affair. It does, however, contain several papers which are valuable responses to Feyerabend.

JOHN M. PRESTON
University of Reading

Realism Rescued: How Scientific Progress is Possible
JEROLD L. ARONSON, ROM HARRÉ & EILEEN CORNELL WAY, 1994
London, Duckworth
vii + 213pp., Hb 0715624768, £30.00

Realism Rescued is a brave, innovative and challenging attempt to defend afresh a substantive form of scientific realism which commits itself to the existence of natural kinds in the world and secures theoretical truth and progress in science. The novelty of Aronson, Harré and Way's approach (henceforth AHW) is their use of the technical device of type-hierarchies to address and reshape a whole bunch of issues that pertain to the debates over realism, but also to more general themes in the philosophy of science (e.g. the nature of models, the analysis of counterfactuals and the status of laws of nature and of physical properties). Since in this short notice I can only focus on some selective aspects of this rich book, I will concentrate on AHW's use of type-hierarchies to represent scientific theories and to explicate the concept of verisimilitude.

1. Type-hierarchies and natural kinds

AHW construe scientific theories as type-hierarchies which intend to capture the structural relationships between natural kinds. Generally, a type-hierarchy (more broadly, a semantic network) is a tree-structured graph of nodes joined by links, where nodes represent objects or concepts and links represent the relations between them. One of the most useful ways to characterise type-hierarchies is as having higher nodes standing for (super)types (e.g. ANIMAL), intermediate nodes standing for subtypes (e.g. DOG) and leaf-nodes standing for tokens/individuals (e.g. Fido). Then, links are taken to exemplify "a-kind-of" or "instance-of" relations between nodes in the hierarchy. This kind of representation is useful because it moves from the less to the more specific (e.g. from type ANIMAL to subtype DOG) and this move is accompanied with *property-inheritance* (e.g. if animals are living organisms, so are dogs). Practically this means that type-hierarchies can facilitate knowledge representation and, in particular, inference (e.g. one can infer that since Fido is an instance of the type DOG it is also an animal etc.).[1] Once this feature of type-hierarchies is taken into account, one can easily see how this device can be employed to characterise the structure of a network of natural kinds. But given that natural kinds are characterised in terms of laws that govern their behaviour, a type-hierarchical account of scientific theories should also include these laws. AHW suggest that each node of the type-hierarchy incorporates a corresponding set of laws governing the behaviour of the system represented by that particular type (p. 40).

A defence of realism, however, needs to go far beyond the defence of this technical device as an adequate representation of scientific theories. For the issue at stake is this: is the world objectively carved up in natural kinds with determined boundaries and relations between them or is the taxonomic representation of "natural" kinds a matter of convention, convenience, context and the like? Since we can carve up nature in a good many different ways—responding to different needs and purposes—isn't the choice

between different type-hierarchies context-dependent and, ultimately, conventional? Why should we think that a given type-hierarchy cuts the world at its joints, or that it's a better representation of the natural ordering of kinds? These are not merely rhetorical questions. For instance, one can read Kuhn's incommensurability thesis as making precisely the point that (a) each theory creates its own taxonomy of natural kinds based on its own concepts and connections between them and (b) the taxonomies (or type-hierarchies) of two competing theories cannot be mapped upon one another.[2]

Realism *vis-à-vis* natural kinds should at least involve the metaphysical claim that the causal structure of the world is given and that different type-hierarchies are to be judged with respect to how well they represent this structure. This kind of position can leave some room for some conventional classification and choice between type-hierarchies, but only on the assumption that this conventional element is causally inert: it contributes nothing to the causal powers and the causal connections between the objects in the world. For instance, a realist can admit that a type-hierarchical classification of harmonic oscillators in terms of colour or manufacturer can be totally conventional, but not so for a classification in terms of the forces to which the oscillators are subjected. For it is precisely those forces (and the concomitant laws) that specify the causally relevant different types of harmonic motion. AHW capture this point by arguing that "choice between competing schemes may be a matter of convention to the extent that each system is compatible with the actual causal structure of the world" (p. 43). This attitude accommodates conventional judgement within a realist understanding of natural kinds. But it generates some natural questions. For instance, how can we determine what the actual causal structure of the world is like? How can we determine which classifications reflect the causal structure of the world and which are causally inert? If we are not able to this, then we cannot possibly judge whether a given type-hierarchy is natural or conventional.

