SCIENCE VS REALITY: A DEBATE

Guest Editors: David Braybrooke, Thomas Vinci

This special issue of *The Dalhousie Review* contains the essential proceedings of the conference on scientific realism held at Dalhousie University in the summer of 1983. Also included are four independent essays on related philosophical topics; these appear after the conference proceedings and before the book reviews. The conference proceedings includes complete papers, summaries or reports of papers, records of discussion that followed papers, and linking commentary by the editors. We have tried to signal the components by changes of typeface.

Not all of the vital areas of philosophical research in the Dalhousie department are here represented. The next issue of *The Review* contains a paper by another member of the department. However, this issue can be thought of as an indicator of the philosophical work at Dalhousie over the past quarter of a century.

IS THE WORLD REALLY WHAT SCIENCE TELLS US IT IS? A SUMMER DEBATE AT DAHOUSIE

Introduction

The Dalhousie department of philosophy has been the scene in recent years of a number of summer conferences. Two, in successive summers, treated vagueness: How is logic, which has assumed exact propositions, to deal with empirical propositions, considering that all of them are vague to a degree? A pilot conference one summer took up the present status of the concept of causality. It was followed another summer by a full-dress conference on causality, bringing in leading contributors to the subject. In the summer of 1984, a three-week institute brought economists and political scientists together with philosophers to discuss the theory of rational choice extended from private choices to public ones. The Social Sciences and Humanities Research Council of Canada, which had joined the university in sponsoring the conference on causality, joined the National Science Foundation of the United States in funding the institute on public choice.

In 1983, with funds from the Council and the university, the department held a conference on a theme broader than any of these, indeed broad enough and important enough to count as a version of the chief question of philosophy:

Is the world really what science tells us it is?

Some philosophers — the scientific realists — confidently say that it is; and argue that if any conflict between scientific findings and common sense ideas crops up, common sense must give way. Other philosophers — the anti-realists — disagree. They point out that even the most successful of past scientific theories — Newton's for example — have been discredited in time. By induction, they say, we may expect our present theories to turn out false, too.

At this point the scientific realists may fall back on the position that some day, at the ideal limit of scientific investigation, science will give a perfectly true account of reality. Meanwhile, they say, it will behoove us to keep abreast of science, changing our views according to its teachings, without trying to cling to common sense.

Some anti-realists in the philosophy of science may say in rejoinder that by and large common sense has already survived a number of revolutions in scientific theory; and can be expected to survive more. It gives us, by and large, true descriptions of phenomena — of the way the world looks to us; and will continue to do so. The task of science, anti-realists say, is to find theories that will "save" these phenomena, that is to say, give us means of deducing predictions of them. To carry out this task the theories need not claim to be true, or even claim that the entities which they postulate exist; and they had better not, because we must be prepared to discard the theories and the entities when science advances.

But what, the realists say, can an advance in science signify if it does not amount at least to coming closer to the truth? Furthermore, realists will insist, common sense does not give us the phenomena without begging questions about their character. Common sense embodies a lot of theories. They are folk theories, and science discredits them.

Such were the positions represented and debated in the conference on scientific realism. The peak week of the conference fell during the first week of August, with most of the participants from outside Dalhousie coming in then. But, as a run-up to the debate of the peak week, two of our visitors, James Brown of the University of Toronto and Kathleen Okruhlik of the University of Western Ontario, canvassed some of the chief texts of the current philosophical literature in a summer session class. Most of the members of the Dalhousie department, along with our graduate students, regularly attended the class. In attendance also were other philosophers teaching in the summer session, Calvin Normore, then of Columbia University, and George Pappas of Ohio State, along with a graduate student from Syracuse University, Carl Matheson, who did his B.A. and his M.A. in philosophy at Dalhousie. All three were, like Brown and Okruhlik, scheduled speakers at the conference.

Fog stranded the first speaker of the peak week in Fredericton. Undaunted, and eager to have the first crack at the subject, Paul

Churchland from the University of Manitoba, the author of a sprightly recent book, Scientific Realism and the Plasticity of Mind *(Cambridge University Press, 1979), rented a car and drove through the night to keep the appointment in Halifax. Here, in synopsis, is what he said:*

Paul M. Churchland
*The University of Manitoba**

On The Global Assessment of Human Cognition

Having slowly emerged as the dominant epistemological theory, scientific realism is now quite properly under intense scrutiny and criticism. Versions of the view proliferate, and criticisms flow from a variety of quarters. We may enter the debate by considering a view that is now over twenty years old, a view I still think of as *the normative core of scientific realism*:

> Global excellence of theory is the ultimate measure of rational belief and rational ontology, in all spheres and at all levels of human cognition, including the perceptual level.

"Belief" here means belief in the literal *truth* of propositions, but this core thesis is nonetheless consistent with the idea that belief should be tentative at best, and with the idea that all or most of our current theories are strictly speaking false, or even entirely devoid of reference. While circumspect in this respect, it is extremely bold in its unrestricted assumption concerning the speculative and theoretical character of all human knowledge. Conjointly, these features have made this view an extremely fertile and attractive view.

First, from this perspective our common-sense conceptual framework for apprehending reality reveals itself as an intricate network of 'folk *theories*': folk mechanics, folk topology, folk psychology, folk thermodynamics, and so on. Second, the structure of common-sense explanations becomes clear, when they are seen as exploiting the humble 'laws' of folk theories. Third, the curious features of certain common-sense concepts—e.g., psychological or intentional concepts—become clear when seen as structural elements in a sophisticated folk psychological theory. Fourth, it smoothly solves certain long-standing

*Now (1985) at the University of California, San Diego.

epistemological and ontological problems: the other-minds problem, for example, and the problem of Eddington's two tables. A belief in other minds is an explanatory hypothesis, beliefworthy to the extent that it is explanatorily and predictively successful. And it is the table-as-described-by the *best theory* (the scientific theory) that is the real table. The common-sense table is real only if the common-sense theory descriptive of it reduces cleanly to the more powerful theory.

Core realism is also very liberating. It frees one from the tyranny of 'a priori' truth: all sentences are elements in a theoretical network, and no sentence is immune from revision as its embedding network is subject to global evaluation. More liberating still—and this has yet to be *fully* appreciated by the philosophical community—core scientific realism frees us from the tyranny of 'The Given' in perception.

On this view, our perceptual judgments themselves emerge as theoretical responses (= judgments expressed within the idiom of some speculative conceptual framework) to the causal impingements of the environment, and thus must also bend to critical assessment as the theories they presuppose are subjected to global evaluation. This frees us not only from certain epistemological theories, such as positivism and instrumentalism: *it frees us to transform the conceptual content of sensory perception itself*, in a potentially endless variety of useful and revealing ways, limited only by the bounds of human theoretical imagination, and by the actual intricacies of our causal interactions with the world. The trick is to be pulled by learning to respond habitually, to one's pre-conceptual sensory processing, with judgments framed in the idiom of successful scientific theories. This holds promise for expanding our sensory perspective not only of the external world, but of the internal world as well. Introspection can be informed by neuroscience, just as vision can be informed by physics.

All of this presents a very enthusiastic picture of core scientific realism. Certainly its positive elements have been sufficient to make me an ardent and vocal exponent of that general view (see my [1]). Even so, this picture also contains the makings of a rather poignant awkwardness about scientific realism. Very briefly, the awkwardness is as follows.

If all of our knowledge and understanding is speculative and theoretical, then so is the conceptual framework we have been calling 'folk psychology'. Accordingly, the familiar run of cognitive concepts—"believes that P", "perceives that P", "infers that P", and so forth—show up as theoretical concepts, concepts whose claim to descriptive integrity is only as good as the global theoretical virtues of folk psychology. Intriguingly, or perhaps distressingly, folk psychology is highly suspect, both as measured by its simple explanatory and predictive success, and as measured by its consilience or coherence with neighbouring domains of science, such as structural and developmental neuroscience and evolutionary biology. Folk psychology, in short, threatens to be empirically false, and its ontology of cognitive states is thus suspect. *But the standard statement of core realism presupposes the existence of those states, and the rough integrity of the internal economy in which they are supposed to figure.* If we cannot speak of beliefs, or of rationality, how can we even formulate, let alone urge, anything like the sort of realism defined earlier?

Few philosophers have yet been moved to find a major problem here. The battle is much noisier over the cut between theory and observation, the approximate truth of mature theories, the limits of rational ontology, and whether science should aim as high as seeking truth. As I see things, however, *here* is where all the bodies are buried. And such other problems as do indeed confront scientific realism—such as the nature of rationality, the relation between our cognitive representations and the world, and the long-term prospects for human cognition—all of these problems will eventually lead us back to the matter of the ontology of folk psychology (and folk semantics, and folk epistemology), and of its ultimate dynamical relevance and categorial integrity.

If I am right in thinking this, then scientific realism might well require a fundamental reformulation, one that reconceives the units of cognition as something other than sentences or propositions, and reconceives the goal of cognition as something other than truth (some goal even higher than truth, perhaps). Such a revolutionary theory might exploit the cognitive categories to be supplied by a matured neuroscience, by a metric tensor theory of cognitive transformations, or by a matured theory of self-organizing systems, or by some information-theoretic version of non-equilibrium thermodynamics, or by something along these lines. Before returning to these themes, let me examine briefly some of the important criticisms of scientific realism recently levelled by Larry Laudan, Bas van Fraassen, and Hilary Putnam.

Laudan's most excellent critique [4] is aimed at certain naive and over-inflated versions of realism. His discussion shows that neither the verdict of history, nor the attempt to explain our scientific success, will support common realist claims about the approximate truth and referential integrity of theories in the 'mature' sciences, or even the claims that science will ultimately converge on truth. I think Laudan is right. But notice that these points are consistent with the highly circumspect core realism defined at the outset of this paper. All cognition remains speculative and theory-laden; and the claim to truth of a successful theory, while very weak in absolute terms, may still be strong relative to the available competitors.

Van Fraassen's critique [3] cuts more deeply into the realist position, or tries to. Consider the distinction between

(1) things observed by some human (with unaided senses),
(2) things thus observable by humans, but not in fact observed,
(3) things not observable by humans at all.

Van Fraassen's position excludes (3) from our rational ontology. He is willing (anxious) to allow (2) into our rational ontology, and simply shrugs his shoulders at Hume's traditional underdetermination problem, complaining that (2) may be problematic, but that (3) is an additional and even deeper problem. He has to do something like this, since it would not be at all plausible to exclude both (3) *and* (2) from our rational ontology.

Van Fraassen thus requires a *principled* distinction between (2) and (3), a distinction *adequate* to the radical difference in epistemic attitude he would have us adopt towards them. However, when one examines the actual grounds of the distinction between (2) and (3), one discovers that it is only very feebly principled, and is wholly inadequate to bear the enormous weight that van Fraassen puts on it. Things in (2) are simply those things that fail to enjoy an appropriate spatio-temporal *position* relative to our native senses. Things in (3) are simply those that fail to enjoy an appropriate spatio-temporal *size* relative to our native senses, or an appropriate *energy*, or an appropriate *wavelength* relative to those senses. There is a minor practical point to the (2)/(3) distinction in ordinary language, since in the absence of technology, we generally have more voluntary control over the spatio-temporal position of our senses, than over their size or constitution. But that is an accidental and idiosyncratic fact about humans, which disappears with technological aids like microscopes and slow-motion photography. And in any case, the logical problem is the same whether we are inferring across a 'gap' in spatio-temporal distance, spatio-temporal size, energy level, or wave length: the problem is underdetermined hypotheses and ampliative inference. Quite aside from its motivation, therefore, van Fraassen's scepticism is arbitrary and unprincipled in the bounds that it sets. Core realism, therefore, need make no concessions here.

Putnam's 'internal realism' [5] presents a different and more oblique challenge to scientific realism. From an 'internal' point of view, Putnam can accept our 'core realism' with equanimity. Putnam's (sceptical) concern is with the concept of truth itself, and with the integrity of that notion, as traditionally conceived. Here I am inclined to sympathy, since I wish myself to reexamine the integrity of our 'folk semantic' and 'folk epistemological' notions.

What this means is that our 'core scientific realism' is not seriously threatened by any of the major criticisms of realism currently urged. The real threat lies in the possible poverty of the common-sense concepts it presupposes.

BIBLIOGRAPHY

[1] Churchland, Paul M., *Scientific Realism and the Plasticity of Mind* (Cambridge, 1979).

[2] --------------, "The Anti-Realist Epistemology of van Fraassen's *The Scientific Image*", *Pacific Philosophical Quarterly*, 63, 3 (July, 1982).

[3] van Fraassen, Bas, *The Scientific Image* (Oxford, 1980).

[4] Laudan, Larry, "A Confutation of Convergent Realism", *Philosophy of Science*, 48, 1 (March, 1981).

