

## Against 'The System'

Nancy Cartwright  
LSE & UCSD

Again and again 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." I do not argue, however, that there is no possibility for one consistent theory that can characterize everything in the universe around us. 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 get let metaphysical issues like these intrude into our scientific practices. 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. 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 in science 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**, 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 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!

The first project 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. I aim to defend pluralism from a realist point of view

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 science 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.

Larry Sklar, in commenting on *The Dappled World*, 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. So 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 which 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*: 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 ways, as guides for the constructions of models of the target situation, models from which we can produce predictions.

*Harry Collins:* 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:* 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 best 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. 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.
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.

How do we figure out what the world is like? I agree with Sklar that we should “rely upon what our best available science provides for us”. 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*. 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 swift vehicle 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 rely 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 we are confirmed at the most abstract and general level possible, we are starting out on a long and difficult task that requires a 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, developed by my LSE colleagues John Worrall and Eli Zahar, now taken on by Stephen French, James Ladyman 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 vs 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 – 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.

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 vs waves tout court. There isn’t even anyone who needs to take a view about whether a particle vs 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 – “rely upon what our best available science provides for us”.

The first I use in the book. 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 the line with Patrick Suppes’s and Ron Giere’s 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 quantities  $\{f, m, a\}$  on those objects s.t. 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 cp 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. Here are familiar examples. .... 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 .... or .... or ....or.... . Our previous question about the cp clause can be recast then. How much of the world can be correctly represented as .... . 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 lots of kinds 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 q.m, classical electromagnetic theory, classical and quantum statistical mechanics, qft, qed, and condensed matter physics. More abstract theories that involve invariances, symmetries and the like, which is where much theoretical works now focusses, all piggyback on these for their application. So the same conclusion will hold of them. For other fields, I have not looked 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.

For my second argument 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 qm – and after all, until we have some major scientific revolution, it is qm that will have to provide a treatment of their behaviours and interactions. The rule in qm – 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  $\psi_1$  and  $\psi_2$ , 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 and invest the time in a model that will necessarily 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 be committing 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  $\psi_1$  and  $\psi_2$ , 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. 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.