is space a container or …? I am personally deeply curious about the answer; but that does not mean that I have to come to a conclusion, to let these matters affect the way I live in and try to change the world.

---

[13] Poincaré, H. (1902) *Science and Hypothesis*, Repr. New York: Dover, 1952.

[14] Worrall, J. (1994) 'How to Remain (Reasonably) Optimistic: Scientific Realism and the "Luminiferous Ether" ', in D. Hull, M. Forbes and R.M. Burian (eds) *PSA 1994*, vol. 1, East Lansing, MI: Philosophy of Science Association. Zahar, E. (1997) 'Poincare's Philosophy of Geometry, or does Geometric Conventionalism Deserve its Name?' *Studies in History and Philosophy of Modern Physics*, 28B(2), pp. 183-218.

[15]French, S. and Ladyman, J. (2003) 'Remodelling Structural Realism: Quantum Physics and the Metaphysics of Structure', *Synthese*, v.136, no. 1, pp. 31-56.

[16] Scientific realism: how science tracks truth, London: Routledge, 1999.

There are, of course, important questions we do need to answer. Which questions they are depends on which level we operate from and what our problems are. I need to form a view about whether the laser will function properly enough before I settle on eye surgery. Laser engineers in Silicon Valley need to form very detailed views about the theory of lasers of specific constructions before they suggest types with expensive design changes; reviewers at the NSF need to form views about the promise of various pieces of proposed research before they divvy out support; and so forth.

Perhaps at some point in its work some research group has to decide on which general propositions to use in constructing their models – those of quantum or of classical physics, those of waves or those of particles. But there is no pressure to form views about particles versus waves *tout court*. There isn't even anyone who needs to take a view about whether a particle versus a field approach is more promising in general for the construction of successful models, let alone people who need to take a view about which set of claims is right.

So we have at least that blessing. Reasonable belief formation about propositions of 'the theories' of modern physics is difficult to achieve, but it is also not called for. So I am not inclined to try to form views about the metaphysical question of the possibility of the one consistent system. But, as I noted, if I had to I would heavily hedge my bets. Not just for our usual epistemic reasons that probably any theories we have will ultimately be found mistaken, but rather because I do not expect that any theory that we would have would be able to inherit upwards the confirmation that successful prediction confers on models constructed using it. Beyond that, though, I think we do have positive, albeit in no way conclusive, evidence that best current theories are theories limited in their domains. I shall present two kinds of considerations that favour a limitation in extent, considerations that I maintain do what Sklar and I both think should be done – "[consult] our best available science".[17]

II.4. Abstract Concepts and Principles that Tie Then to Concrete Descriptions

The first argument about what the best science teaches me I use in *The Dappled World*. The underlying strategy is the one I just described: I look at how laws are used when they are instantiated in models that make successful predictions. We could put the basic thesis without all the 'ifs' and 'buts' this way. The equations of physics all have a specific kind of ceteris paribus clause in front: *So long as nothing relevant occurs that cannot be described within the concepts of the theory, then…* The dramatic example I use is "F=ma": for any object so long as nothing happens that affects its motion *other than things that can be correctly represented as forces*, then its acceleration will equal the force exerted on it times its inertial mass.

This is in line with Patrick Suppes's and Ronald Giere's[18] versions of the semantic view of theories. According to Suppes, a theory is a set-theoretical structure. Consider Newtonian theory: a (proto) Newtonian system is a set of objects and a set of

---

[17] Sklar, 2003, p.424.

[18] Patrick Suppes: Representation and Invariance of Scientific Structures (Stanford, California: CSLI Publications, 2002) and The Structure of Scientific Theories (ed., University of Illinois Press, Illinois, 1977, second edition.). Ronald Giere, Science without Laws. (Chicago: University of Chicago press, 1999).

quantities {f, m, a} on those objects such that for every object, *f* stands for force, *m* for inertial mass, *a* for acceleration and f=ma. So the theory tells us the characteristics of a Newtonian system. It does not tell us which systems, if any, are Newtonian systems. That we find out as we come to identify types of systems for which we have strong evidence that f=ma holds.

