Ratio (new series) XVIII 4 December 2005 0034–0006

PHYSICAL REALISM

Brian Ellis

Abstract
Physical realism is the thesis that the world is more or less as present-day physical theory says it is, i.e. a mind-independent reality, that consists fundamentally of physical objects that have causal powers, are located in space and time, belong to natural kinds, and interact causally with each other in various natural kinds of ways. It is thus a modern form of physicalism that takes due account of the natural kinds structure of the world. It is a thesis that many present-day scientific realists would surely accept. Indeed, some might say that physical realism just is scientific realism, but under another name. However, the argument that is presented for physical realism is not the standard one for scientific realism. It is not a two-stage argument from the success of science to the truth of scientific theories to the reality of the entities postulated in these theories. It is more powerful than this, because it is more direct, and its premisses are more secure. It is more direct, because it develops what is basically a physicalist ontology as the only plausible metaphysical explanation of the new scientific image of the world. It is more secure, in that it does not depend, as the standard argument does, on any doubtful generalisations about the nature or role of scientific theory.

1. Scientific realism

I suppose scientific realism to be a thesis about the nature of reality. It is, therefore, primarily a metaphysical thesis. Nevertheless, there is a philosophical programme known as 'scientific realism', which is as much about the nature and role of scientific theory, and the epistemic status and semantic implications of its laws and theories, as it is about metaphysics. This paper is critical of the programme, but not primarily of the thesis. I call my own version of the thesis 'physical realism', because the metaphysical thesis that I wish to defend is a sophisticated physicalist one that is inspired more by 1960s physicalism and the new essentialism, than by the programme of scientific realism.

The programme is admirably presented and discussed in Stathis Psillos's (1999) book on the subject. This book presents an overall perspective on scientific realism that is comprehensive, fair and lucid. Its major defect is one that it shares with most other justifications of scientific realism, viz. that it presents the case for realism as a two-stage argument from the empirical success of science, to the truth, or approximate truth, of its dominant theories, to the reality of the things and processes that these theories appear to describe. Formally this argument would be sound, if one had an adequate theory of truth to carry the metaphysical burden. But no such theory of truth is developed in Psillos's book (or anywhere else, to my knowledge), and one is left to speculate on what ontology might be implied by the truth of science's well-established theories.

Rather than try to develop such a theory of truth here, I shall tackle the metaphysical task directly – my aim being to give an account of the nature of reality that will adequately explain why science has been able to construct the scientific image that it has. In presenting my argument, I shall presuppose no concept of truth other than that of epistemic rightness (as developed in Ellis 1990). For it is useful to separate the epistemic issues from the metaphysical ones. To address the epistemic issue of whether a proposition is true or not, we have only to consider whether the grounds for claiming it to be true are sufficient to justify the claim. If the proposition is an empirical one, then we might consider the empirical evidence that is available to decide whether or not it puts the issue beyond doubt. If it is a proposition of mathematics, then we might consider its proof, or, if that is beyond us, the merits of the claim that it has been conclusively proven. The epistemic concept of truth is thus straightforwardly applicable to propositions of all sorts, whatever their field, and we do not need a metaphysical theory of what would make a proposition of a given kind true or false in order to use it. For truth, in the minimalist sense of epistemic rightness, is metaphysically neutral.

I take it that the following propositions are all true in the minimalist epistemic sense:

(a) Sugar is soluble in water
(b) It is impossible to produce a perpetual motion machine of the second kind

(c) To a first approximation, every body attracts every other body in the universe with a force that is inversely proportional to the square of the distance between them
(d) The speed of light is the same with respect to all inertial systems
(e) $e^{i\theta} = cos\theta + isin\theta$
(f) There are just five regular polyhedra
(g) The efficiency of a Carnot engine working between the Absolute temperatures T_2 and T_1 is $(T_2 - T_1)/T_2$
(h) Ideal markets ensure Pareto-optimality
(i) Paracetamol relieves pain
(j) The subject is thinking about a horse

But none of these statements is metaphysically transparent. For it is not clear what ontology is required to accommodate them. Are there any real causal powers, necessities, forces, or inertial systems in the world? What do mathematical theorems describe? What are the truthmakers of ideal theories? To what categories do mental events and states belong?