Realism Rescued contains only the seeds of an answer to these questions. According to AHW the stable, invariant, elements of a type-hierarchy, the types and laws that remain invariant when embedded in different type-hierarchies, are precisely those that represent the objective causal structure of the world (p. 43). This is a rather promising line, the essence of which, i.e. the idea of invariance, has been recently popular among realist philosophers of science.[3] It boils down to the claim that the objective causal structure of the world—the joints of nature—is reflected in the stable elements of the sequence of type-hierarchies by means of which we attempt to comprehend the world. And this sounds akin to a kind of Piercean view that the causal structure of the world is what will be recorded in the ideal or in-the-limit-of-science type-hierarchy. But AHW leave this line of thought undeveloped. It is arguable that this position is the best defensible realist account of natural kinds and that, brought in contact with the broader view that science aims at unification, it can be used to support the objectivity of science and the idea of scientific progress. These central issues are only sketched and not adequately explored in *Realism Rescued.*

2. Type-hierarchies and verisimilitude

AHW want to defend a full-blown version of scientific realism which incorporates and defends some notion of verisimilitude. They rightly think that scientific realism without verisimilitude is like salt that has lost its savour. (To be sure, their position is that scientific realism is primarily a metaphysical doctrine "but truth, realism and verisimilitude are all parts of a single picture", (p. 123). Armed with the representation of

theories as discourses about type-hierarchies, they attempt to give an account of verisimilitude in terms of similarity relations within type-hierarchies. AHW's novelty consists in reversing the problem-situation concerning the characterisation of verisimilitude. Instead of defining verisimilitude in terms of the distance from the whole truth, they define truth in terms of verisimilitude. Truth, for AHW is "a limiting case of verisimilitude" (p. 123). The rough idea is that verisimilitude pertains to cases where the type picked out is similar to the actual type, whereas truth pertains to cases of real match between two types. More precisely, two types (of objects) are said to be similar when they are represented as subtypes of the same supertype (e.g. WHALE and DOG are both subtypes of the supertype MAMMAL). But similarity is a notion that admits of degrees and AHW capture this feature rather nicely by suggesting that the degree of similarity between two types is a function of their relative locations within the type-hierarchy (p. 129). AHW define verisimilitude by means of a distance function (borrowed from Amos Tversky) such that the distance (degree of similarity) between two types is the weighted difference between (a) the properties that the two types share in common and (b) the properties with respect to which the two types are different (pp. 132–133). Then the degree of verisimilitude of a truth-claim is calculated by finding the distance of the type picked by this object from the type of the actual object. For instance, assuming that the actual object is a humpback whale, the verisimilitude of the claim that the object is a dolphin depends on the distance (defined as above) between the type DOLPHIN and the type HUMPBACK WHALE in a type-hierarchy. On this account, a claim C (e.g. that a given object is a dolphin) can be more or less verisimilar than another one C^* (e.g. that it is a haddock), depending on whether C's distance from the actual object (e.g. a humpback whale) is smaller or greater than C^*'s distance from it. And similarly, truth is presented as a limiting case of verisimilitude to the extent that the type picked out is identical to the type corresponding to the actual object.

It is impossible to do full justice to such a complex and innovative attempt to characterise verisimilitude within the confines of this short piece. So I shall restrict myself to some general critical comments. For I think this attempt to characterise verisimilitude in terms of locations in type-hierarchies creates no less problems than it solves. For one, there seems to be a general problem with the use of similarity to account for verisimilitude. As we saw, AHW construe similarity in terms of locations in a type-hierarchy. But what determines a type-hierarchy in the first place, if not some similarity relation between the types chosen to stratify the hierarchy? If this is so, then it seems rather trivial—and not explanatory—that type-hierarchies determine similarity relations; they are meant to do so. AHW insist that they do not take the notion of similarity as primitive but that they try to analyse it in terms of locations in type-hierarchies (pp. 127ff.). I agree that given that type-hierarchies are somehow given, the AHW approach to similarity is cogent and really casts new light to the nature and significance of similarity judgements. But, I think, it is not enough to declare the ordering of natural kinds as a "metaphysical primitive" (p. 123). For this does not advance our understanding of how hierarchies of natural kinds come to being.

For another, AHW's theory of verisimilitude is contextualist. The creation of type-hierarchies is, to some extent at least, subject to a judgement as to what factors (or classification) count(s) as relevant and this judgement may vary from context to context. (Here one may only need to note the so-called *qua* problem: two animals may be similar *qua* aquatic beings, but dissimilar *qua* the way they feed their young.) But if similarity relations are based on locations in type-hierarchies, then similarity relations will vary from one type-hierarchy to another. AHW seem happy with this since they say "what is

similar to what and in what degree is a meaningful question only relative to some type hierarchy or other" (p. 129). This sounds fine with regard to similarity. But is someone right to conclude that the question of whether a claim is *verisimilar* is also meaningful only relative to some type-hierarchy or other? If we follow AHW's advice, that is, if verisimilitude depends on locations in type-hierarchies, then the verisimilitude of a given claim will vary from one type-hierarchy to another. So, for instance, the claim that a given object is a dolphin may end up having different verisimilitude depending on whether the embedding type-hierarchy is that of mammals or that of aquatic animals. And similarly, the claim that an object is a dog, given that the actual object is a dolphin, may be more verisimilar in the type-hierarchy of mammals, less verisimilar in the type-hierarchy of animals with tails, but simply false in the type-hierarchy of aquatic animals. My view is that although, strictly speaking, AHW's characterisation avoids the standard objection of language-dependence, it generates a new problem, namely that of verisimilitude being context-dependent. On AHW's theory, claims of verisimilitude can only arise in particular contexts, expressed in terms of certain type-hierarchies, and they can be questioned in other contexts, where different type-hierarchies are chosen. But contextualism cannot capture a realist understanding of verisimilitude. From a realist point of view, it makes perfect sense to say that a claim C is more verisimilar than C^*. Yet it does not make good sense to say that a claim C has such-and-such verisimilitude according to *this* type-hierarchy but different verisimilitude according to *that* type-hierarchy. I therefore think this type-hierarchical approach to verisimilitude needs to be freed from the context-dependence of our ways to classify in hierarchies and be such that it tracks the causal structure of the world. If, as AHW suggest, "there is no way of answering the question of what the real ordering of biological types is", and if "this [ordering] will depend on the purpose for which this type or that type hierarchy is to be put" (p. 134), I cannot see how the type-hierarchical approach to scientific theories can offer a realist account of verisimilitude and (if truth is a limiting case of verisimilitude) of truth.