[5] Putnam, Hilary, *Reason, Truth, and History* (Cambridge, 1981).

In the discussion of Churchland's paper, Duncan MacIntosh (a graduate student from the University of Toronto) led off by objecting that it was both implausible and incoherent to maintain that the folk concept of truth was going to be superseded. How could you express as a candidate for belief the doctrine that there is no place for the concept of truth? Churchland replied, "What I'm imagining is that the current family of cognitive notions would be replaced by another family of cognitive notions which stood to the old ones the way a 747 stands to a flying carpet." As for incoherency, the incoherency is no worse than that involved in arguing that there is no such thing as vital fluid with someone who believes that nothing can be a living being who does not have it. If that person believes that only living things can put forward assertions, he may claim that to assert there is not vital fluid is incoherent, since the very act of asserting this is proof that there is. But this is absurd.

Simon Blackburn, from Oxford, objected that the core realism which Churchland was defending does not sit very happily with the continuity between science and common sense. Many common beliefs are quite stable and their stability can be explained by holding that they are true. The historical induction on which Churchland relied is spurious. Not all beliefs are in danger of being superseded. Churchland maintained, against this objection, that all common sense beliefs were suspect. There are, for instance, no properties in reality answering to the hot and cold of common sense because that view sees one dimension of properties where reality gives you three: degree of heat, amount of heat energy, and rate of flow of heat energy. It may be admitted that some common sense notions will resist supercession more than others. Functional notions, for example, like "chair" and "table" can survive many changes in views as to the physical constitution of things. But are these theoretically interesting cases? Are statements asserting the existence of single items of this kind, or of the city of Rome, theoretically interesting?

William Seager, from the University of Toronto, pressed the point that meaning is relative to human beings and hence to the senses which we happen to have. What is observable must be distinguished and settled with respect to them. Churchland held that this distinction is not itself hard and fast, but shifts with the introduction of new instruments of observation. It's not true, as van Fraassen assumes, that if you see something, then there's no problem about it. All the cognitive work remains for the mind to do, and it's as imperilled there, even in its ontology, as it is anywhere else.

"I'm an old realist," Churchland went on, "seeing the need for and wanting to become some new kind of realist but still wanting the resources that would enable me to do this . . . . I'm in between a rock and a hard place here. I don't think that signals anything essentially wrong or poverty-stricken about my position. It just signals a painful point in the transition of what I hope is a rational evolution of belief."

Churchland's paper tried to extricate scientific realism from current controversies about it and to begin formulating it in new, less vulnerable terms. The next speaker, Wilfrid Sellars of the University of Pittsburgh, one of the chief spokesmen for scientific realism in the past quarter of a century, went back behind current controversies to discuss the roots of scientific realism in epistemology and the theory of meaning. The following is a précis, by the editors, of Sellars's talk:

Wilfrid Sellars
University of Pittsburgh

Meaning, Truth, and Realism

You can't see what doesn't exist. So if I accept it that I am seeing the color of a red fire truck, I must inevitably accept it that the color of the truck exists. On the other hand, when I see that the fire truck is red, I am accepting it (perhaps without seeing the color itself) that "red" is properly predicated of the fire truck. Is what is predicated—"redness"—

something different from the color that the fire truck has? I would say "Yes", the empiricist says, "No." He would not (as have some philosophers following Plato) identify redness with an abstract object itself invisible, though he would have to admit that there are some abstract properties that cannot be simply seen, for instance, magnetism. We thus get a bifurcation between predicates; some will signify objects that can be seen, some will signify objects that cannot.

For empiricists, objects of the second kind can't claim first-class existence; hence propositions involving them can't claim first-class truth. At best, they possess a kind of warranted assertibility that falls short of truth itself. This is the empiricist position, expressed in the view that instrumentalists and other opponents of scientific realism take of theoretical predicates and propositions.

All of this presupposes that we have a workable theory of perception that divides perceptible features from imperceptible ones. Suppose we do. Why should imperceptible features or objects—objects imperceptible and theoretical—be deemed second-class? One reason is that perceptible features, besides being actually sometimes present to observation, are features of which we can form images. But are images necessary or even suitable as vehicles for abstruse thinking? Scientific realists must hold that we can have thoughts in a full-bodied sense of objects both imperceptible and unimaginable. A second reason holds that the domain of inference is separate from and dependent on the domain of observation. But this is not so. All observations involve inferences.

Phenomenalism, which pretended to define theoretical predicates in terms of perceptible features, has vanished as such, but instrumentalism, which survives, is a variation upon it. Ernest Nagel, for example, grants that theoretical terms are not definable by observational ones, but holds that the spheres of application of the first can be bounded by citing observational terms.

Suppose we succeed in eliminating the ontological significances of the distinction between theoretical predicates and observational predicates. Might we not have to worry that now all predicates fail to designate genuine, first class property-existents? The first worry gives rise to a second: If truth as a property of a sentence 'Fa' implies exemplification as a relation between the designatum of 'a' and a *first-class* property-designatum of "F", are we not committed to denying that *any* form of predication would be true? This poses one of the central challenges to scientific realism.

Truth is an issue somewhere in the same old thicket of issues about what is really in the world and on what conditions we can relate language to reality so composed. Such a leading anti-realist as van

Fraassen gives astonishingly little attention to truth. If he had looked into it, he might not have found it so easy to distinguish his own position from realism. Let's take a look ourselves, asking as our central question, "How does truth connect with meaning?"

A correct theory of meaning does away with abstract objects. The Platonic notion of an abstract object corresponding to properties (sensible or otherwise) is a mistake about the use of distributive singular terms. "The lion" in "The lion is tawny" does not refer to lionhood, but to any of the set of standard lions; of any of these it is said that it is tawny. But we can equally say, "Tawny is true of the lion," or of any particular lion Leo. And here "tawny" functions as a distributive singular term, referring not to an abstract object tawniness, but to any of the tokens in English of the word-type "tawny" in their standard use, and hence related in a standard way, psychologically, sociologically, and historically, to objects in the world. Similarly, "red" in "Red is true of fire-trucks" (or of this fire-truck) refers to any of the tokens in English of the word-type "red," but not only to these. It refers also to the tokens, in corresponding standard uses, of the words for "red" in other languages.

However, as the medieval anti-Platonist "breath of air" (*flatus voci*) theory of universals held, predicates are not names. The use just made of "refers" is strictly unacceptable. Predicates, that is to say, predicated-symbols, are auxiliary symbols the function of which is, literally, to bring other symbols into visible (or audible) relations with one another. There are possible conventions that could dispense with them. Thus "*a* is next to *b*", a complex sign, could be rewritten, under a suitable convention, as "a_b; and "Red (*a*)" or "*a* is red" could be rewritten, in another system of conventions, as one sign with a special design, such as that of a Gothic $\mathfrak{A}$.

Though meaning is not a relation in any ordinary sense, and in particular not a relation between a word and the world, words stand in the relation of being linguistic representatives of objects in the world. In "*a* is red" the symbol "*a*" is a linguistic representative of a particular object *a*. The sentence as a whole—that complex sign, which could be rewritten as various other signs, complex or simple—invests "*a*" with the character of belonging to a class of linguistic representatives of red things. Put another way, "*a*" the linguistic representative of *a* and "... is red" (where . . is to be filled with "*a*" or some other name of a particular object) is the linguistic representative of the class of red things.

It must be understood, in this connection, that classes themselves require careful interpretation. They are produced by classification, and classification operates with flexible criteria, changing between criteria as classification changes between serving one function to another. Classification is a means of picturing the world so that we can get about in it. Picturing is a kind of mapping. So, in the end, on my account of meaning, is truth.

The generic notion of truth is the notion of *semantic assertibility*. To say that "2+2=4" is true or "a is F" is true is to say that we have a right, in accord with the rules relevant to the kind of discourse in question, to assert the sentence "2+2=4" or "a is F". Truth, then, is a ticket of permission, not a property. In the case of "factual discourse", however, the correctness of the performance is defined in terms of the correctness of the mapping. Hence, factual truth is more than just warranted assertibility. This answers the challenge to realism posed above and so our trip from the frying pan ends not in the fire, but on the solid ground of realism.

In the discussion of Sellars's paper, the question was raised, in just what sense was Sellars offering a theory of meaning. Sellars maintained, in reply, that he was not so much offering to explain what "meaning" means, as to explain, on the basis of data about semantics and syntax, what semantic assertibility amounted to and how the same thing could be shown to be semantically assertible in different languages. This was as much or more a theory of truth as a theory of meaning. Truth is semantic assertibility; and semantic assertibility rests on causal conditions and pragmatic conditions fulfilled by speakers as much as on anything isolated under the head of meaning. "What," Churchland asked, "is the relation between semantic assertibility and empirical assertibility—assertibility about the world?" "They are the same thing," said Sellars.

Crispin Wright, from St. Andrews, supported by Blackburn, asked why one couldn't argue the other way about sign designs. Should one argue that elimination (by the Gothic $\mathfrak{A}$) of a separate symbol for the predicate "red" showed that it was just an auxiliary symbol without anything of its own to name? Why couldn't one argue that the single symbol, Gothic $\mathfrak{A}$, was more complicated than it looked, in effect containing a name or referring term for redness that in a fully adequate expression of logical form would be given a separate symbol? Sellars was not impressed by the suggestion that these were equally good arguments.

Someone asked for further light on the bearing of Sellars's theory of semantic assertibility on the issues of scientific realism. In answer, Sellars gave an *ad hominem* illustration: If Bas van Fraassen (once a student of Sellars's) had a more adequate conception of truth—closer to semantic assertibility—van Fraassen would be driven toward scientific realism.

The next speaker was David Gooding, from the University of Bath—a Dalhousie graduate in philosophy and a Dalhousie M.A., who afterwards took a D.Phil. in the philosophy of science at Oxford. He is an expert on Faraday's significance for the history and philosophy of science and drew specific examples for his paper from the researches of Faraday's time and milieu.

David Gooding
The University of Bath

Experimentation, Representation and Realism

This paper addresses an implication for realism of the constructedness of observational evidence. Scientists make observations and produce data in circumstances that are contrived and often difficult to repeat. The difficulty of getting experiments to work is a commonplace. Yet scientists remain confident that they often do succeed in discriminating aspects of nature from artefacts of their instruments and techniques. According to realists such confidence is to be expected: the existence of an independent natural world makes scientific practice intelligible and is the only plausible and *scientific* explanation of the success of science, which they define in terms of cognitive growth, empirical adequacy, predictive success and convergence—the formation of a durable consensus about what sorts of things there are in the world.[1]

Recent social and historical studies of science, on the other hand, have emphasized the discontinuity of scientists' beliefs about what there is. They point out that the very phenomena about which scientists theorize are produced by observational, mensurative, interpretative and other techniques. They argue the wholly conventional basis of experience, of its representation, and of judgements about it. This position—*constructivism*—treats science as an enculturated system of conventions whose contact with the natural world is infrequent, tenuous at best, and possibly non-existent. Constructivism is a socialized version of idealism which purports to show that realism is neither tenable as an *interpretation* of scientific practice nor necessary as an *explanation* of its success.[2]

Constructivist studies do face up to the practical problems at the coal face where natural knowledge is mined, and they recognize the social aspects of that process. Their significance for realism lies less in their epistemological conclusions[3] than in their exposure of a weakness in realist assumptions about how scientists obtain, communicate and reach agreement about new information about the natural world. Realists tend to consider observation only as a passive noticing of entities (e.g. electrons, nebulae) and qualities (e.g. colour, charge)—of just those things in experience from which semantic ascent is readily made. But this excludes a multitude of activities necessitated by the fact that scientists must invent and communicate representations of effects, processes and relationships.[4]

I argue that constructivism disposes of the Boyd-Putnam 'mystery of convergence' by providing a plausible (and well-researched) alternative. But I am more interested in its implications for the argument that local consensus about particular phenomena would be miraculous if all observers are "optimally informed" in a language, *none* of whose terms actually refers.[5] Constructivists could argue that, far from being miraculous (on all except realistic assumptions), local consensus is to be expected. Scientists' ability to communicate about phenomena presupposes training in ways of making observations, interpretations and explanations. Being "optimally informed" is the norm. This challenge is instructive. I develop a realistic construal of success in producing new facts and phenomenological relationships through an analysis of *experimentation*.

Experimentation is the exploratory form of empirical inquiry which has been very productive of results that were 'crucial' in the Baconian sense of 'experiments of the fingerpost'. Such experiments introduced information about nature which (despite the constructedness of its representations) brought about important changes in existing views of the way the world is. Room had to be made for previously unrecognized empirical possibilities. But how do we know that the anomalies introduce information about nature? Where it can be shown that scientists' agreement cannot be explained solely in terms of a pre-existing theoretical consensus, we may infer that they have in such cases succeeded in discerning an aspect of the natural.