Within this kind of framework the big question of course is, what reasons do we have, pro or con, for thinking that the *ceteris paribus* clause does serious work. Do things often happen that cannot be correctly represented by the concepts of the theory? Here is an aid in thinking about the question.

Many of the central concepts in many physics theories are abstract. By *abstract* I mean something very specific: they are only deemed to be correctly applied when some more concrete descriptions apply. Consider force. If we are going to make use of the formula f=ma, we need to write down some expression for *f*. How do we do so? We have a handful of specific descriptions each of which licenses us to write down a specific force function. The principles that express this license are called 'bridge principles' – they bridge between the abstract language of high theory and a more concrete language that gives content to the abstract terms.

The most well-known example of a bridge principle is Newton's law of gravitation, which tells us what form the abstract concept *force* will take – Gm_1m_2/r^2 – when we the more concrete description "one massive object is located a distance *r* from another" holds; the law of electromagnetic attraction and repulsion links the form $\varepsilon_0 q_1 q_2/r^2$ for the force with the concrete description, "a charge is located a distance *r* from a second"; and so forth. When we cannot legitimate a particular force function in this way, we deem the treatment *ad hoc*; and a successful prediction from a model with an *ad hoc* force function does not confirm the theory, nor would the theory give us sufficient reason to rely on it.

So force functions legitimately apply to systems only when they can otherwise be described as masses in the vicinity of other masses or charges in the vicinity of other charges any of the other (more) concrete descriptions available from our bridge principles. Our previous question about the *ceteris paribus* clause can be recast then. How much of the world can be correctly represented in these specific ways. Not much on the face of it. How much 'underneath'? The great physicist Lord Kelvin despised Newtonian mechanics. He thought that the Newtonian models of finite numbers of point masses, rigid rods and springs, in general of inextendable, unbendable stiff things can never simulate much of the soft, continuous, flexible and friction-full world around us.

My own view sides with neither Kelvin nor the Newtonians. I think we do not have very good evidence either way. We have very good reason to think that the planets are a Newtonian system, cannonballs more or less, and so forth. Also that inside the casing of a battery we find a classical Maxwell electromagnetic system, and so forth. But about the vast array of cases that have resisted treatment by this or that theory so far, we just do not know. Nor would I want to generalize from successes at extending our theories into new domains, for I'm sure we have had a vast number more failures than successes. My advice here is what I urged earlier. Don't bet. If you have to, hedge heavily. And remember that the dappled world is every bit as possible as a unified homogenous one.

I used here the example of Newtonian mechanics. But there are a number of other fields that I have looked at that work in the same way: the central concepts are abstract and legitimately apply only where some small set of more concrete descriptions can be applied. This I claim is true of classical quantum mechanics, classical electromagnetic theory, classical and quantum statistical mechanics, quantum field theory, quantum electrodynamics and condensed matter physics. More abstract theories that involve invariances, symmetries and the like, which is where much theoretical works now focuses, all piggyback on these for their application. So the same conclusion will hold of them.

What about other fields? I have not looked in detail elsewhere to see how theory is used in successful application. But that is what I maintain we must do. We are after all interested in the extent of those theories *that we have very good reason to believe are true.* That means theories *as used for successful application and prediction.* What follows from any metaphysical overage that is not empirically confirmed is not a sound basis for belief.

II.5. Does Theory Tell us How Far it Stretches? The Case of Quantum Mechanics

For my second argument about what physics teaches us about its scope I turn to a specific claim of Sklar that our theories of fundamental particles themselves say that they apply to everything that these particles make up. I want to consider what this claim looks like in quantum mechanics – and after all, until we have some major scientific revolution, it is quantum mechanics that will have to provide a treatment of their behaviours and interactions. The rule in quantum mechanics – both in standard axiomatizations, which as I said I don't take to provide very sure evidence about what the world is like, and in practice in the cases I've looked at it is this: If two systems are in states ψ_1 and ψ_2, the composite is in $\psi_1 \times \psi_2$. I have two remarks about this rule.