AHW's defence of realism culminates with offering an argument that "increasing plausibility is an inductive ground for increasing verisimilitude" (p. 191), where plausibility is construed in terms of (a) empirical adequacy and (b) metaphysical adequacy: the model that is used to represent unobserved processes must be embedded in a hierarchy of established natural kinds. This argument is based on causal manipulability and can be sketched as follows. Models guide the manipulation of objects and structures to which we do not have direct access. But in most cases of accidentally and contingently unexamined structures (i.e. structures which either no one had, as a matter of fact, looked into, or no one at the moment knows how to carry out an examination) this model-guided manipulation led to theoretical success in that, when examined, the structure was found to resemble the model. Hence, the history of accidentally and contingently unexaminable cases inductively links successful model-guided manipulation with verisimilitude (cf. p. 200).

AHW are aware of the point that this argument needs further premisses if it is to be successfully extended to what they call "necessarily unexaminable" structures, i.e. structures that involve genuine unobservable entities. For what justifies the move from the case of accidentally and contingently *unexamined* structures to the disputed case of necessarily *unexaminable* ones? Isn't this, after all, the point of the dispute between sceptics *à la* van Fraassen and realists? AHW say that "we can leave out the 'observation' link [between successfully manipulable models and unexamined-but-examinable structures] from this point on. We have established inductively that there is a direct link

between manipulability and verisimilitude" (p. 200). But all this hangs on offering further arguments against the alleged relevance of observability to one's epistemic attitude towards theoretical claims. AHW rightly put a lot of weight on the "principle of epistemic invariance" (p. 194), which, in effect, states that if manipulability offers evidence at all, this evidence is equally effective, no matter whether we are dealing with observables, possible observables or unobservables. I think this is sound and captures what the realists should argue for. But AHW need to say more in defence of the principle of epistemic invariance than they do in *Realism Rescued.*

In short, realism is not quite yet rescued from all the challenges and attacks that have been thrown at it. Realist philosophers of science need to do more work. But *Realism Rescued* suggests some new directions that the defence of realism may take, explores new conceptual territories, and offers a brave, if not a knockdown, argument for realism. No one seriously interested in the realism debate should miss this book.

Notes

1. For a detail account of the use of semantic networks in computer science cf. T. MacRandal (1988) Semantic networks, in G. Ringland & D. Duce (Eds) *Approaches to Knowledge Representation* (Research Studies Press).
2. For more on this reading of T. Kuhn's thesis cf. my review (1994) of P. Hoyningen-Huene's *Reconstructing Scientific Revolutions, British Journal for the Philosophy of Science,* 45, pp. 923–926.
3. cf. P. Kitcher *The Advancement of Science,* Oxford University Press, 1993, pp. 171–172.

STATHIS PSILLOS
Kings College London

Scientific Nihilism: On the Loss and Recovery of Physical Explanation
DANIEL ATHEARN, 1994
State University of New York Press, Albany
ix + 387 pp., Hb ISBN 0–7914–1807–3, £52

This is a book of two parts, the first critical, the second constructive. Athearn's basic critical thesis is that modern physics is non-explanatory and that philosophers of scientific explanation have mistakenly taken mathematical physics to be a prime example of successful science. The constructive project is a bold attempt to give an alternative, and, in Athearn's terms, truly explanatory account of radiation.

What, then, is it for something to be truly explanatory? For Athearn, a genuine explanation gives a causal narrative account "identifying and comprehending the physical processes ... that generate observed phenomena".

Why, according to Athearn, have physicists neglected the causal narrative, or explanatory aspect of science? One defence might be that modern physics, and quantum mechanics in particular, has shown that the world is simply not governed by causal processes. According to Athearn, this is a fundamental mistake. The advent of modern physics showed that traditional concepts of mechanism were to be discarded, but not that causality itself should be abandoned—the mistake of the physicists was to conflate these two notions. Mechanical causes are those involving vibrating matter or moving objects—one billiard ball colliding with another, and causing the latter to move. Athearn's notion of causality is broader than this, and becomes important in the constructive part of his book, in which he gives, amongst other things, a non-mechanical account of non-local causation.

This explanatory project may sound familiar to philosophers of physics, who