This defeats the premise from which constructivists would infer that we cannot have knowledge of an independent natural world. But it raises a further question: How does a new phenomenon, noticed perhaps by just one individual, enter the social domain of public experience and discourse, so that it can be seen and talked about by others? Consider an example, the experimental activity engendered by Oersted's discovery in 1820 of the magnetic effect of a current. This shows that phenomenal novelty is represented by trial interpretations which I shall call *construals*. The Oersted effect had an anomalous aspect, the apparently circuital or non-central character of electromagnetic interactions. This was both unanticipated and incompatible with prevailing beliefs, it made everyone a novice for a time: scientists had to learn to produce, modify and utilize the phenomenon without benefit of unambiguous, self-evident, authoritative guides to interpretation. Such cases (not uncommon in the history of science) are not discussed in the philosophical literature. There, typically, a single novice is initiated into linguistic practices well-understood by the rest (Wittgenstein, Quine) or a single scientist simultaneously discovers and baptizes new information on behalf of the rest of us (Putnam).[6]

In view of the fact that Oersted's contemporaries began with terms which were neither precise nor definitive, it is striking that consensus was reached so quickly about the existence of certain aspects of the phenomenon and about derivative effects: that this consensus did not foreclose upon possibilities of further new experience: that it was not affected by scientists' very different background assumptions (about, e.g. the nature of force, the priority of the needs of mathematical analysis [central forces] over appearance [skew forces]), and that agreement about observational possibilities remained even after divergent theoretical interpretations had been formulated. A striking example is Ampère's theoretical reduction in 1820 of the circuital aspect to a (somewhat unconventional) combination of conventional forces. This enabled him to deny the *reality* of the circuital action. But it did not prevent him experimenting successfully in 1822 with a construal of the phenomenon as circuital. These points show that phenomena can be represented so as to inform theory whilst retaining a measure of independence from it, and can even lead to new effects incompatible with it. (This exposes a problem with the usual notion of prediction, but that is another story).

These scientists were not working with 'observations' or 'theories' as these are usually understood. They worked with *construals* of the phenomena which enabled each member of the various research groups to do three sorts of thing: (1) to structure the new phenomenon in terms of one of its aspects and remember how to recognize this despite variations in other aspects of the perceptual field (e.g. *Oersted's 'conflict' operationalized in terms of the manipulation of magnetized sensors, of conventions about the direction of the current etc.*); (2) to make sense of their own and of each other's experimental manipulations and discourse about these, and to communicate this sense to others outside the laboratory (e.g. *mnemonic devices and images published by Brande, Faraday, Schweigger, and Brewster*); (3) to suggest further exploration which enlarged both phenomenal and instrumental domains (e.g. *'lateral' or 'tangential' action realized in*

Schweigger's multiplier and applied in Ampère's galvanometer; 'new magnetic motions' realized in Faraday's prototype dynamo and electric motor; Wollaston's 'rotations' realized by Ampère in 1822.) The two construals that became definitive of the phenomena (in the manner of Kuhn's exemplars) performed these social and cognitive functions more effectively than others. They embodied more successful solutions to the practical problems of representation, communication, and exploitation engendered by phenomenal novelty, solutions on which theoretical interpretation and explanation depended. This example suggests that observation is much less important than experimentation, the activity which produces construals.[7]

I conclude:

1) that the need for new construals cannot be explained by non-realists at all and that the success of those construals that were incorporated into theory is best explained by assuming that reality was made to collaborate by creating the problem and by limiting the range of possible solutions;

2) that these solutions were discovered through the simultaneous exploration of nature (the structure of electromagnetic interactions) and of culture (the possibilities of representation inherent in the existing theoretical consensus); thus construals became definitive of experience just where they successfully met both natural and cultural constraints;

3) and that the study establishes a *prima facie* case for *experimental realism* based on the commensurability of experimental practices, by showing that knowledge obtained in local, empirical contexts is necessary to theoretical knowledge of the sort postulated by more ambitious forms of realism.

Experimental realism claims that interaction with nature is a necessary condition of the communicatory and experimental successes described in the case study.

NOTES

1. The position is summarized in L. Laudan, "A confutation of convergent realism", *Philosophy of Science 48*, 1981: 19-49
2. Latour, B. and Woolgar, S., *Laboratory Life: The social construction of scientific facts*, Sage (Beverly Hills and London) 1979; K. Knorr-Cetina. *The Manufacture of Knowledge: An essay on the constructivist and contextual nature of science*. Pergamon (Oxford) 1981. Barnes, B. and Shapin, S., *Natural Order: historical studies of scientific culture*. Sage (Beverly Hills and London) 1980.
3. K. Knorr-Cetina and G. Mulkay restate the position in terms of a distinction between epistemic and judgemental relativism, in their introduction to *Science Observed: Perspectives on the Social Study of Science*, Sage (Beverly Hills and London), 1983, pp. 5-6. Epistemic relativism denies any form of realism: judgemental relativism is a-realistic in the manner of N. Goodman. "On Starmaking". *Synthese* 45, 1980; 211-215.

4. One of the most significant inventions in the history of science --the calculus—was a quantitative method of representing not entitles or qualities, but *change*.
5. N. Jardine, "Realistic realism and the progress of science';, in C. Hookway and P. Pettit, eds. *Action and Interpretation*. Cambridge UP (Cambridge), 1978.
6. D. Gooding, "How do scientists reach agreement about novel observations?", *Studies in History and Philosophy of Science*, 17, 1986: 000-000, and "In Nature's School: Faraday as an experimentalist", in D. Gooding and F. James, eds., *Faraday Rediscovered*, MacMillan (London), 1985 (in press).
7. Cp. I. Macking, *Representing and Intervening*, Cambridge UP (Cambridge), 1983 and D. Gooding, "Experiment and concept-formation in electromagnetic science and technology", *History and Technology*, forthcoming, 2. 1985.

Churchland began the discussion of Gooding's paper by asking whether Thomas Kuhn (author of the *Structure of Scientific Revolutions* [Chicago: University of Chicago Press, 1962, 2d ed., 1970], one of the landmarks in the current history and philosophy of science) counted amoung the "constructivists." Gooding replied that Kuhn was not a thorough-going constructivist, but his use of "gestalt" notions had greatly influenced the school, as was manifest in their current effort to rewrite the historiography of experimentation.

Some participants offered at least a partial defense of the constructivists. MacIntosh claimed that they would have different explanations of novelties, predictive success, and so forth, which would give a relativistic account of the stir caused by Oersted's results. Okruhlik maintained that constructivists, at least of the Edinburgh school, would not deny that experience plays a role. They emphasized the social component, but accepted "local" interaction with nature. Blackburn pointed out that if they were in some sense idealists, being such wouldn't make them anti-empirical. Idealists are not anti-empirical. They hold that if you want to know whether the tide is rising you go look.

Churchland, on the opposite tack, said that what's on the books in science at any one moment reflects all sorts of things, social structure among them. Why should it not reflect nature, too? The best strategy in the philosphy of science would be to develop a theory of how the impact of nature filters through the cultural screen, but this would be an enormous task. Gooding commented that in the last section of his paper he had taken a first step toward carrying out that very task.

Mary Frances Egan, a graduate student at the University of Western Ontario, followed Gooding with a paper challenging a position recently enunciated by her former teacher, Churchland.

Mary Frances Egan
The University of Western Ontario

Natural Kinds and The Pragmatics of Reference

Paul Churchland has argued[1] that the Kripke/Putnam theory of reference for natural kind terms involves commitments which are incompatible with scientific realism. I articulate a 'pragmatic' construal of the causal theory of reference as developed by Putnam (which I'll set forth in a moment), according to which the theory avoids the commitments in question and indeed comports well with scientific realism.[2] I then consider some of the connections between my pragmatic account and issues in the philosophy of science.

Churchland argues that the causal theory involves an implicit commitment to the ontology of common sense. If it *does* then it is ill-suited to a scientific realist program, because there is good reason to believe that theoretical developments will continue to suggest taxonomies which cut across the categories of common sense rather than smoothly reduce them. Commitment to the causal theory of reference for natural kind terms, then, appears to preclude taking a realist stance toward any system of categories which is incompatible with the 'manifest' framework of common sense.

I argue that Churchland is mistaken about the causal theory's alleged commitment to common sense ontology. Properly construed, the theory is ontologically neutral: it involves no commitment to any particular ontology.

The causal theory provides the following scheme for fixing the reference of kind terms:

> Something is an x (or is in the extension of "x", if "x" is a mass noun) if and only if it bears similarity relation R to *this* (stuff).

Two conditions must be satisfied for an instantiation of the schema to specify an extension for a term: (A) (the indexical element) The speaker must be suitably related to the object (sample) indicated by the demonstrative. Typically, the relation will be 'causal', in some sense to be specified. (B) The similarity relation R must be specified. The theory, therefore, provides an account of our reference-fixing practices for general terms, not just for natural kind terms. If R = 'has the same micro-structure as', the schema will typically pick out a natural kind; if R = 'serves the same purpose as' or 'has the same causal antecedents and consequents as', it will specify different sorts of functional kinds.

The second condition provides an answer to Churchland's worry. Unless R is specified, use of the schema fails to pick out a similarity class. The specification of R is often only tacit—contextual features will serve to isolate the relevant respect of similarity. These contextual features will typically include social practices. Putnam's twin earth example illustrates the point. It is an empirical fact about our long-standing social practices that we take the relevant respect of similarity for putative samples of water to be composition or micro-structure, not phenomenal indistinguishability, nor any set of functionally characterized properties. Thus, *given* these referential practices, if a sample isn't H_20, it isn't water. What referential practices *do* get established is primarily a pragmatic matter—a function of the interests and goals of the relevant linguistic community.

The point I urge against Churchland is that only in conjunction with some particular referential intentions (i.e. an assignment of a value to R) does the reference-fixing schema pick out a class at all. I then argue that our referential practices are in turn a function of our theories, in conjunction with our interests and goals, and may change in response to developments in theory, so that the specification of R for natural kinds, or for any kinds, may vary over time. The pragmatic construal of the causal theory, then, preserves a central insight of descriptivist accounts—that our access to natural kinds is mediated by theory—without requiring that we have a correct theory of a kind in order to refer to that kind. What *is* required for successful reference, on the pragmatic construal, is that our referential practices prove to be *appropriate in the long run*. Appropriateness is a pragmatic notion—our referential practices in science must be appropriate to the goals of theorizing (predictive success, explanatory power, etc.). Continued comparative success of a (set of) referential practice(s) in meeting cognitive goals is evidence that the terms of theories based on these practices do refer.

In summary, I have argued that the causal theory of reference for natural kind terms is properly construed as an account of our reference-fixing practices. On such a construal, it is empirical theory, not the theory of reference itself, which specifies ontology (which is exactly what the scientific realist wants). Since reference is a relation between language *use* and the world, the intentions of language users and other features of the context, including the relative virtues of background theories and assumptions, will partly determine whether a term refers (and what it refers to) on an occasion of use. In addition, the pragmatic account leaves open the possibility that two (or more) sets of practices might prove to be equally appropriate, and thus that the terms in both theories might successfully refer, even though they may cross-classify what we intuitively take to be the same domain.

NOTES

1. "Laws, Conceptual Progress and Representational Media: In Search of the Essence of Natural Kinds" (delivered at Simon Fraser University, February 1983), forthcoming.
2. My reading of the theory is, I believe, the interpretation implicit in the recent work of Putnam, although I suspect that there may be substantial points of disagreement with Kripke's version of the theory. In any case, I do not provide textual support for my construal, since my concerns in this paper are not primarily exegetical.

Churchland told Egan that her version of the causal theory of reference was much more palatable than Putnam's. But he wondered how one would succeed in preserving continuity of reference? How is something's being "of the same nature" to be construed across changes in scientific theories?

Matheson commented that the theory which Egan had adopted as a revision of Putnam's was Putnam's theory in one of his early articles. How does it allow for any referential continuity at all? How can two different theories have the same concept of microstructure? Egan conceded that there was a problem about relativism afflicting her contention respecting continuity; but she said that it was not an insoluble problem.

Blackburn asked, "Don't my children refer to water even not knowing microstructure? Aren't there two different questions—succeeding in referring to water; establishing just what water is?" Egan said, "I disagree that the children are referring to water. Consider their liability to mistakes about colorless liquids."

"Isn't it odd," commented David Braybrooke of Dalhousie, "that it might turn out we don't now succeed in referring? And, joined with Churchland's view of the onward march of science, this you seem to be committed to." "But I," Egan replied, "unlike Churchland, allow that some common sense theories may survive at the ideal limit." In answer to a question from Thomas Vinci, also of Dalhousie, about the larger significance of her concern with reference, Egan added that she wanted to connect reference with truth-in-the-long-run.