First, I have no idea about the range of cases over which it is used. It is used in lots and lots of successful model that I have looked at. But is there anything special about the set of cases in which it is successful? This is a question of how far we can stretch our inductions. I do not have, nor think others have, any clear answer. But why make big leaps over little leaps, sweeping inductions over small ones? As I indicated earlier, I do not even understand the drive to do so, let alone the epistemic justification.

Again, I can see lots of situations in which it is reasonable to bet on the rule. Imagine I am very gifted at applying quantum techniques in a particular domain, say some special set of problems concerning superconductors, and a new phenomenon is discovered in that domain. I think I have a good idea how to model it. Should I take it on? If I do so, I know I will be using this particular quantum rule for constructing states of composite quantum systems. Well, perhaps I should. These after all are the skills I have; I couldn't do it any other way. And from all I already know in this domain I feel my idea is a good one.

Then, probably I should go ahead an invest the time in a model that will use this rule. But this justification for doing so neither requires nor uses the big broad generalization that this rule always works. In order for this choice to be rational it doesn't even require compelling confirmation for the belief that it will work in this case. It does require me to have reason to think it does not fail, but how strong that reason must be will depend very much on how heavily I am investing and what the payoffs and losses would be. We can construct different stories for others who might commit themselves to the rule in different circumstances and in different ways. But I never see what we lose by making cautious rather than bold inductions – except of course the possibility of believing in the truth – if it is true – by faith, without strong empirical ground.

The second remark is to focus your attention on the *if: if* two systems are in states ψ_1 and ψ_2, then …. Under what conditions do the little systems that compose bigger ones have quantum states? We know it is very difficult to get systems into quantum states, or at least into known quantum states. For years quantum physicist stressed the importance of preparation and some still do. Willis Lamb (who won the Nobel prize for the Lamb shift and who developed the first quantum theory of the laser) is explicit in his claims that quantum states must be prepared (as I say, either by us or by nature); they only occur in very special circumstances.[19] Is Lamb right? Or does every little thing have a quantum state? You know by now what I will say about this. I think it can serve as the overall conclusion to my remarks about the one big consistent scientific system:

The grounds for sweeping generalizations are always weak. And we don't lose by not overbidding our cards. Quantum mechanics is a fine example. It has impressive empirical successes. This gives us good reason to believe that the quantum models that generate these successes, and the propositions therein, are true. It also gives us good reason to believe that they are likely to work in very similar circumstances to those in which they have worked before. That is the belief we are entitled to work out from in cases where it matters – and happily, it is all the belief we need to get out of science what we would like to.

III. What Allows for Consistency between the System and the Empirical World

III.1. Kant, Friedman and the new synthetic a priori

Laying questions of scope aside, let us now turn to questions of fit: What ensures the consistency of our scientific image of the world with the world itself? Perhaps, following the teachings of the Vienna Circle, we should take that question to be a senseless one. There is no conceivable way to compare our accounts of the world with the world itself, so there is no proper way to answer the question. Even without this rather operationalist approach to what is and is not sensible one might doubt that the world comes, ala Wittgenstein's *Tractatus*, neatly divided into objects and properties in the right way to underwrite the idea of a match between descriptions and facts.

[19] Willis E. Lamb, Jr. 1955 Nobel Prize Address, *Science* 123 (1956).

This worry about talking about the consistency of our scientific image with the world itself certainly accords with a concern expressed throughout this volume, that we cannot sensibly talk about the consistency of one thing with another but rather must consider the consistency of the two under a specific description of each, or as modelled in specific ways. I am happy then to abandon the question in this form and repose it in terms more acceptable: Is our scientific image of the world consistent with the world as we construct our experience of it? Indeed I agree that we should do so. This however makes my criticisms of the Kantian approach easier.