I shall not try to answer all of these questions here. My more modest aim is to argue that scientific realists need to be able to answer such questions, if their position is to be tenable. It is not enough for them to argue that the established laws and theories of science are mostly true, or at least approximately true, as though this were the end of the matter. In my view, this is just the beginning. The real work has yet to be done in spelling out the metaphysical implications of this conclusion.

Psillos's account, like that of so many others, leaves the real work to others. For the common assumption seems to be that the principal worry for scientific realists is the question of whether established scientific theories are true. Certainly, the main challenge to scientific realism has come from this direction. But, even if the question of the truth of established theories could be settled decisively in favour of scientific realism, the ontological question would remain. For the correspondence theory of truth on which realists usually rely is far too weak and indecisive to carry the metaphysical burden of the argument for scientific realism. In putting forward his semantic thesis, Psillos advises us to take scientific theories 'at face value', and see them as 'truth-conditioned descriptions of their intended domains, both observable and unobservable'. But how, I ask, is one supposed to do this? I have no trouble in taking 'The cat is on the mat' as a truth-conditioned assertion about the domestic scene at home. Nor am I puzzled about how to understand the statement 'Oxygen has atomic number eight.' The ontological implications of these claims seem clear enough. But many, if not most, scientific claims are metaphysically much more obscure than these, and one cannot just read their ontological commitments off the page.

Originally, a scientific realist was simply one who believed the world to be more or less as the scientific image implies it is. Scientific realists, such as Smart (1963) and Sellars (1963), believed it to be a mind-independent reality, whose content and structure is gradually, if imperfectly, being revealed to us by the methods of empirical science. No scientific realist at any time has believed that the world is exactly as current science depicts, because they would all concede that some of our currently accepted theories are bound to be superseded, and lead to significant changes in our beliefs about the world. Nevertheless, scientific realists have generally been persuaded that such changes are now unlikely to lead to the wholesale rejection of the current scientific image. Some big changes may occur, they concede, but they think that radical changes, such as those that have historically occurred in astronomy, dynamics, chemistry, heat theory, geology, and biology, are unlikely to recur in any of these well-established areas. And no scientific realist believes that the very existence of a mind-independent reality is ever likely to be seriously challenged by scientific advances.

But even if this optimism should prove to be ill founded, and profound scientific revolutions were later to occur in most fields, belief in the ontology that seems to be required for current science might still be the most rational metaphysical position. For such a metaphysic would at least be the best explanation that is currently available of the empirical successes and failures of science, even if it were to prove to be inadequate for ultimate science.

Most scientific realists of the 1960s would probably have called themselves 'physicalists'. For the physicalists of this period (e.g., Smart 1963; Sellars 1963) were those who believed the world to be essentially a physical world that is more or less as the physicists of that era believed it to be – and nothing more besides. Thus, they accepted realism about the scientific image, and combined it with physical reductionism. Consequently, the early physicalists rejected the view that there are any essentially mental events or processes, i.e. events or processes occurring in people's minds, that could not, even in principle, be reduced to physical ones. The world, they argued, is really just a physical world, and all mental events are really just physical events. This position seems to me to be basically correct – but only as far as it goes. The physicalism of the 1960s needs to be supplemented in various ways to account for causal laws, natural necessities (and impossibilities), and truths about natural kinds, properties, and relationships. For without a much richer ontology than Smart or Sellars ever envisaged, many of the various kinds of truths encountered in the sciences would lack truthmakers. A richer ontology of the kind required was argued for in Ellis (1987), and developed in *Scientific Essentialism* (Ellis 2001).

2. The new scientific image

The new scientific image of the world is an elaboration and development of the old one. Like the old one, the new scientific image is of a basically *physical world*. It is a world in which all objects are really physical objects, all events and processes are physical, and in which physical objects can have only physical properties (Ellis 1976). To elaborate this position, I now propose the following definitions:

1. A *physical object* is anything that has energy, or consists of things that have energy.
2. A *physical event* is any change of energy distribution in the universe.
3. A *physical process* is any causally or inertially connected sequence of physical events
4. A *physical property* is any real property, possession of which would make a difference to at least one kind of physical process involving that object.