The speaker following Egan was James Brown of the University of Toronto. He had taught several classes in the Dalhousie department a *few years ago and was back during the summer to lead (with Kathleen Okruhlik) the class in the philosophy of science that got everybody on the scene ready for the peak week of the conference. The following is a summary by the author:*

James Robert Brown
The University of Toronto

Realism, Anti-Realism, and The Quantum Theory

For more than half a century philosophers and physicists have been struggling with quantum mechanics: How should the formalism, which is so successful in application, be understood? The question I am concerned with here is intimately related to the issue of interpretation. It is the issue of realism versus anti-realism and I am concerned with how versions of each fare with the quantum theory. The conclusion will be that all extant versions of realism and anti-realism are failures; everyone is on the wrong track.

The problem stems from trying to understand the ψ-function which represents the state of a sytem and is governed by the Schroedinger equation. Schroedinger thought it was a smeared out electron; Born thought it was a probability and that it represented our knowledge. If either interpretation had worked, things would be easy for us today. But the split screen type experiments show that neither is plausible. Instead, in the reigning Copenhagen interpretation, we get something quite bizarre. Observations create; what is, is somehow connected to what we can know. A traditional realist would understand the uncertainty principle, $\Delta p \Delta q \geq \hbar$, by saying that the Δ represents knowledge: knowledge of p implies ignornace of q. But the Copenhagen interpretation, on the other hand, says Δ is ontological: knowledge of p implies the non-existence of q. Einstein and other traditional realists (for the most part) have denied that the theory is complete. But this view has been seriously undermined by the recent Bell results[1].

Newton-Smith has a radical version of realism.[2] He has an ontological part (theories are true or false independent of us) and an epistemological part (we can have good reasons for choosing one theory over another). To save the epistemological part from underdetermination problems, he denies the existence of inaccessible facts, and thus abandons excluded middle. Ironically, this makes Bohr a (Newton-Smith) realist. But, this runs afoul of the ontological part since, on a whim, by making the appropriate measurement, we can create position (or momentum) facts-of-the-matter, thereby undermining the independence clause in realism.

While versions of realism are failures, anti-realist accounts fare no better at making sense of the quantum theory. Here I consider von Neumann's idealism[3] and Putnam's new anti-metaphysical realism.[4] Their respective views become clear in von Neumann chains. This is a series of physical systems such as an atom, a Geiger counter, a TV set, a human eye, a human brain, etc. Which is the system being measured and which is the measuring device? If the quantum theory correctly describes macro-objects then how do we describe the measuring process? To end this regress von Neumann posited a *mind* which was non-physical and had the remarkable power to "collapse the wave function" into a definite eigen-state. Putnam claims there is no objective cut in the von Neumann chain. Any cut is arbitrary and when it is made then things really are in superposition on the non-measuring side of the cut.

To defeat both von Neumann and Putnam I use a separated system consisting of two cameras taking two initially undeveloped photos of Schroedinger's cat. Since the two photos, when developed at a distance, are both pictures of a cat which is alive (or dead) the cat must have been alive (dead) before the undeveloped pictures were separated, otherwise we would have a miraculous correlation. (But this amounts to the view known as local hidden variables which, as I already noted, is implausible. What the counter-example really does is show that von Neumann and Putnam can't account for a remarkable correlation.)

A response to a related matter which is gaining currency is that there are superluminal connections between the separated systems, and that these signals, which are compatible with special relativity (understood operationally), account for the correlations in EPR and similar situations. To counter this I show that such signals must move backwards in time. This has the consequence that a system would be in an eigen-state *before* it was measured. Thus, a measurement would in effect *discover* and *not create* the thing it measured. So versions of anti-realism undermine themselves; they are no better off than their realist rivals. The quantum theory, I am afraid to say, has defeated all.

NOTES

1. See d'Espagnet, "Quantum Theory and Reality", *Scientific American*, 1979, for a good popular account. My attempt to cope with this, Brown, "The Miracle of Science", *Philosophical Quarterly*, 1982, I now think is not successful.
2. See *The Rationality of Science*, Routledge and Kegan Paul, 1981.

3. See *Mathematical Foundations of Physics*, Princeton, 1955.
4. See "Quantum Mechanics and Observation", *Erkenntnis* 1981.

In the discussion of Brown's paper, Matheson maintained that it was not quantum mechanics, but special relativity and Bell's inequality that defeated traditional realism. With Terrance Tomkow of the Dalhousie department, he was ready to maintain that quantum mechanics had nothing to do with scientific realism one way or another. Brown, of course, stoutly denied this.

Tomkow said that idealists would object that minds are not at places in space-time, as Brown assumed. Brown replied that idealists still had to deal with physical problems involved in time-reversing processes.

Roland Puccetti of the Dalhousie department asked whether the many-worlds interpretation of quantum mechanics was compatible with traditional realism. Brown said "Yes" and that this was its only advantage.

Peter Clark of St. Andrews wondered whether the simultaneous running of the two cameras (one for each of the undeveloped photos) was feasible. Brown said that it could easily be imagined if one introduced the idea of a mechanical arm. Clark went on to recall that Putnam has pointed out that the conjunction of two assertible theories is not itself necessarily assertible. Wasn't there something like this problem here, as a consequence of partitioning the situation? Brown considered this amounted to suggesting that separating the cameras already makes a quantum mechanical difference to the wave-functions; and took the suggestion under advisement.

Dave Davies, a graduate student from the University of Western Ontario, was the first speaker on the third day of the conference.

Dave Davies
The University of Western Ontario

Putnam's "Narrow Path"

In the first of his *Howison Lectures*, Hilary Putnam maintains that certain contemporary tendencies in philosophy press upon us the question whether "a narrow path can indeed be found between the swamps of metaphysics and the quicksands of cultural relativism and historicism". He further claims that "internal realism' can provide us with just such a 'narrow path'. These claims invite critical examination. I grant, for the sake of argument, Putnam's case against 'metaphysical realism'; I argue, however, that he fails to establish either (i) that 'cultural relativism' is an untenable position, or (ii) that 'internal realism' offers a genuine alternative to the doctrines which he wishes to eschew.

In section *I*, I focus on Putnam's 'internal realist' construal of truth. The 'Idealisation Theory' of truth takes the latter to be "an idealisation of rational acceptability": a 'true' statement is one which would be 'rationally acceptable' under "epistemically ideal conditions"[1]. I argue that Putnam's formulation of the 'Idealisation Theory' is defective in certain important respects, and that these defects admit of no obvious remedy which does not also cast serious doubts on the viability of Putnam's more general project. I note, firstly, Putnam's claim that his account of truth preserves the principle that a statement and its negation cannot both be true.[2] I argue that he must restrict the application of this principle to the sets of statements contained *within* particular 'versions';[3] for, if the 'Idealisation Theory' incorporates the principle in an *un*restricted form, it is not compatible with the conjunction of two further doctrines which are integral elements in Putnam's 'internal realism'—namely, scientific pluralism and scientific realism. I further argue that Putnam's more general pluralism—a pluralism which countenances true statements in sundry versions not 'reducible' to the versions of science—will commit him to some form of cultural relativism unless he avails himself of a principled distinction between those versions ('correct' versions) which are capable of generating or sustaining truths, and those versions which are not so capable. Such a distinction, I suggest, will require the notion of an 'ideal collection of versions', a 'collection of versions' which would itself be 'rationally acceptable' under 'epistemically ideal conditions'.

I then argue that Putnam faces the following dilemma: (1) If 'truth', construed as "an idealisation of rational acceptability", is to be a notion that has any 'cash value' (Putnam's phrase) for us, then it seems clear that 'idealisation', for the statements belonging to a given mode of discourse, must involve some form of projection from existing standards of 'rational acceptability' within that mode of discourse. But (2) if 'idealisation' *is* understood in this way, then the resultant theory of truth seems to commit its proponent (i.e. Putnam) to either cognitive relativism or cognitive 'imperialism', albeit of a somewhat 'sophisticated' nature. In arguing for (2), I maintain that Putnam's account is unclear with respect to both (i) the type of 'projection' from existing practices which is to count as an 'idealisation', and (ii) the standards which are properly applicable in assessing the 'rational acceptability' of beliefs in cultures other than our own. I further maintain that there is no plausible resolution of these matters which would enable Putnam to evade the presented dilemma. Nor can he appeal to his distinction between 'criterial' conceptions of rationality, on the one hand, and a 'normatively important' sense of rationality, on the other; for, so I argue, the latter conception is equally available to the relativist or the imperialist.

I begin section *II* by sketching a taxonomy of possible relativist accounts of truth. The taxonomy utilises a number of distinctions which are explicit or implicit in Putnam's formulation and defence of 'internal realism'. I distinguish:

(a) 'Naive' ('*V*') and 'Sophisticated' ('*S*') forms of relativism—the latter employ the notion of an 'idealisation';

(b) 'I_1' and 'I_2' forms of *S* relativism—these differ in their interpretation of the notion of an 'idealisation';

(c) 'Descriptive' ('*D*') and 'Normative' ('*N*') forms of relativism—the latter construe 'rational acceptability' in terms of Putnam's 'normatively important' sense of 'rational';

(d) 'Personal' ('*P*') and 'Cultural' ('*C*') forms of relativism—these differ as to the nature of the standards of 'rational acceptability' of which truth is a function, and

(e) 'Total' ('*T*') and 'Objective' ('*O*') forms of relativism—the latter maintain that the truth of a statement relative to a given set of standards of 'rational acceptability' is *not itself relative*.

I can now evaluate the force of Putnam's explicit arguments against relativism when applied to the various forms of relativism allowed by the taxonomy. I grant to Putnam certain arguments which effectively rule out any *V* or *D* forms of relativism and any I_1 form of *S* relativism. I then examine a further argument which purports to rule out *any* form of *S* relativism. If this argument is valid, and if the *V*/*S* dichotomy is an exclusive one, then we seem to have a conclusive refutation of cognitive relativism. The argument draws upon Putnam's more general claim that the relativist must be able to distinguish between 'being right' and 'thinking one is right' if relativism is to be a tenable position, and that this distinction requires some notion of 'objective fit'. Putnam then argues that the *S* relativist cannot avail himself of the latter notion by appealing to the notion of an 'idealisation' of 'rational acceptability', because such a strategy undermines the 'relativist' nature of his position. I point out, firstly, that this argument carries no weight against *O* forms of *S* relativism, but I also accept (and offer further argument for) Putnam's unargued contention that a *cognitive* relativist must be a *T* relativist. Putnam is incorrect, however, in his claim that *T* relativism is "unintelligible". And, so I further argue, his argument against *S* relativism equivocates on the notion of 'true-for-X' and imposes the unreasonable demand that the relativist notion of truth be explicated in terms of a *non*-relativist notion of 'objective fit'. I conclude, on the basis of the above considerations, that Putnam's explicit anti-relativist arguments fail to refute at least one form of relativism, namely, the I_2 form of *TSNC* relativism.

There is a further anti-relativistic argument which I find implicitly contained in certain of Putnam's pronouncements. This argument attempts to ground the claim that there is "objectivity humanly speaking" in general features of human biology and culture, and, if sound, it would both (i) provide Putnam with a refutation of I_2 *TSNC* relativism, and (ii) allow him to evade the dilemma posed in section *I*. I suggest, however, that the soundness of the argument is highly questionable, and that, even if it were sound, its employment as adjudicator in the dispute between 'internal realism' and 'relativism' presupposes a certain understanding of the relationship between philosophy and other modes of discourse—an understanding the compatibility of which with 'internal realism' is itself distinctly problematic.

NOTES

1. Hilary Putnam, *Reason, Truth and History* (Cambridge: C.U.P., 1981), pp. 55-56.
2. *Ibid.*
3. Here, and below, I employ the term 'version' in the Goodmanian sense recently appropriated by Putnam. See Nelson Goodman, *Ways of Worldmaking* (Indianapolis: Hackett, 1978), and Putnam, *op. cit.*

Blackburn began the discussion of Davies's paper by questioning whether it was legitimate for an internalist like Putnam to talk about there being several versions (of the world-picture) at the ideal limit of inquiry. Wasn't this possible only in a God's-eye perspective? Davies replied that the internalist position essentially embodies the notion of a collection of versions. Blackburn retorted, "Already I'm supposed to be persuaded that there will be a variety of equally eligible views at the ideal limit. But why should I admit this? . . . It seems, oddly, to be a condition on the ideal limit as conceived by the internalist that we not be able to say a contradiction between theories means that at least one theory is wrong."

In a comment moving in the opposite direction, Churchland declared that he was inclined toward relativism, not by Putnam's arguments, but by the idea of selection in an evolutionary process. But why suppose that the process was ever going to come to a stop? Will there ever be an ideal limit?

A Dalhousie student, Dave Jennex, asked whether there were not some constraints from reality that Putnam would admit. Davies replied that Putnam thought there was something apart from our versions of the world-picture, but it was not something that we can express outside a version.