Kant thought that reality is preadapted to our forms of cognition in so far as our forms of cognition require a rational structure and that our experience of the world must already be in the terms and patterns of a rational structure. Perhaps so, if we have a loose enough sense of 'pattern' and of 'rational structure". The concepts with which we describe the world as experienced are indeed repeatable: they apply again and again, in a large variety of loosely systematic ways and there is a good deal of what Hugh Mellor calls 'connectivity' among them.[20] But that is a far cry from the rational structure of our advanced mathematical physics. I have long argued that the rational structures of contemporary mathematical physics do not fit the world as experienced.[21] They may fit loosely, in limited domains, if we do not look too closely. But there is a trade-off: the more accurate we wish our accounts to be to what occurs in experience, the less they will fit into the rational structure – the System – of modern physics. And this is true even if we allow highly sophisticated concepts into our descriptions of experience.

For the Kantian, part of the job of ensuring that our empirical knowledge fits the world of experience is done by the synthetic a priori. This provides the rational framework within which we experience the world. As Michael Friedman puts it, "synthetic a priori knowledge (typified by geometry and mechanics) …functions as the presupposition or condition of possibility of all properly empirical knowledge." (*Dynamics of Reason*, 1999, Stanford, California: CSLI Publications, p. 26) Michael Friedman is keen to resurrect the role of the synthetic a priori, but not as a once-and-for-all framework necessary for empirical experience. Rather each proper theory in modern physics has its own framework that is held, relative to it, as a priori and that makes possible the genuinely empirical knowledge within that theory.

One of Friedman's principal examples are the three laws of Newtonian mechanics, which are a priori in his sense in the Newtonian scheme as currently understood. The law of universal gravitation – "that there is a force of attraction or approach, directly proportional to the two masses and inversely proportional to the square of the distance between them, between any two pieces of matter in the universe" (*Dynamics*, p. 36) – is the one empirical law in the scheme. This, he points out, talks about *acceleration*. Newton defined acceleration relative to absolute space. Since we do not believe in absolute space, we cannot do this. We say rather that the law of universal gravitation holds in any *inertial frame*, "where an inertial frame of reference is simply one in which the Newtonian laws hold (the center of mass frame of the solar system, for example, is a very close approximation to such a frame." (p. 36) This is why in our current rendering of Newtonian theory Newton's three laws must be taken as a priori:

---

[20] D.H. Mellor "Connectivity, Chance and Ignorance." *British Journal for the Philosophy of Science* (November 1967), 18(3), pp. 235-238.
[21] Cf. How the Laws of Physics Lie Oxford : Clarendon Press, 1983.

> It follows that without the Newtonian laws of mechanics the law of universal gravitation would not even make empirical sense, let alone give a correct account of the empirical phenomena. For the concept of universal acceleration that figures essentially in this law would then have no empirical meaning or application: we would simply have no idea what the relevant frame of reference might be in relation to which such accelerations are defined. (p. 36)

In calling these laws *a priori* Friedman points to a difference he claims in our empirical warrant for them. These are not testable in the same sense in which the law of universal gravitation is because they must already be supposed in order for the concepts in the law of gravitation to be empirically meaningful, which is surely a precondition for being testable. I have two comments on this doctrine.

First if he is right it makes a huge problem for those of us who feel that where the stakes are at all high we should only use laws that have survived severe empirical tests in similar applications. On Friedman's scheme the law of gravity can be tested – and that's a good thing since this is an important law used daily in millions of applications. But Newton's second law, $f = ma$, cannot be tested; it functions as a constitutive principle to give meaning to the concept of acceleration due to gravity. Yet $f = ma$ is at the heart of the applications of Newtonian mechanics. It is how we calculate motions, which are the object of study of mechanics. Moreover, the law we can test is almost useless unless it is deployed in conjunction with the second law. So Friedman's scheme seems to be just back-to-front with respect to empirical warrant: the second law, which we can hardly do without in putting Newtonian mechanics to work, is untested, but the law of gravity, which helps us calculate one universal but often insignificant component of the force function needed for the second law, can be severely tested.[22]

This kind of situation was not so bad for Kant, whose elaborate scheme was meant to show that synthetic a priori knowledge is indeed knowledge. But Friedman has no such guarantees. There are of course reasons for adopting a given set of constitutive principles and some constitutive principles are clearly better than others. But on Friedman's view the reasons for the far more widely used second law must be markedly less good than those for using the special-case law of gravity since the law