I think that most physicalists of the 1960s would have accepted this picture, or something very like it.

The new scientific image differs from the old one in a number of important respects. Firstly, it embraces the idea that the world is a highly structured reality (Ellis 2001, 2002), in which there are objective hierarchies of distinct kinds of things in each of the two principal categories of existence.

1. There is an objective hierarchy of distinct kinds of objects or substances generated by the species relation. This hierarchy is nowhere more evident than in the field of chemistry. For there are known to be literally tens, if not hundreds, of thousands of categorically distinct kinds of chemical substances, with laws at all levels of generality relating to them.
2. There is an objective hierarchy of distinct kinds of events or processes. For example, there is a hierarchy of kinds of causal processes that is formally like the hierarchy of kinds of substances, since every different chemical equation describes a distinct kind of process, and the kinds of processes that may be described are evidently related as species in a hierarchy.

Secondly, the instances of the kinds in these two hierarchies all have intrinsic physical properties in virtue of which they are things of the kinds they are. Consequently, the new scientific image of the world implies that reality has a definite modal structure (Ellis 2001, 2002). For the natural properties that exist in the world must be supposed to include a range of causal powers, and, realism about these causal powers should be as much part of the new scientific image of the world as realism about the physical objects or events that possess them. Consequently, the new scientific image of reality is not that of a Humean world of logically independent events, as Smart's image of reality undoubtedly was, but rather that of a world of objects of categorically distinct kinds necessarily involved in the natural events and processes that are the appropriate displays of their genuine causal powers in the contingent circumstances of their existence.

3. Do the laws and theories of science truly describe reality?

According to the proposed definitions of the physical, the category of physical objects includes all of the fundamental particles and fields of the kinds that are recognised in physics, all of the atoms and molecules that chemists talk about, all of the cells and organisms of biology; in short, just about all of the things that most scientific realists think they should believe in. However, the category of physical objects does not include any of the Platonic objects of mathematics, logic, or modal semantics. Nor does it include the idealised objects of abstract model theories. It does not, for example, include geometrical points, perfect gases, perfectly reversible heat engines, inertial systems, or ideal incompressible fluids in steady flow in uniform gravitational fields, even though there are laws of physics that seemingly tell us about such things. Nor does the category of physical objects include Newtonian extrinsic forces, although there are laws of action, combination and distribution of such forces. In my view, these things are not real, and the propositions that are supposed to be true of them need not be even approximately true of real things. Therefore, a scientific realist should not, in my view, be required to believe in them.

Yet the argument from the success of science, to the approximate truth of its dominant theories, to the reality of the theoretical entities seemingly referred to in these theories, would appear to lead to a different conclusion. For it seems to require belief in the Platonic entities of mathematics, the theoretical entities of abstract model theories, and the forces of Newtonian physics. There are, of course, many scientific realists who do believe in some or all of these things. But a physical realist has no good cause to do so.

There is probably no more successful or well-established branch of scientific knowledge than arithmetic, which I take to be the theory of numerical relationships. But scientific realists are surely not required to believe in the reality of any abstract particulars, such as numbers. There are, it is true, some scientific realists, who do believe in numbers, e.g. those who rely what I call 'the strong argument for scientific realism' (see my 1992 critical notice of Bigelow and Pargetter's *Science and Necessity*) to argue from the predictive success of number theory to the conclusion that numbers exist. Nevertheless, I do not think that we should accept them. For there is no plausible ontology that would accommodate them. Similarly, there is the case of geometry. The theories of spatial and spatiotemporal relationships are certainly of fundamental importance in physics. And such relationships undoubtedly exist. But the primary theoretical entities of the geometry, viz. geometrical points, are neither physical objects, nor universals, nor members of any other plausible ontological category. Therefore, we should not have to believe in them.