The speaker succeeding Davies was Crispin Wright, Professor of Logic at the University of St. Andrews. A summary by David Braybrooke from his notes of Wright's presentation follows:

Crispin Wright
University of St. Andrews

Two Arguments Against Scientific Realism

(Wright contended that the notion which "won't marry with scientific realism" is the unacceptable notion that "what people say is dependent on theoretical background."

(His first argument began with a three-part distinction between (1) objectivity of truth for statements (2) objectivity of meaning (3) objectivity of judgement. Weakest of these, objectivity of judgement still presupposes a distinction between fact-stating discourse—how things are—and projective discourse—how they strike a set of people. In fact-stating discourse, disagreements have to be explained in certain ways: misunderstanding; ignorance; vagueness; mistake; prejudice. But this list will not suffice if all statements, including observation ones, are theory-laden. This opens up the possibility of intractable disputes, and fact-stating discourse collapses.

ways: misunderstanding; ignorance; vagueness; mistake; prejudice. But this list will not suffice if all statements, including observation ones, are theory-laden. This opens up the possibility of intractable disputes, and fact-stating discourse collapses.

(Wright's second argument pointed out that we have lots of beliefs—theories—associated with "red," but that one can still competently pick out instances of red without them, without, e.g., the belief that red things remain so in the dark. "That is red," "That is a tomato," "That is a torque converter," are all judgements that might be made in response to observations. The first is distinctive in that mastery of the distinctive appearance constitutes mastery of the concept. Thus it is more primitively obvservational. We must still distinguish between "looks red" and "is red." The former might be taken as basic and immune to theoretical change. We would not have to go so far perhaps with certain other judgements, about, e.g., shape and texture.)

At the outset of the discussion of Wright's paper, Tomkow wondered whether Wright had made the notion of objectivity of judgement clear. Couldn't one construct a rationally compelling argument concluding that Margaret Thatcher was boring? Wright replied that if objectivity could not be made out, all responses might be regarded as equally cognitive or non-cognitive.

Realists including himself, according to Churchland, acknowledged something like Wright's first argument about theory-ladenness undermining objectivity for observers with different theories, but relied on global excellence to save objectivity for a theory. Another way out of Feyerabend's semantic solipsism would be just to pick out certain statements as semantically privileged.

MacIntosh maintained that Wright reduced the observer to a looks-red detector without internal qualia. Matheson asked how the definition of "looks red" was to be made public?

Wright conceded that "looks red" might be mistaken, and that on occasion theoretical principles came into establishing its correctness, when it was a correct judgement. Was it to be supposed that theoretical principles came into play in every case?

The next speaker—the final speaker on the third day—was Duncan MacIntosh, a graduate student at the University of Toronto. The following is a synopsis by the author:

Duncan MacIntosh
University of Toronto

How to Put Reality Into Language*

To be a realist about a given expression is, among other things, to hold that it has a literal and determinate meaning given in its truth-conditions, their satisfaction making it true (otherwise false), these conditions involving only the states of an objective and independent world—those logically figuring in its meaning. A *scientific* realist holds, further, that the statements which express typical scientific theories are of this sort, and that in general, their truth-values can be discovered using the method of inference to the best explanation of the data in the production of the most globally adequate theories—those with the greatest possible simplicity, predictive power, and world-controlling capacity.

Realism simpliciter involves a theory of what expressions might mean. As a theory of meaning it is the view that an expression's meaning consists in what makes it true (rather than what might verify its truth, for instance, insofar as this is not logically component in its being true). There is a theory of meaningfulness—of that in virtue of which an expression is significant—which would naturally go with a realist theory of meaning. In general, an expression has a meaning in virtue of having conditions of appropriate and inappropriate use. (Symbols are arbitrary relative to what they signify, so no sign intrinsically signifies. Neither does anything else—e.g., mental states and platonic abstracta. Meaning is the result of an association between a sign and what it signifies.) If its meaning consists in its truth conditions, and using it consists in assigning it a truth-value, then the appropriate theory of meaningfulness would be that in ideal verbal behavior, the use of an expression is controlled by its truth-conditions, their obtaining making it appropriate to call it true.

Unfortunately, there is no extant combination of epistemologies and theories of the relation between language and world compatible with realism and the realist theories of meaning and meaningfulness. The most plausible theories of meaningfulness tend to be reductionist in restricting what could be meant to what could be told true in direct

 I am indebted to Victoria McGeer, Sheldon Wein, Jan Zwicky, Lynd Forguson, Danny Goldstick, and James Young for their helpful discussions with me on these matters. I am grateful to the Social Sciences and Humanities Research Council of Canada, whose doctoral fellowship support I enjoyed during the writing of this paper.

experience, while the most intuitive analyses of what is meant in a given expression make it a mystery how one could know its meaning. I here give a critical survey of such combinations in their standard incarnations as positivism, relativism, coherentism, and scientific realism. In the process I defend realism, and try to show where the other doctrines went wrong.

The positivists believed that only sentences uniquely and consistently truth-valuable in principle by the totality of immediate experience (*qua* sense-data) could be meaningful. This seems to limit what is to what could be known in topic-restricted experience, *contra* realism But they also held that as it happened, every significant contingent sentence was equivalent to a construction from such sentences, leaving over no truly meaningful, inadjudicable claims. However, this requires that every sentence be demonstrated logically equivalent to sentences about immediate experience. Unfortunately, as Quine and others have noted, no sentence about what is entails, in isolation from other sentences, claims about what should seem to be so in experience. Furthermore, no claim presumptively about experience is free of theoretical entailment or presupposition, so sentences make claims about experience only in conjunction with theory-sentences which seem not to. And even though experience seems able to bear on the truth-valuing of any sentence, given the fixing of suitable ancillary assumptions, it appears impossible to trace each assumption back to experiences which could give its meaning and decide its truth. Only complex networks of expressions relate to experiences, rather than each sentence individually. And networks incompatible in internal detail could be equally compatible with total experience, choices between them being therefore arbitrary, with experience underdetermining theory-truth. Michael Dummett's revival of positivism will similarly suffer from the difficulty of it being impossible to assign verification-conditions to expressions taken singly, making them thereby individually determinate in meaning. Thus, experience is inadequate to the truth-valuation of expressions, and if sentences are individually meaningful, as the realist would have it, this cannot be by an individual association with the experiential conditions by which they could be told true, because they just don't have such conditions individually. The positivist method of identicating the control conditions upon sentences with their truth-conditions, then, seems a failure.

How then is it that expressions are meaningful? Some scientific realists have no theory of meaningfulness (e.g. van Fraassen), but they must come to grips with it sooner or later. Others (e.g. Sellars and Churchland) have held, unlike the positivists, that there are in fact no systematic limitations on the conditions which we can recognize, because condition-recognition is not limited to direct experiential recognition. Every possible condition will either figure or not in the best possible explanation of the data, where the data consists in what our most globally adequate theory of perception says our physiologies make us competent to detect unaided in the physical universe, this being extendable indefinitely by measurement technologies and observational aids. So, more things than the positivists thought are directly recognizable, and the method of explanational inference makes everything else mediately recognizable. If some condition does figure as data, or in the best explanation, an expression remarking it is true; otherwise false. This seems to offer some promise of making every possible condition in principle a recognizable one, concordant with both realist scruple and a plausible theory of meaningfulness.

Unfortunately, there is nothing in the concept of explanation that says there must be a best one—rather, there could be several, incompatible with each other, and equally good by all global criteria of theory-evaluation. Choosing one among them as true would be arbitrary, and would make the world dependent upon this choice, rather than objective and independent, making realism false again. Also, since it is indeterminate what is true, and so indeterminate what in truth is the condition under which someone is using an expression, it would be indeterminate what the actual use-conditions of sentences allegedly expressing certain truth-conditions by this method really are. Thus, their meaning would in fact be indeterminate, again contra-realism. And the problem of Quinean holism remains: Global criteria of theory-adequacy apply globally, not to individual sentences independently. Likewise, as we have seen, for experiential data. How then is each expression in a theory individually meaningful?

Quine's response to the problem was to give up the notion that sentences *have* a determinate meaning, holding that adequacy to experience is just a boundary-condition on the global adequacy of theory, there being no objective choosing between equally adequate theories, nor indeed, between alternative ascriptions of meaning to individual expressions within theories. Claiming that one sentence gives the meaning of another is just the arbitrary regimenting of which sentences to hold co-assentible. This is conditioned as much by back-ground linguistic habit as by empirical fact. Such choices account for our intuitions as to the synonymy of certain sentence-couples, and the antonymy of others. Likewise, truth-claiming is just the partly arbitrary adoption of a verbal disposition supposed congenial to global integration. So it is not the obtaining of a truth-condition which makes an expression true, but linguistic fiat partly constrained by the dictates of empirical adequacy.

Rorty has gone so far as to hold that as theory infects all claims, there are not even objective empirical constraints on total theories. Theory and language-deployment are a function of nothing but arbitrary socio-linguistic convention about the use of expressions relative to each other, while truth is just maximized intra-linguistic coherence in the largest possible collocation of expressions. He drops as completely unworkable the project of conceiving meaning and truth as relations between expressions and an independent, extra-linguistic world, observable or otherwise, losing too, the world itself. Now, since for Quine and Rorty, the significance of expressions does not issue from their individual association with the empirical and extra-linguistic, neither does the empirical any longer serve as a constraint upon their possible subject-matters. Similarly for the relativists, who hold that no facts are brutely obvious, so that neither truth, meaning, nor belief may issue from the obvious. Again, expressions can now be unreductively about anything, though now, beliefs about it are to be explained by their real—usually socio-political—causes, truth becoming relative to those causes, as the sole determiners of truth-value ascription.

Now, against Quine and the Quinean elements in Rorty, I hold that one can't make sense of there being equally adequate alternative theories of things unless one can distinguish their component claims from each other in point of determinate meaning, and thus to have real alternatives. And expressions can't get such a sense from either being arbitrarily used (this being the very paradigm of meaninglessness in a sentence), or by having a use determined solely by the relations sign-designs might have to each other, rather than to things extra-linguistic in an independent world. One might know all the intra-linguistic usage-rules for the expressions in a foreign language, for instance, and still not have a clue as to the meanings of such expressions either singly or taken as a whole. Against Rorty's updated coherentism, I hold with the venerable objection that many incompatible collocations of sentences could be equally large and internally coherent, neither criterion then being adequate to the specifying of which collocation gives the truth, and neither concept being serviceable as the provisioner of meaning to expressions. (What does network size have to do with meaning, and how much can one infer about expression-meanings from the knowledge that a given set of expressions is coherent?) Against the relativists, I hold that if the control-conditions upon a sentence's use are really, in general, things logically irrelevant to its truth or content (like political ambition), it is a mystery how a sentence ever gets to have its meaning, thence to be even relativistically truth-valued. If the usage-condition for a given sentence is a political ambition, why isn't it *about* a political ambition? Similarly, where the control-conditions are thought to be explanatory criteria (Sellars), other sign-designs (Quine, Rorty), experiences conceived as epiphenomenal to possible realities (sense-data positivists), and so on—call this the irrelevant conditions objection. I accept too the objection against the Sellarsians that explanatory adequacy underdetermines theory-truth and theory-identity.

Though I think the positivists were wrong about the nature and objects of experience, I accept their implicit constraints upon an adequate theory of meaning: Sentences must have, as their determinate meanings, the conditions which control their use, and these must all be in principle recognizable in experience upon their obtaining. But I conciliate semantic and epistemic empiricism with realism, thus: I think, with Quine, that inchoate experience (*qua* sensible, proximate environmental stimuli) can bear on any expression given sufficient ancillary assumptions. Fortunately, this means that even the most abstruse but truly meaningful claim of science or metaphysics has some empirical significance in its association with the other expressions which set it up for experiential test in a given moment. This also affords a recipe for isolating the experiential significance of any expression, in spite of its having such significance only in a network of its fellows: For any network of expressions, there will be some set of experiential moments with which it is compatible, and another with which it is not. The experiential significance of a given expression within the network consists in those experiences which would exchange sets (i.e., become either newly compatible or newly incompatible with the network), upon a change in the truth-value of the given sentence, as a result of the operation of the intra-linguistic usage-rules which define the structure of the network (which rules are the invariant rules of logic). Two expressions in different languages or theories thus have the same meaning if they have the same effect on the experiential associations of their respective networks. In arguing this I hold that, far from every claim being theoretic in the sense of holding independently of experience, every claim can in fact be about the experienceable (not reductionistically, but *sui generis*—I am expanding the experienceable), given suitable contextualizing assumptions, in the right circumstances. This should seem all the more plausible when we take to heart the failure of sense-data theory as the proper explication of the nature of experience, and substitute for it the naturalistic account in which experience is of the world, not of epistemically minimal, subjective states of the self. It will seem more plausible still when we accept the scientific realist's notion that technology indefinitely expands the range of the recognizable (as, for example, when, assuming the truth of particle-theory, and the normal operation of a geiger-counter, we take its beeping to disclose to us the sub-visible event of a particle emission).