---

[22] The light principle and the equivalence principle are constitutive principles for GTR according to Friedman. In defending the claim that these are not testable in GTR he points out that the putative tests can have, indeed have had, alternative interpretations. But that seems to be true of any test, whether of a law he counts as constitutive or of one he counts as empirical. He also sometimes speaks of them not being tested *as playing the role they do* in GTR. This remark may hold the answer to my worries. But so far as I understand it, I would dispute it. When we discuss warrant, it is application where it matters, and here I would argue that a responsible view demands that we have warrant, as much a possible, for exactly what is presupposed in the application. If that includes something about the general role of the principle in the theory, then we had better have warrant for that. (For instance, in my own work I treat the law of gravity as an ascription of a capacity and I point out that we regularly test our claims about the strength of this capacity that the law describes. It is, however, a far more difficult thing to test the claim that it is *capacity*. i.e. that what masses do in attracting other masses when no other forces are acting (the conventional test situation) is the same as what they will contribute to the total force when other forces are present. Nevertheless, this claim too must be tested if it is to be used, whether or not it serves to show how the law is treated in the theory.

of gravity benefits from all the reasons in favour of the Newtonian system and in addition it has been severely tested.[23]

My second comment is that the need for the kind of definitions that Friedman's constitutive principles underwrite is generated by his vision of the theories in question as rational structures – that is, as a small set of consistent, unambiguous claims (plus their rational consequences) involving precise, well-defined and unambiguous concepts – the system. So, for instance, *acceleration* needs a once-for-all characterization, a single characterization that we cling to across all applications and all tests of the law of gravity; from that follows the need for one single reference frame that is always referred to when the law of gravity comes into play.

But that does not seem to be how it happens. Each use of the concept of acceleration is necessarily tied to a specific frame within which the acceleration is supposed to be measured. This is naturally different when one computes the trajectory of a cannonball from when one computes the trajectory of an electron being bumped up along the Stanford linear accelerator, and different again in every real circumstance. On the view that sees one concept of acceleration defined in a consistent, unified rational structure, the use of each specific occasion-dependent frame requires a series of justifying assumptions. First, acceleration is defined relative to those frames in which Newton's laws are true – of which there aren't any. Next we suppose that defining acceleration relative to the frame of the fixed stars will give a good enough approximation to the values the acceleration would have were it measured in a (non-existent) frame in which Newton's laws are true.[24] Then we suppose, on each occasion, that the frame we have picked for that occasion should yield values for the acceleration that are close enough to those yielded in the frame of the fixed stars.

I think instead we only make, only need to make and are generally only justified in making one assumption for each occasion of use: in the frame we choose on that occasion, Newton's laws are true enough to support the conclusions we want to draw. That assumption I take it will almost always be more secure than the series of assumptions needed for use of the univocal concept of acceleration required by the rational structure. And where warrant matters, it is to be hoped that it will be supported by reasons far stronger than those supporting that series.[25] Among these – if we are to trust to our conclusions – had better be a great deal of experience in similar situations and also enough experience to argue that those situations are

---

[23] Nor does he want guarantees that they are genuine knowledge. Friedman explicitly separates his views from scientific realism; what he looks for is a historical convergence of rational community opinion. Nevertheless we should have warrant for our views, especially, I wish to stress, warrant for their use. Perhaps Kant's synthetic a priori, being something without which we cannot perceive or reason, did have a claim to warrant. The question still remains for Friedman as to how it can be acceptable that our more widely used principles are less warranted than those more narrowly used.

[24] Note that it is not good enough to assume simply that in the frame of the fixed stars Newton's laws are approximately true, since the closeness of the three laws to the truth is not enough, without further assumption, to guarantee that the values of any particular quantity in the two frames will be approximately the same. The closeness of the first approximation does not guarantee that of the second without more ado.