Geometrical points in space are not to be confused with the mass-points of Newtonian gravitational theory, which are abstract entities of a different kind. For Newtonian mass-points are supposed to be located in space, rather than be elements of it. As such, they are supposed to be capable of moving about in space, and consequently of occupying different geometrical points at different times. But mass-points are not physical objects either. Nor are they members of any other acceptable ontological category. Rather, they are just the idealised objects of an abstract model theory. It is true that mass-points are supposed to have mass, and therefore energy. But mass-points cannot be accepted as physical entities any more than geometrical points can be. For mass-points exist only as abstract representations of physical objects, idealised as having point-like locations, and postulated as having masses so that they can, theoretically at least, interact with each other. Of course, a theory that employs such abstract entities is not a realistic one. Presumably, it was never intended to be. Nevertheless, physicists, who wish to understand how clouds of real particles behave, might do very well to consider how clouds of point-masses would behave. This is just the sort of thing that any abstract model theory does. It does not, and is not intended to, describe the world, but to model it, and so reveal its underlying structure. One would have to be very naïve to suppose that Newton, or anyone else, ever believed in the point-masses they postulated in their theories.

Consider also Sadi Carnot's theory of the heat engine (1824). After 180 years, Carnot's theory is still the fundamental one in the area, and every student of thermodynamics has to have a good understanding of it. But Carnot's model of the heat engine is consciously not realistic. Carnot sought to explain the workings of the heat engine by abstracting from the heat losses, and other causes of inefficiency, that occur in all real heat engines, in order that he might consider the fundamental nature of the processes involved in producing useful work. Carnot was a caloricist, and his hypothesis was that heat produces useful work in the process of falling from a higher to a lower temperature level, just as water in a water mill does in falling from a higher reservoir to a lower one. We now know that Carnot was wrong about this, and, according to Psillos, there is evidence that Carnot himself had doubts about the nature of the process (although you would never know this from Carnot's original paper). Whether this was so or not, Carnot was not wrong about the essentials. Essentially, work is produced by maintaining a gas at a high temperature while allowing it to expand isothermally, allowing it to expand further adiabatically, so that its temperature drops, compressing it isothermally at this lower temperature, and then compressing it further adiabatically to restore it to its original pressurised state at the higher temperature. The work done in a given cycle is the excess of the work done by the gas in the expansion phase over the work required in the compression phase.

Carnot's model of the heat engine was an idealisation of this process. The temperature of the source of heat required for the isothermal expansion was assumed to be the same as that of the expanding gas, so that no temperature differences would appear, and hence no heat losses would occur, during this phase. Likewise, the temperature of the heat sink was assumed to be the same as that of the working substance throughout the isothermal compression phase. The adiabatic expansion and compression phases of the cycle were supposed to occur in a perfectly insulated cylinder, so that no heat losses would occur during these phases either. The whole process was supposed to involve a perfectly lubricated cylinder to eliminate the possibility of work losses due to friction. Obviously, no real heat engine even approximates to this ideal, and the efficiencies of real heat engines are orders of magnitude less efficient than Carnot's theory implies they could be. But Carnot's theory is nevertheless the fundamental one in the theory of heat engines. For, using Carnot's model, and what is now known as the Second Law of Thermodynamics, it is not hard to demonstrate that:

a) The efficiency of a Carnot engine is independent of the nature of the working substance, and is a function only of the temperature limits through which it operates. Specifically, $e = (T_2 - T_1)/T_2$, where T_2 is the Absolute temperature of the source, and T_1 that of the sink,
b) The Carnot engine is the most efficient possible for any heat engine operating between these temperature limits.

But neither Carnot, nor anyone else at the time, ever thought that Carnot's engine was anything other than a theoretical fiction. So, Psillos's semantic thesis, which would require us to take such well-established theories such as Carnot's 'at face value', and see them as 'truth-conditioned descriptions of their intended domains', does not seem to be good advice. Good theories in established fields in the mature sciences, even seminal ones, need not be realistic.