Since every condition on this construal becomes one that one can in principle come to recognize to obtain (if it does), all other criteria of theoretical adequacy or truth than empirical adequacy become mere heuristics in its service. And now, total empirical adequacy (over the infinitude of data, with underdetermination remaining for finite ranges of data) uniquely identifies the true total theory.

Quinean holism, rather than being the death of positivism and the obstacle between empiricism and realism, ultimately negotiates their union. Theory of meaning no longer offends realism, because it imposes no *a priori* material restriction on what there might be. It merely explains that possibilities get into language by there coming to be a construal under which extended experience bears on claims respecting them, while what can be experienced (*qua* recognized) is in principle open-ended, just so with what might be.

All of Quine's indeterminacy challenges are answered, and the principal impetus to relativism and to Rortyanism is destroyed. There returns obvious truth, about a world directly evidenceable.

"Don't realists claim that there are inaccessible facts?" Puccetti asked as the discussion of MacIntosh's paper began. MacIntosh said that he was a sort of realist who doesn't admit such inaccessible facts. Puccetti objected, "What about black holes?" MacIntosh maintained that the world looks different because of black holes.

To questioning by Tomkow, MacIntosh replied that theories may be underdetermined by any finite set of data. Matheson argued that if so MacIntosh had, ironically, given a nice coherent rationale for relativism. But MacIntosh maintained that if two theories have the same observational consequences, over the infinite totality of all possible observations they are the same theory; otherwise we have no acceptable theory of meaning.

This seemed very doubtful to Churchland. Moreover, couldn't we want to change our ideas of how the world looks? MacIntosh said that he was willing to be liberal on this point.

Blackburn asked whether MacIntosh held that sentences are determined as assertible only by the totality of data. MacIntosh said, "Yes. But the hypothetical aspect of my theory is no objection, given at least that I have unlike Quine logically relevant facts of the matter." Tomkow commented that according to MacIntosh disputes about truth or meaning will only be settled in the long run. "But consider," Tomkow went on, "what people meant thousands of years ago. Wasn't that determinate?"

Wright asked whether MacIntosh was not propounding a sort of phenomenalist reductionism to the way the world looks. MacIntosh answered that "looks red" implied that the world is red on the hypothesis that the observer is normal. Blackburn commented that perfect satisfaction of the condition of looking like a fire hydrant implied that there was a hydrant.

The organizers of the conference had been watching the weather day by day, trying to choose the day that promised most sunshine for an outing to the beach. With near perfect success, they called for a day off after the third day of discussions. In an assortment of vehicles, the participants jitneyed down to Bayswater for a day of picnicking, field sports (football and frisbee), and ocean swimming. Cavorting in the surf, they gave new meaning to the expression, "a school of philosophers," and returned to business the next day much refreshed and as argumentative as ever.

The first speaker after this break was William Seager, a newminted Ph. D. from the University of Toronto, and now an assistant professor there.

William Seager
University of Toronto

Credibility, Confirmation and Explanation

One who thinks, as I do, that to give an explanation you must believe in the truth of the explanation, must face a dilemma that runs somewhat as follows. Any theory T can be split into two parts so that T = T′ + E, where T′ is the assertion that T is empirically adequate and E is the set of theoretical claims that T makes. It is clear that any explanation which goes beyond the assertion of some empirical corrrelation deducible from T (and hence also from T′) will make reference to E in some way or other. It will thus employ information in excess of that in T′, hence an increase in explanatory power would appear to necessitate a corresponding decrease in probability (degree of confirmation).

This is the crux of the dilemma: generally, an increase in explanatory power is thought to represent an improvement in a theory, but it seems we have just shown that this cannot be a confirmational improvement—not something which gives us more reason to believe the theory. But it also seems from the foregoing that without belief one cannot appeal to a theory for serious explanation (or, perhaps more realistically, without the expectation that at least some of our theories, maybe with some modifications, will eventually be belief-worthy, we could not offer any theoretical explanations). Are we thus doomed either to explanatory silence or the steady erosion of our epistemic position?

Perhaps we can escape without springing the trap. Consider for a moment the sort of probability which is assigned to theories. Without a wholesale subscription to a subjectivist interpretation there can nonetheless be little doubt that this probability is not a matter of frequencies or propensities. It is, rather, the degree of belief which a rational agent ought to maintain with regard to the theory in question. We can suppose that for each rational agent, a, there is an associated "degree-of-belief" function (which must, for familiar reasons be a probability function), P_a. A theory, T, can be represented by one (rather long, no doubt) statement which is assigned some value by P_a. New evidence will shift this value according to some conditionalization scheme. Since P_a is a probability function, it is unquestionable that adding information to T will reduce the value P_a assigns to it no matter what the evidence (save for the case where we add probabilifying evidence to T, but this would hardly count as a theoretical extension.)

However, the question which arises here is whether all learning of this sort can be modelled by some sort of conditionalization, within a single probability function. As an example, consider a lover given to jealousy, but who at the moment has almost complete trust in her lover. But suppose a pattern of events which are somewhat puzzling on the hypothesis of faithfulness but are by no means evidence of unfaithfulness. However, if the jealous lover discovers a consistent story which explains these puzzling occurrences and which involves the unfaithfulness of her lover, she may suddenly take up the belief in her lover's inconstancy. She will take up belief in her story (or at least increase its probability which was, by hypothesis, initially very low). Could it ever be rational for her to do so? I think so. Her new "theory" has more explanatory power than her old (which had none with respect to the relevant phenomena—there was no explanation of the puzzling behaviour) but is taken to be more probable.

But this cannot be unless we postulate not only a change in probability-value-on-evidence, but also the occasional alteration of the probability function itself. When the evidence suddenly falls into place, as it were, and we say "now I see!" we at least seem to have a better understanding of the world. And this seems to be an epistemic matter as well. Further, it is the articulation of the new theory rather than the accumulation of new evidence that occasions this response. The sudden new way of looking at the world reorders one's plausibility scale.

This new plausibility scale is no more than a new subjective probability function, within which the values assigned to various statements may have radically changed from the values assigned by the previous function. The change may be radical in that it does not follow the rule of conditionalization—the change may occur without the receipt of any new evidence, merely with the articulation of a theory. The change may also not follow the rule of so-called "Jeffrey's" conditionalization; that is, the values of conditional probabilities may change as well as that of absolute probabilities—the new theory may well suggest that evidence bears in ways that differ from those codified by the original probability function.

This last sort of change is one that seems to occur very frequently in science. One simply does not see the relevance of, say, increased pollution to the colour of moths until one has articulated some sort of selection theory. A detailed working out of one example of this can be found in [1]. There it is shown how the unification brought about by the acceptance of Newton's theory of gravitation suddenly made such previously irrelevant facts as the motions of Jupiter's satellites evidentially relevant to Galileo's laws of projectile motion.

Perhaps the change in plausibility ordering which occurs upon the articulation of a new theory is like cases of suddenly reordered perception. This is the more likely if, as the cognitive psychologists suggest, perception is in part a matter of hypothesis articulation. It happens often enough that a scene changes its appearance radically without the input of new evidence or data but simply because a new interpretation strikes us. This view is reminiscent of Hanson's in [2]. The book is better remembered for the notions of observation and theory-ladenness, but perhaps it is more important to consider its examination of occasions of sudden changes in our way of looking at the world brought about by theory articulation. In fact, Hanson explicitly likens the apprehension of a new theory to re-interpretation of the world. Talking about Kepler he says: "the difference between "librations vs. ellipse" and "librations = ellipse" is like the difference between the bird and the antelope. . . ." [2], p. 83 (the last remark refers to one of those now philosphically familiar ambiguous figures.)

This sort of change in plausibility assignment is also something like literary criticism. A new interpretation of a character may suddenly strike us and be plausible simply because of the way everything "fits in" on that interpretation. Needless to say, literary critics hope to convince us that their interpretation is correct, not merely that it "saves the phenomena" (for when the novel is laid out for you when you start, you start with the phenomena being saved or else your interpretation does not start at all).

This change may also be illustrated by the case of cryptograms. Suppose a complex signal was received by one of our radio telescopes. The proposition that this signal came from an intelligent source would be given low probability (since any natural interpretation would be preferred—the reasons for this are probably complex and interesting having to do with the perceived importance of the other hypothesis). If someone could articulate a theory or code which made the signal sensible then the probability that it was indeed a message would go up substantially (given that the theory was reasonable) even in the absence of further information.

The view I am defending is that we engage in science in order to interpret a puzzling world. Like the literary critic we start with and must abide by what the "book of the world" provides us—we must indeed save the phenomena. But as we know that the characters in a novel must be more complex than mere producers of the behaviour ascribed to them and no more, we know that there is more to the world than what we observe. Our interpretations take this into account, and, strikingly, put into order the world of appearance and go some way toward maintaining this order into future appearances. A new interpretation may strike us as providing a better way to order the world and when it strikes us thus it ipso facto becomes more plausible. This is why a theory can seem more likely than a proper part of it for until the theory is articulated that part may be assigned a very low probability (perhaps, as an example, the bald assertion in 1400 that the earth moves around the sun). But, of course, once the new plausibility scale emerges, this part will partake in the general increase in probability.

If reinterpretation changes probability assignments and if explanatory power increases with reinterpretation then explanatory power is a confirmational virtue, as well as an informational virtue. But if it is true that new theories or new interpretations actually alter rational individual's probability functions then the question of justification naturally arises. For while it is true that after the reinterpretation things "hang together better", and we have more evidence for certain of our pet propositions than we had before the reinterpretation (we are more secure, as Harper would put it), how do we know that we won't reinterpret away all this security at a later date? Hume asked how one could justify a certain inductive practice and suggested that there was no real answer—that this inductive practice was built into the human mind.

Just why certain views seem more plausible than others to us is a difficult question to answer. There was a time when "demon-possession" was a plausible explanation of certain sorts of behaviour. Why does this seem no longer plausible to us? The answer would seem to be that our present theories provide no room for demons in our world. Our present interpretation of the world precludes them. But this is an instance of a plausibility scale being reordered via the introduction of theory not a justification for it.

However, the full scale solution to the problem of the "reliability of reinterpretation" cannot be attempted here. I will rest content with the reconciliation of credibility and confirmation which "theory driven plausibility reordering" can provide. I think there can be little doubt that plausibility scales do alter with the articulation of new theories. It also seems clear that what we judge to be plausible we judge to be worthy of belief (more or less). Accepting a theory is then, in part, accepting a way of looking at the world. A new way of looking may suggest the modification of old likelihoods—and this is a matter of what we take to be true, or as likely to be true.

NOTES

1. Harper, William. "Consilience and Natural Kinds", Xerox.
2. Hanson, N.R., *Patterns of Discovery*, Cambridge University Press, 1958.

Richmond Campbell of the Dalhousie department won Seager's approval for a suggestion that the conflict between confirmation and credibility be treated as analogous to the conflict between maximizing true belief and minimizing false belief. Some sort of balance had to be struck between the two. Seager thought that van Fraassen might agree and, agreeing, suggest that the balance should be stuck at the point of empirical adequacy.

Vinci asserted that there were three notions to consider: degree of confirmation; probability; and degree of credibility or rational warrant. Carnap had identified degree of confirmation with rational warrant. He also identified degree of rational warrant with degree of probability. Why this? Seager replied that you are not more rationally warranted unless you have more probability;

to which Vinci retorted that perceptual statements may be rationally warranted without having an assignable probability.

Normore asked, "Does lack of belief vitiate explanation? In the 14th Century people didn't believe there were points in the continuum, but constantly cited them in explanations." Seager rejoined, "Could they have had the real answer?" Normore said, "They would have held, 'a real answer', but not the true one."

Matheson took up the same theme. If explanations had to have true premisses, all scientific theories, along with all folk theories, have been false, and hence could hardly serve. Seager replied that either explanations don't have to be more than empirically adequate or not all past explanations have been false.

Clark said that the history of scientific practice made it implausible to suppose that probabilities have figured importantly in accepting theories. Connectedness has been much more important than increases in probability.

George Pappas, of Ohio State University, who had come to Halifax not only for the conference, but also to teach in the Dalhousie summer session, was the next speaker.

George Pappas
The Ohio State University

Explanationism

Taken in a very broad way, explanationism is the view that explanation is importantly related to other matters. Thus in methodology we are familiar with Occam's Razor which may be construed as the injunction not to posit entities that are not needed for some explanations. There is also the positive side of the Razor: *do* posit all that is needed for explanation.