[25] Of course some of the reasons we might try to offer at the last step of the series might, cast somewhat differently, be much the same as some of the reasons we give for the immediate assumption that Newton's laws will be well enough satisfied. For instance, that we are in a frame whose rotation will have negligible (for these purposes) effect, an assumption that can be made without having to fix the rotation to the fixed stars.

relevantly similar and to do so without resorting to a highly abstract label that catches, or seems to catch, a lot of diverse phenomena by being thin in content.[26]

III.2. The Rational System: Instrument versus Description

Friedman explicitly contrasts his view with what he describes as an *instrumental* view of theories. I am not sure what he means by 'instrumental' but surely my anti-theory theory of theories falls under this description. Theories for the most part do not make claims about empirical reality. They may in their core principles – like Newton's laws and the law of universal gravitation – take the linguistic form of claims. But the claims that constitute the rational structure of physics describe only a world of our imagination, a fictional world that, when interpreted in the right ways can provide a rough template (to use a concept of William Wimsatt's[27]) for organizing bits and pieces of the real world. When it comes to empirical reality, this rational structure is best viewed as a tool for constructing the million-and-one genuine claims that make up the body of our physics and engineering knowledge.[28]

In so far as we should think of physics theories as making claims, the claims – or, the *warranted* claims – are severally complex and not simple to state even given a mix of various theoretical and practical languages (say those of quantum and classical physics and of laser engineering) and of mathematical forms of representation. They constitute a vast interrelated network – a tangle – with hugely overlapping family resemblances. These resemblances depend on shared vocabulary; shared basic structural forms, especially those provided by what are usually identified as 'the fundamental laws' of the theory; shared techniques; and analogies not only of form but of content. But the claims do not fall properly under the 'fundamental laws' that constitute the rational structure, nor are the same terms understood in the same way

This is what we see when we look at what happens when physics successfully treats the world, and it is a huge stretch – a leap of faith, I would say – to think that somehow our successful treatments could be made to fit the rational structure if only we knew more, could calculate better… (Here we see the idea that the failures of fit are all due to post-lapsarian sinfulness rather than to genuine diversity in the world.) The picture I paint is far closer to what I see whenever I look at physics theories in use. This matters because it is only the theory as used and tested that can be counted as warranted.

III.3. Warrant (Again) and the Claims of Theory

I speak about the *warranted* claims of physics; it is the demand for warrant, stringent warrant where stakes are high, that drives my image of theory in physics. Recall the conventional view about warrant for claims in the theory-net what I called the *ghost-*

[26] "Force" for instance is such an abstract label. It gets (more) concrete content only when it is fleshed in with a description of what constitutes the source of the force, like a particular arrangement of charges or the presence of a massive planet. See 4.1 below.
[27] "False Models as Means to Truer Theories" in M. Nitecki, and A. Hoffman, eds., Neutral Models in Biology, London: Oxford University Press, 1987, pp. 23-55.
[28] For a somewhat lengthier discussion, see Cartwright, Suarez and Shomar 1995.

train view: We have inductive evidence of various kinds for the principles of the theory. The abstract language in which these principles are cast carries the warrant silently and swiftly to new cases far away and far different from the original inductive base. This seems to me a dangerous picture and one that we do not take seriously when there is much to lose if our verdicts about the new case are wrong.

Consider a well-known historical case. Newton insisted that mass is mass, whether terrestrial or celestial. He thereby allowed us to apply the same abstract concept to both. This was a huge conceptual breakthrough that provided the tools – the concepts, the basic equation forms, the mathematical techniques and so forth – for treating planets and cannonballs in analogous ways. But calculations of the trajectories of cannonballs are hardly warranted by observations on billiard balls, let alone by observations on the locations of the planets. If it really matters to us where our cannonballs will land, our confidence in the calculations for their trajectories should be based on a great number of successes, as well as our failures and our corrections for them, on vast numbers of trials. It is only then that we have good (enough) reason to suppose that *this* is a kind of case where Newton's laws and Newton's concepts provide a central tool and that we have made the right use of it.

My view on warrant goes along with my earlier arguments about the choice of frames of reference. If we are going to use Newton's law then we need good empirical evidence that they are accurate enough in the frame we have chosen, and accurate under just the specific detailed interpretation we are giving them in this case. The kind of empirical evidence we need, and the kind we generally insist on for practical belief, belief where it matters, is evidence in situations very like the one to hand, with very much the same kind of claim at stake.