The case of forces, conceived of in the manner of Newton as extrinsic to the bodies on which they operate, is different again. According to my proposed definition of physicality, Newtonian forces are not physical objects, since they do not have energy. The action of a Newtonian force is a physical event, since it necessarily involves some change in the object on which it acts. But strangely, the production of such a force is not. The object that produces the force is not affected by the fact of its production; it is affected only by the reaction that its action produces. So, if Newtonian forces are objects, then they are very curious objects. As Robert Mayer once noted: 'If gravity be called a force, a cause is supposed which produces effects without itself diminishing, and incorrect conceptions of the causal connexion of things are thereby fostered' (Magie 1935: 199). Theoretically, Newtonian forces are always eliminable from physics. We need only combine the laws of their production with those of their combination and action to obtain laws of distribution of their effects. But, in practice, we cannot always do this, because the resultant forces are often the products of too many and too complex factors, and the measurements and calculations that would be required to eliminate the forces would be beyond us, even if there were some point in trying to do it. So, should we be realists about forces of the kind that Newton described? The standard arguments for scientific realism clearly say 'yes'. They are, after all, theoretical entities in causal roles in highly successful theories in one of the maturest of all of the sciences. But a physical realist should say 'no'. Forces are not physical entities. Therefore, however useful they may be as theoretical entities, they must be rejected in ontology.

There are, therefore, some very good reasons for thinking that established physical theory is not ontologically transparent. Firstly, the ontological implications of established propositions about Platonic entities, such as numbers, geometrical points, Euclidean planes, propositions and possible worlds, simply cannot be read off from the page, as though they were unproblematic. There are, of course, realists about all of these things. But according to physical realism, if one wishes to be a realist about any of them, then one must construe them as universals of some kind. For these are the only kinds of physical entities in a physical realist's ontology that they could possibly be. But mathematical entities do not appear to be universals either.[1] Secondly, the ontological implications of laws about theoretical ideals are not obvious. Certainly, there are well-known and well-established laws that purport to be about ideal things of one kind or another. For example, there are established laws of ideal gases, inertial systems, black body radiators, ideally free markets, perfect competition, and so on. And since these laws are all presumably true in the minimalist epistemic sense, a scientific realist who accepts Psillos's semantic thesis would appear to be committed to the reality of all of them. But according to physical realism these entities do not exist as physical objects. They are, rather, the objects of ideal theories, which are not intended to describe the world, as Psillos's semantic thesis requires. Thirdly, there are objects of conceptual convenience, such as forces. There is good reason to think that forces are not physical entities. Certainly, they are not physical objects, events or properties, as I have defined them. Therefore, there is good reason not to include them in a scientific ontology. Therefore, scientific realists should not, qua scientific realists, be committed to their existence. Psillos's semantic thesis should therefore be rejected.

4. Explaining science metaphysically

The best argument for physical realism is not an argument from the empirical success of science to the truth of it laws and theories to the reality of its theoretical entities. For no such argument is able to be as metaphysically discriminating as it should be about the different kinds of theories in science, or the different kinds of theoretical entities that are postulated. The best argument is the one that derives from the extraordinary nature of the new scientific image or the world, and the attempt to explain it metaphysically. For the question that needs to be addressed is this: How is the sophisticated, relatively stable, scientific image of the world that is the result of the last two or three centuries of scientific work to be explained? Don't look at it theory by theory, I say, and seek to justify the ontologies of the most successful ones in terms of what these theories are able to predict. Look at the picture as a whole. This is what no hypothesis other than a global metaphysical theory such as physical realism can possibly explain. The image that the most successful scientific theories have systemati-

[1] My positive account includes a range of numerical and spatio-temporal relationships as genuine universals, but does not embrace numbers or geometrical points. See Ellis (1987).

cally constructed for us is an extraordinary one. It is an image of a world that consists entirely of things belonging to an elaborate, strongly interconnected, hierarchical structure of categorically distinct kinds (of chemical substances, particles, fields, etc.), and involved in natural processes which themselves are organised in a natural hierarchy of categorically distinct kinds. Moreover, the ways in which they are involved are seen as depending on their specific causal powers, capacities and propensities, and the spatiotemporal relations between them. It is a picture, not of a mechanistic world, but of a world that is every bit as tightly organised and structured as Paley's watch.

According to the Maxwell-Bridgman theory of physical reality, a real thing always manifests itself in more than one way. A basketball not only looks round; it feels round, and it rolls. A basketball that manifested itself only visually, or only from one point of view, might well be dismissed as a chimera. But one that you can catch and throw into a basket is not a chimera – it is the real thing. Much the same is true of the theoretical entities of science that most of us believe in. The same entities crop up again and again in the explanations we offer, and are seen to be involved in a great many different natural processes. A theoretical entity that had only one role in explanation, e.g. to conserve energy in some specific kind of process, would be highly suspect. But one that has as many different properties as a copper atom does, and manifests itself in as many different ways as copper atoms do, cannot plausibly be supposed to be a theoretical fiction. The theoretical picture of modern chemistry, to which the theory of copper belongs, is too tightly interconnected for it to be anything other than what it purports to be.