Explanationism, though, is not confined to methodology. Occam's Razor has epistemic overtones if it is stated as: If some entity, X, is not needed for a genuine explanatory purpose, then there are no X's—or, at least we have reason to believe that there are no X's. There is also a corresponding positive claim: If X's are needed for a genuine explanatory purpose, then there are X's—or, at least we have reason to believe that there are X's. These two epistemic versions of Occam's Razor lie behind a fair number of the claims distinctive of scientific realism. For instance, scientific realists (hereafter: SR's) often maintain that (some) scientific theories are literally true; that the theoretical terms of such theories succeed in denoting; that the denotata of such terms are theoretical entities; and that such scientific theories provide the best descriptions of the physical world—or, more tendentiously stated, the *only* correct descriptions of the physical world. If we add what SR's regard as a truism, namely, that the physical world is all the world there is, then the tendentious claim lately noted becomes the view that,

> . . . in the dimension of describing and explaining the world, science is the measure of all things, of what is, that it is, and of what is not, that it is not.[1]

Each of these SR theses has received a great deal of attention, and interest in them has intensified in the wake of the critical evaluation of the SR position by van Fraassen. Comparatively little attention, by contrast, has been focused on the epistemic claims lying behind SR, perhaps because those epistemic claims seem so plausible as to border on the truistic. Nevertheless, it is worth considering the background epistemology of SR.

We might try expressing what I referred to as the epistemic versions of Occam's Razor as:

> (1) Ultimately, the only justification one has for believing that there are (are not) X's is that X's are (are not) needed to explain some data D.

Here I have added the key word 'only,' since the epistemic versions of Occam's Razor stated above are obviously not strong enough to serve as support for the bold form of SR suggested by the quote from Sellars nor for the other SR doctrines listed earlier.

Actually, (1) is too broad a thesis for SR purposes. For, (1) puts no restrictions on the sorts of explanations involved, and no restrictions on data D. Thus, (1) might serve to justify the existence of religious entities, contrary to the spirit of SR. Hence, we would want something such as,

> (2) Ultimately, the only justification one has for believing that there are (are not) X's is that the postulation of X's is needed in (is not needed in any) some scientific theory to help explain some data D.

Statement (1) expresses a version of what I call 'epistemological explanationism.' It tells us that epistemic justification in positive or negative existential statements is tied to, is *essentially* tied to, explanation. Statement (2), by contrast, expresses a version of *scientific* epistemological explanationism and so (2) is a much stronger thesis, since it is considerably more narrow.

Both (1) and (2) seem initially quite implausible. Against (2), consider the statement "There are holes in the ground next to the fence." The postulation of holes next to some fence is surely not needed in any scientific theory, yet people are often justified in believing such a statement. Or consider "There are spots which I see right before my eyes." This statement *might* explain some data (e.g., the fact that I have blinked), but it *need* not do so (I might not have blinked nor even been disposed to do so). So, we have reason to reject (1), for people are often justified in believing that there are spots right before their eyes. It seems easy to raise problems for the background epistemology of SR.

There is, of course, more to be said. Statements (1) and (2) are both instances of *coherence* theories of epistemic justification. The basic form of such theories is this:

> (3) A statement, P, is epistemically justified for a person, S, if and only if P is a member of a system of statements, K, and the system K is coherent. coherent.

The concept of coherence needed here can be, and has been, explicated in different ways. Minimally we can say two things: (i) the system K is the system of statements understood and believed by the person S, and (ii) K is logically consistent. The actual members of any such system of statements will thus differ from person to person, and from time to time for the same person.

Historically, idealist philosphers accepted (3) and explicated coherence as some sort of logical relation between the members of K.[2] Another possibility would be to explicate coherence as an inductive relation: roughly, K would be coherent (or, to use a term of C.I. Lewis, congruent) if (i) and (ii) hold and each member of K is strongly inductively supported by the conjunction of the other members of K.[3] There is also the subjective coherence theory defended by Lehrer: on such an account a system K is coherent, given (i) and (ii) just in case each statement in K is believed (by the person S) to have a better chance of being true than any competitor of that statement has.[4] However, none of these ways of filling in details for (3) is useful to the SR. What is needed, instead, is a notion of *explanatory* coherence. I know of three ways this idea might be analyzed. The first, which is sometimes attributed to Quine, can be expressed this way:

> (4) A system of statements, K, understood and believed by person S is coherent just in case K is consistent and each member of K is explained by the conjunction of the remaining members of K.

Another, once discussed but later rejected by Keith Lehrer,[5] is this:

> (5) A system of statements K understood and believed by person S is coherent just in case K is consistent and each statement in K is *either* (a) a member of some sub-group of K which serves to explain some member of K not in that sub-group, *or* (b) is a statement which is explained by some sub-group of K but which is itself not a member of any sub-group of K which explains other members of K.

These claims are complex in wording, but simple in content. (4) just amounts to saying that in addition to consistency, each member of K must be an explain*er* in order for K to be coherent; and, (5) just amounts to the idea that in addition to consistency, each member of K must be either an explainer or an explain*ee* in order for K to be coherent.

The claim with which we began, namely (1), is readily seen as a version of the combination of (3) and (4). Moreover, it seems that (1) has no more plausibility than does that combination. Unfortunately, the combination of (3) and (5) is considerably more plausible than (3) and (4), and the reason is quite simple: some statements are justified because they are (mere) explainees and this eventually is not covered by (3) and (4).

The combination of (3) and (5), however, does not fare much better. Imagine an individual, such as a child, whose system of beliefs is "impoverished" in the following way: there are no statements which he believes and understands which explain his belief that wasps sting people. Nevertheless he is justified in this belief based simply on the fact that his mother told him that wasps sting people.

A third way of capturing the notion of explanatory coherence results if we replace (3) with,

> (6) A statement, P, is epistemically justified for person S just in case P is an explaining member of, or an explaining member in, a system of statements that has maximal explanatory coherence for S.

Maximal explanatory coherence, in turn, is defined:

> (7) A system of statements, L, has maximal explanatory coherence for person S just in case among the consistent systems of statements understood and believed by S, L explains more of what is to be explained than any other system of such statements.[6]

The idea behind (7) can be illustrated by noting that a person will typically believe a great many different statements, and by noting that the total number of believed statements can be broken up into different systems or sub-groups each of which is consistent. That sub-group that explains more of whatever is to be explained has maximal explanatory coherence for that person.

The theory given by (6) and (7) has been extensively criticized; I focus here on just one of these criticisms.[7] Relative to (7), how does one pick the statements *to be explained*? One approach is to select those statements that are individually confirmed by experience; thus, statements to be explained would be observation statements such as "I see small spots right before my eyes." Other statements believed by a person would be justified provided they serve to explain the statement about the small spots. However, it is clear that this option is closed to the SR who defends explanatory coherence, since the justification for the statement concerning the spots would derive not from explanation but rather from experience. Another approach is "social;" statements to be explained are those that would be assented to by nearly everyone under specific circumstances. The statement concerning the spot would be affirmed by nearly anyone in the right circumstances—when spots are present. Thus, what would serve to justify such an observation statement would be community agreement rather than experience.

To see a problem here, let us label the spot statement 'x'. Now consider the further statement "The statement 'x' would be assented to by nearly everyone in appropriate circumstances." The foregoing method of picking statements to be explained is adequate only if we are justified in believing the latter statement, but what would its justification be? Either it would be an explaining member of a maximally coherent system of statements or it, too, would be a statement to be explained. Presumably it is not the latter, so it would have to be an explaining member of the system. But what does it explain? We can *give* it an explanatory role. First we label it,

> *Y:* The statement "x" would be assented to by nearly everyone in appropriate circumstances.

and then we note that Y helps to explain another statement, namely,

> *z:* Person A assents to statement "x".

Of course, now our problem breaks out again with regard to *z*. We could solve it by stipulating that *another* statement, Y-#2 (= "The statement 'z' would be assented to by nearly everyone in appropriate circumstances.") has an explanatory role because it helps to explain yet another statement,

> *w:* Person B assents to statement "z".

It is thus clear that the method of community agreement, for selecting statements to be explained, leads to an infinite regress; hence, we need to reject this method of explainee selection.[8]

I know of but one other attempt to pick out statements to be explained. Consider a person at a moment in time, and assume that the person already has a stock of beliefs B. Suppose as well that at that very moment the person acquires three new beliefs. At that moment the statements to be explained are those which form the content of the three newly acquired beliefs. This selection is itself based on the observation that every belief, no matter what its content and no matter how it is acquired, has some small degree of initial credibility at the moment it is acquired. It may not survive after the moment: it may prove inconsistent, or its denial may be explained by the statements making up the person's standing beliefs. But this fact is immaterial. Which statements are to be explained is a matter which is relativized to a moment, and to whichever statements are acquired at that moment.[9]

This proposal faces two problems. First, the general principle it incorporates ("Every newly acquired belief has some initial positive credibility.") demands justification if it is a contingent truth. Any such justification, however, would proceed in terms of explanatory coherence, and it is hard to see how circularity or a regress would be then avoidable. Second, and perhaps more importantly, there are certainly many times when a person does not acquire any new beliefs. The present proposal would force us to say, quite implausibly, that at such times none of a person's standing beliefs are justified, no matter how well they functioned as explainers at earlier moments. An account with such a consequence is clearly not adequate.

We have been exploring different versions of explanatory coherence theories, each of which is a species of what I earlier called 'epistemological explanationism' Nothing has been said, however, about *scientific* epistemological explanationism, except in regard to statement (2) which is restricted to existential sentences. I will not here examine all of the different more general versions of (2) which might be available. Instead, we need only consider an analog of (3), namely:

> (8) A statement, P, is epistemically justified for a person, S, if and only if P is a member of a system of *scientific* statements, K, and the system K is coherent.

By the term 'scientific statements' used in (8), we will understand all those statements entailed by the theoretical statements of all currently accepted scientific theories. Any account which is based on (8) will have the immediate consequence that a great many people have very few justified beliefs. For most of the statements we ordinarily suppose people are justified in believing are not scientific statements. Moreover, should science wither away and be replaced by nothing, no statements would be justified for any person given such an account. A theory with such counterintuitive consequences has little initial plausibility.

Scientific realism may yet be shown to have secure epistemic underpinnings. But if I am right in claiming that some form of explanatory coherence theory provides the background epistemology for scientific realism, and also correct in alleging that no extant version of an explanatory coherence theory is plausible, then the conclusion to draw is that to the extent that scientific realism rests on the background epistemology, scientific realism is resting on sand.

NOTES

1. Wilfrid Sellars, *Science, Perception and Reality*, (New York: Humanities, 1963), p. 173.
2. For discussion of this, see Keith Lehrer, *Knowledge*, (London: Oxford University Press, 1974), chapter 7.
3. See C.I. Lewis, *Analysis of Knowledge and Valuation*, (La Salle, Illinois: Open Court, 1946).
4. Lehrer, op. cit., chapter 8.
5. Keith Lehrer, "Justification, Explanation and Induction," in M. Swain, ed., *Induction, Acceptance and Rational Belief*, (Dordrecht: Reidel, 1970).
6. These definitions derive from Lehrer, *Knowledge*, op. cit., chapter 7.
7. Criticisms may be found in Lehrer, *Knowledge*, and in James Cornman, "Foundation versus non-foundational Theories of Empirical Justification," in G. Pappas and M. Swain, eds., *Essays on Knowledge and Justification*, (Ithaca, New York: Cornell University Press, 1978). In what follows I make use of Cornman's ideas to some extent.
8. This problem is similar to a problem facing a foundationalist theory of justification, namely, that of accounting for the justification of foundational principles of inference from *within* the foundational theory.
9. This account derives from William Lycan, "Conservatism and the Data Base," unpublished manuscript. At the conference, Duncan MacIntosh, Calvin Normore and David Braybrooke raised points which were quite similar to this account.

Much of the discussion of Pappas's paper consisted of doubts variously expressed, about whether explanatory coherentism was not after all an adequate theory of epistemic justification. MacIntosh maintained several times over, that all that was needed for "cognitive control" was having some means of criticizing propositions. Normore suggested that something was epistemically warranted if it survives a certain process, taking beliefs in any order. He said, "This alternative has advantages for the realist, who doesn't need to hold that epistemic warrant matches the way things are, whereas an anti-realist needs an explanationist account. A realist can accept divergent maximally coherent sets of belief, arrived at from different starting points. For him, this just shows our limitations as epistemic engines."

Churchland spoke up. "I rise to defend foundationalism. Why could the principles of foundational epistemology not be justified on the same grounds as anything else? Is it self-refuting to try?" Pappas repeated some arguments against foundationalism, including the objection that it would be circular to justify the principles of foundationalism on foundational grounds. Churchland was unmoved. "That the principles should pass their own test," he commented, "is a virtue, not something to condemn."