I do agree with the conventional picture that warrant is carried from old cases to new cases by inductive generalizations. But I want to keep firmly in view that small inductions are more secure than grander ones and detailed local experience – trial and error, with both successes and failures – really matters.

One of the results of these truisms about induction is that we do not really know what is in the theory-net – what is a warranted claim of the theory – and what is not. Knowledge of what constitutes the theory does not lie with the theorists. Instead it is dispersed among the vast number of those practitioners who put the theory to use to model concrete situations. Only those who grapple with the attempt to construct local models will know which are warranted and which are not, hence severally what are and are not acceptable claims of the theory. This is a far different story from that of the rational structure, the axiomatizable or even axiomatized structure, that is there for all with the requisite mathematical skills to know. What the theory says is no easy matter to find out.

My own image of our attempts to carry warrant in the conventional way is very different from transport via the ghost-train. First, the inductive climb up is insecure; for our best theories there are a lot of routes, but none of them is all that stable or secure. Then the route across even within the theory itself can be shaky. And getting down again to the concrete details of the new case is usually far more insecure and slippery even than the inductive climb up into the theory.

Defenders of the System may seize on this picture in their own defence. It is, or so many say, not the rational structure that is insecure but rather the links to more concrete, more usable, descriptions of situations. This is a response that I do not understand. When the links are insecure – we are not sure of them, we make do, we need to make *ad hoc* adjustments, approximations, improvements, often even just guesses – then so too is the conduit for warrant. The principles of the System cannot be warranted by an induction on behaviours that are not clearly known to fall under them. The situation is even worse for warranting claims that we are not really sure follow from this System.

Friedman is explicit that he is not interested in realism. But the disagreement between us about what theory is has a familiar form from the realism debate. The philosophical thesis that the rational System is useful in just the way we find it to be in constructing claims that can be highly warranted is weaker than the thesis that the rational System itself makes highly warranted claims, and it has all the same evidence in its favour without having to make excuses that lack solid evidence for many and pervasive failures.

III.4. An alternative locus for constitutive principles

According to Friedman constitutive principles, like those defining the frame of reference of the concept of acceleration, make the "empirical application of the theories in question first possible".[29](*Dynamics*, p. 49) I think this is a misdescription. The principles that Friedman calls 'constitutive' make the concepts *intelligible*, not *empirically applicable*. They provide univocal and precise definitions that fit the concepts into the relevant rational structure but they are not the principles that make possible the empirical application of these concepts.

When it come to the presuppositions for empirical knowledge using the concepts of physics, as opposed to presuppositions for fitting these concepts into the rational structure, I think there are three better candidates:

Bridge principles. Recall, many concepts in physics are abstract in a very specific sense: given the way the theory itself works, these concepts never apply to an empirical situation except via some one or another of a collection of more concrete descriptions. It is the bridge principles that spell out the link between the two. These principles provide empirical content to concepts that are otherwise free-floating, that are otherwise given sense only by their relations to other equally abstract concepts laid out in the laws they figure in. so bridge principles do just what Friedman requires of constitutive principles: they make the empirical application of the abstract laws possible.

We should note though that this view turns Friedman's picture of Newton's theory on its head. Remember from section I.3 that the law of gravitation is a prime example of a bridge principle. It tells us what form the abstract concept *force* will take – Gm_1m_2/r^2 – when the more concrete description "one massive object is located a

[29] This of course, as he points out, does not guarantee that the empirical principles that we formulate using them will be true, just that they are candidates for truth or falsity.

distance *r* from another" holds. This is just opposite to Friedman's account. He takes Newton's three laws to be constitutive, the law of gravitation as empirical. I take the three laws to be empirical and the law of gravitation to provide empirical content for them.[30]

Measurement auxiliaries. When we tie our physics concepts to empirical reality by measuring them in a specific circumstance, our procedures presuppose the accuracy of a host of laws about the situation that justify the measurement techniques. This is already a much-discussed topic in science studies and I won't pursue it here, except for one caveat.