The emergence of this scientific image of the world really has only one plausible explanation, viz. that the world is, in reality, structured more or less as it appears to be, and, consequently, that the kinds distinguished in it (the chemical substances, particles, fields, biological species, etc.) are natural kinds of one sort or another, and that the causal powers they appear to have are genuine. They may not all be natural kinds in the strict sense in which I use this term in *Scientific Essentialism*. But nor are they just arbitrary classifications that lack adequate foundations in the real world.[2] Any other hypothesis would make the appearance of all

[2] Biological species, for example, are what I call 'generic cluster kinds'. See my reply to Stephen Mumford in Ellis (2005).

this structure in the scientific image astounding. The hypothesis that the appearance of structure arises from our manner of perceiving or thinking about the world has no plausibility at all. It does not even begin to explain the structures that we actually find in chemistry, for example. Moreover, the hypothesis that our mental processing systematically distorts reality in some way, so that the real structures are not the same as the apparent ones, is simply gratuitous. It explains nothing, and the doubt that the structures are as they are represented as being is merely sceptical.

This is the real argument for the proposed ontology of scientific realism. It is a powerful argument, and it is independent of any theory of reference or truth. It does not proceed from any premisses about the truth, or approximate truth, of the laws or theories of science. Nor does it depend on any semantic theory about what makes a law or theory true, or approximately true. It does not even depend on our acceptance of the truth or approximate truth of most established scientific theories. It is enough if we accept that the scientific image is the most rational picture to have of the nature of reality. The argument by-passes all questions about the language of science, and gets down to the crucial question, which is: How is the emergence of a scientific image of the world, consisting of a multiply connected, hierarchical structure of categorically distinct kinds of physical systems that are involved in a range of categorically distinct kinds of processes, to be explained? What gives rise to this image? The image is clearly a human construct. But it is a stable and revealing image that accommodates and explains just about everything in the relevant fields that scientific investigation has demonstrated, and excludes nothing that seems to be indispensable. This is what we should expect, if the new scientific image were, for the most part, descriptive of reality, as physical realism assumes it to be. Its emergence otherwise has no plausible explanation.

La Trobe University
Australia 3083
and
University of Melbourne
Australia 3010
phibde@bigpond.com

References

Bigelow, J. C. and Pargetter, R. J. (1990). *Science and Necessity*. Cambridge: Cambridge University Press.
Carnot, S. (1824). 'Reflexions sur la Puissance Motrice du Feu', reprinted in Magie (1935).
Ellis, B. D. (1976). 'The Existence of Forces', *Studies in the History and Philosophy of Science* 7: 171–85.
—— (1987). 'The Ontology of Scientific Realism', in *Metaphysics and Morality: Essays in Honour of J.J.C. Smart*, ed. P. Pettit, R. Sylvan and J. Norman. Oxford: Blackwell: 50–70.
—— (1990). *Truth and Objectivity*. Oxford: Basil Blackwell.
—— (1992). 'Scientific Platonism. An Essay Review of Bigelow and Pargetter's *Science and Necessity', Studies in the History and Philosophy of Science* 23: 665–79.
—— (2001). *Scientific Essentialism*. Cambridge: Cambridge University Press.
—— (2002). *The Philosophy of Nature: A Guide to the New Essentialism*. Chesham: Acumen.
—— (2005). 'Universals, the Essential Problem and Categorical Properties', this volume.
Magie, W. F. ed. (1935). *Source Book in Physics*. New York and London: McGraw-Hill.
Psillos, S. (1999). *Scientific Realism: How Science Tracks the Truth*. London and New York: Routledge.
Sellars, W. (1963). *Science, Perception and Reality*. London: Routledge and Kegan Paul.
Smart, J. J. C. (1963). *Philosophy and Scientific Realism*. London: Routledge and Kegan Paul.