One might want to argue about measurement auxiliaries that they are not constitutive in Friedman's sense: they are not presuppositions whose correctness is necessary for the *applicability* of the concepts in question, but only for our ability, case-by-case, to measure them. The concepts still apply, but we would be very limited in our ability to measure them. This contrasts with the case of bridge principles since, without bridge principles abstract concepts – i.e. concepts that require more concrete descriptions to hold – would not apply to empirical reality but only have a place in the rational structure. The third candidate, to which I now turn, also has more claim to the title "constitutive" than do measurement auxiliaries.

Representation theorems. We represent features of the empirical world with specific mathematical forms that have specific properties. These forms tend to be far more universal across applications than is any (univocal) interpretation or definition of the related concept. *Acceleration* for instance: no matter what frame of reference we define it relative to, we almost always represent it as d^2x/dt^2. So length itself must be represented as a quantity twice-differentiable with respect to time. But it also has a number of other built-in features as well. Probably the simplest is that length is represented by an additive measure. Can a mathematical representation with these characteristics adequately represent the phenomena to be associated with "length"?

The answer depends on the structure of the phenomena to be represented. In the case of length, the phenomena might include what happens to sets of measuring rods. For instance, as Patrick Suppes puts it[31], the collection A of rods is longer than B "if and only if the set A of rods, when laid end-to-end in a straight line, is judged longer than the set B of rods also so laid out."[32] Formalizing that fact[33], along with a couple of other obvious features we ascribe to the empirical concept of length (for instance, that any collection of rods is at least as long as the empty set), we can characterize the structure consisting of the set of rods and the longer-than relation as a *finite equally-spaced extensive structure*.

[30] Note though that the more concrete vocabulary that fits out the abstract concepts of high theory is not all the vocabulary available for the formulation of empirical knowledge about the systems that physics helps treat. It does not, for instance, include names of materials nor much of the many technical vocabularies of engineering and technology. It should be no surprise that the most accurate expressions of empirical knowledge are in a rich mixed vocabulary.

[31] *Representation and Invariance of Scientific Structures,* Stanford, California: CSLI Publications, 2002.
[32] P. 64
[33] Recall again our concerns throughout this volume that it only makes sense to ask if two things are consistent once we have a model of both.

Now we are in a position to show that an additive measure is an appropriate representation for length by proving a *representation theorem.* In this case the theorem tells us that for any finite equally-spaced extensive structure, there is an additive measure μ such that for every pair of sets of rods, A and B, $\mu(A) \geq \mu(B)$ iff A is longer than B. That is only a start of course. In order to guarantee the empirical applicability of the concept of length as we represent it in mathematical physics, we need a representation theorem relating all the qualitative features we assign to length to its mathematical representation. And similarly for all the quantities of empirical reality and their features for which we provide mathematical representations.

It is, then, I urge, in the representation theorems for the mathematical representations we offer in physics that we find our best candidates for "constitutive principles". This is true for both my image of the theory-net (or "tangle") as well as for Friedman's image of the rational structure of theory. Representation theorems are the preconditions for the preconditions for the application of our concepts to empirical reality. Our representations are consistent with the features we ascribe to empirical reality only if the appropriate representation theorems are true.

IV. Conclusion

Rationality requires consistency within the system. Warrant for applying a system requires consistency between it and the empirical world. Kant thought that reality is pre-adapted to our forms of cognition. As an empiricist looking at the successful and unsuccessful tests and applications of contemporary mathematical physics, I do not find that claim well supported. We have, I am afraid, two demands for consistency that pull in opposite directions.

The same is true when it comes to scope. We want the laws of our system to be exact and exceptionless and to make very narrow, precise predictions. In physics we have happily been able to find a set of concepts that we can deploy to make precise predictions. But even allowing for the kind of ambiguity that the adherents of the rational System abhor, these concepts do not even (to all appearances at least) describe all of the causes of the effects that physics studies. Again we have a trade-off. The system can consist of a small coherent set of exact claims but its scope seems severely constrained. Consistency internal to the System is bought at the cost of both scope and fit.