What followed Pappas and the discussion of Pappas's paper was a double bill, Peter Clark of St. Andrews and Calvin Normore of Columbia, both speaking on the implications of the Skolem-Loewenheim Theorem for scientific realism. Clark was first at the lectern.

Peter Clark
University of St. Andrews

Skolem-Loewenheim Theorem

1 Realism and Conjectural Realism

Conjectural Realism is the idea that scientific and mathematical theories are attempts at correctly describing a structured determinate reality. This conception has two components, a metaphysical and a semantic one. The former component is just the idea that the Universe (of sets or of physics) has a definite determinate structure quite independent of anything we may say, prove, think or wish about it. The semantic component says that this structure is the *intended* interpretation for our theories (*i.e.* they refer uniquely to the physical or mathematical Universe) and we will have grasped the content of those theories, what they say about this intended interpretation, when we understand what structure the intended interpretation must have in order that the theories be true. Now since our theories are true or false descriptions of reality on this view it is perfectly possible that we could have an ideal theory which captures even all possible observational evidence correctly and satisfies every methodological norm, yet nevertheless that theory be false. This follows as a possibility just from the strongly transcendent character of scientific theory over experimental data. Similarly in the case of mathematics because mathematical activity and mathematical reality are quite independent on such a (Platonic) conception it follows that a mathematical theory which was perfect with respect to all possible data (i.e. all possible computational results of say number theory or even analysis) could still not correspond to the true mathematical Universe. Both conceptions are I think *incoherent*, that is that no sense whatever can be made of these possibilities either in the scientific or the mathematical case. My reasons for thinking so are based on those developed by Putnam in his "Models and Reality".[1] Since this argument of Putnam's, which is an extension of an older argument of Skolem's[2] has independent interest in the philosophy of mathematics and arrives independently at similar anti-realist conclusions obtained by Dummett and Wright[3], I think it worth systematic scrutiny.

2 *An Old Argument of Skolem's*

In 1919 Skolem extended a result first obtained by Loewenheim to the effect that any first order theory expressed in a countable language which has a model (an interpretation in which the theory is true) has a *countable* model. Now set theory (say that of Zermelo-Fraenkel) is a first order theory expressed in a countable language and so must have a *countable* model. But a famous theorem (the basic result) of set theory due to Cantor asserts that there exists an uncountable set. Being a theorem, the countable model of set theory must make this assertion true. But the domain of the model is countable so how can it satisfy the assertion of the existence of an uncountable set, for that would mean that there was an uncountable subset of a countable domain, problem! This 'paradox' is resolved by pointing out that all this amounts to is the discovery of yet another limitative result. By that it is meant that all the paradox shows is that *within* the countable model no function exists which will count (*i.e.* put into a 1-1 correspondence) a certain subset of the domain. Since no counting function for that certain subset exists, it is indeed uncountable. Of course in full mathematical reality such a counting function exists, but truth in a model is not truth in *the* mathematical universe and no-one ever thought it was. All this shows is that the 'intuitive concept of a set' is not captured by the axioms of set theory, nor will it ever be because "unintended" interpretations (like the countable model whose existence is guaranteed by Skolem's theorem) will always be around. Hence Skolem's result is usually regarded as another 'limitative' theorem of first order logic because it shows that no formal system can capture the intuitive notion of a set.

3 *The Inadequacy of the Limitative Conception*

But this 'limitative' concept of what Skolem's result shows is quite inadequate. This is so, simply because if the postulates of any formal system (first order and countable) cannot capture the notion of a set *nothing* can. Even if we set about formalising all our beliefs about sets, indeed all our mathematical beliefs and their relations to set theory (preserving consistency) this would not serve to rule out 'unintended' interpretations in particular denumerable interpretations. Now perhaps the mathematical realist will reply that the mind just has a capacity to grasp the intended reference, the real universe of sets, quite independently of what we *assert* in set theory. But it is a consequence of his position that this grasp cannot be expressed or communicated. For nothing by the Skolem argument he can say can serve uniquely to capture that conception and what sort of conception is it which is supposed to be attributed to someone or some community, if nothing even in principle they can say can serve to inform one of what that conception might be? I claim that no such conception is coherent. Note that one cannot have any evidence for a conjecture one might form as to what that conception is—because anything which may be construed as evidence, namely what the Platonist does in working on actual mathematics, in proving theorems for example on the basis of postulates just constitutes more Skolemisable sets of sentences and that, we have just seen, will not fix the "intended" reference of the concept set. Nor is it the case that appeal to second order formulations of set theory will help here. For this argument is not a case of first-order disease, there simply aren't any categorical axioms for set theory.

In summary then mathematical realism (of the kind which supposes a fixed domain—the mathematical universe—to which our conjectures correspond and which allows for a truth conditional account of the meaning of mathematical statements) is untenable. It is so precisely because of the ineliminable existence of unintended interpretations which means that nothing we can say can ever serve to specify the intended reference of the sentences. That we grasp the intended reference is crucial to the realist's account of how we understood those sentences. But there is no way in which this grasp can be achieved. What the Skolem 'paradox' really points to then is a need for a generalised constructivist account of the content of mathematics in terms of our capacity to recognise what would constitute a proof of a given mathematical statement rather than in our (non-existent) capacity to recognise what structure the mathematical Universe would have to have in order to make the statement true.

4 *The Application of the Argument to Science*

While the above considerations may seem plausible with respect to mathematics in its Platonistic construal, it would seem perhaps implausible in the case of natural science, where we causally interact with the intended interpretation of scientific theories, *i.e.* with the physical Universe. However, exactly the same considerations here apply. Every physical theory of significance embodies set theory (or a non-trivial fragment thereof). As such unintended interpretations of the set of such physical theories exist, nothing we can say suffices to pick out one unique intended interpretation. Thus Putnam's extension of Skolem's argument can be put like this: imagine we have got all possible data in and we have a scientific theory T which gets this data right. Imagine further that T satisfies all the methodological requirements that we care to impose. Now the possibility which exists according to realism (with which I began this paper) is that T, while it gets all the data right, could still be false, at the 'theoretical' level, *i.e.* what it says at the highest 'theoretical' level about the intended interpretation, the physical Universe, could well be incorrect. We have agreed that T satisfies all empirical and methodological constraints. But satisfying all empirical and other constraints means that it satisfies all articulable and other constraints on what is to constitute an intended model of the theory. Then if it is true (is satisfied by) all intended interpretations it is *true* and so the possibility that it in fact be false is foreclosed. In order to maintain the possibility that it is 'really' false one has to be able to specify what structure physical reality would need to have in order to make it false. But that constitutes simply laying (without our being able to write) down more operational constraints than these which we have agreed T already satisfies. Further, as we know from the Skolem argument no such set of constraints will specify uniquely the intended interpretation. In short the realist can never specify that he means 'true in this structure, *i.e.* physical reality' by laying down postulates true for the structure. But how else is he able to say 'T may well be empirically adequate but it is false here', (where 'here' means in physical reality) other than by describing what he means by 'here'. But no matter what his description it will *not* fix what he intends by 'here'.

If these considerations are correct then nothing can be made of that possibility which is an immediate consequence of conjectural realism in mathematics and physics that our theories satisfy all operational constraints and somehow miss the truth. But this possibility follows directly from the conjectural realist's contention that understanding scientific theories consists in having grasped what structure the mathematical and physical Universe must have in order that those theories be true. It is the truth conditional account of understanding which is really at fault here.

NOTES

1. H. Putnam, 'Models and Reality', *Journal of Symbolic Logic*, *45*, pp. 464-82.
2. Th. Skolem: 'Some Remarks on Axiomatised Set Theory', reprinted in J. Van Heijenoort (ed.): *From Frege to Godel*, pp. 290-301.
3. M. Dummett: 'Realism', *Synthese*, *52*, pp. 55-112 and C. Wright: *Wittgenstein on the Foundation of Mathematics*.

Calvin Normore took over when Clark had finished.

Calvin Normore
*Columbia University**

The Dialectics of Realism

The 'anti-realist' position recently articulated by Hilary Putnam and defended here by Peter Clark is founded in a view of language which I shall call 'pure holism' and is buttressed by arguments of a type first developed by T. Skolem. This position is by no means evident and the arguments which have been employed in its defence can be as readily used to undermine it. The result appears to be a standoff, but, I suggest, a standoff signals a realist victory.

Pure holism treats language on the analogy of a map without a key. To interpret the map one forms hypotheses about what its different features represent and uses these hypotheses in practice. Hypotheses which lead to practical difficulties are discarded. Of course a small map might contain too little information uniquely to determine a terrain, but larger maps would be compatible with fewer hypotheses

*Now at University of Toronto.

and one might hope for the limiting case of a map so large that only one interpretation of its feature is not ruled out. In that case the map and a Principle of Charity for mapmakers would together uniquely determine a terrain. It is this hope which Skolemite arguments undermine.

The success of quantification theory has led many to believe that the central parts of natural languages and most of our beliefs can be presented as theories within it. Quantification theory itself is a *logic* and so it is not surprising that it has many models but one might hope that the representation within it of all the true sentences of a natural language would have a unique model. This hope is dashed on the Loewenheim-Skolem theorem. It and a family of related results show that any theory formulable in quantification theory has models of every non-finite cardinality.

The claim that these results are metaphysically significant derives its plausibiity from the assumption that foundational theories in mathematics, set theory for example, must be understood as axiomatic theories formalized within quantification theory. But, as Paul Benacerraf has shown, Skolem himself originally understood his results as an argument against precisely that claim. It is easy to construct formal languages so weak that they provide the same representation for very different claims. (For example, a version of quantification theory which had only one atomic sentence would give the same translation to every truth). We do not usually take this as evidence that all truths are one but rather as evidence that the language is inadequate to express the difference between them. Why not adopt the same attitude toward quantification theory? Why not regard the fact that there are non-isomorphic models of every first-order language as a defect in the expressive power of first-order languages rather than as evidence that concepts like denumerability are not absolute.

Although Skolem's arguments can be read this way, Peter Clark takes them rather as showing that since there is no way of formalizing set theory within a first-order language so that the resulting theory has only models with uncountable domains, we have no way of saying and meaning what we think we say and mean when we make such claims as that there are uncountably many sets. All that we can coherently claim, according to Clark, is that the domain of discourse contains no enumerator set of the set of its subsets.

This claim, if true, would gut set theory of its intuitive content. A model for a first-order language is an ordered n-tuple one of whose members is a set—the domain of discourse of the model. If set theory is formalized in such a language and interpreted in this model then all of the sets postulated by the theory are members of the domain of discourse of the model. Thus there is a set containing *all* sets, the domain of discourse. But then if Clark is right there is no enumerator set of the set of all subsets of the domain. There is no such set within the domain, and, since all sets are in the domain, there is none without either. But then Cantor's claim that the cardinality of the power set of a given set is greater than that of the set itself cannot be expressed. So Clark's view forces us to conclude that perhaps the most fundamental result in set theory is not expressed by its 'natural' formulation. Indeed *it* cannot be said at all.

Perhaps we can only reject Clark's view in favour of another. One promising alternative is the suggestion that higher-order theories may be expressively more powerful than first-order theories in the desired ways. There are categoricity results for sufficiently rich higher-order logics. But Clark and others object that higher-order theories are only categorical relative to models which are themselves characterized within set theory. If set theory is a first-order theory then there will still be non-isomorphic models of the higher-order logics in the non-isomorphic models of the underlying set theory. At this point the realist about mathematical concepts should simply deny that set theory is formalizable as a first-order theory and take Skolem's results as a significant contribution to our understanding of the limits of formalization. Of course the non-realist may then insist that set theoretic concepts are only made clear within a formalized theory and recall for us the spectre of Russell's paradox. If the realist offers a formalization, the non-realist can then interpret it as a first-order theory. And so on. The result seems to be a standoff.

Matters get better when they seem to be getting worse. The problems for realism which are produced by treating languages as maps without keys can be raised even within categorical theories. They require not that theories have non-isomorphic models but only that they have distinct models. But that all first order theories have distinct models is a consequence not of the Loewenheim-Skolem theorem but of the traditional truth definition for formalized languages. Such a definition lays down structural conditions which must be met by the sentences of the language which are true in a model. For example if P is a monadic predicate and *a* an individual constatnt than "Pa" is true in a model only if the element of the domain assigned to *a* is among those assigned to P. It is easy to show that if Pa is true in a model M, it can, with a clever choice of assignment of objects to *a* and P, be made true in any domain as large as the domain of M. It doesn't matter *what* is assigned to the constants and predicates of the language, only that the set theoretic relations among the assignments be preserved.

How then is formalization ever possible? Interestingly, when a bit of language is being made to do logical work the set of structures accepted as models is constrained not by imposing the structural constraint that some new sentences be made true by every model, but by insisting that the models meet some new condition which is expressed metalinguistically. For example when identity theory is added to quantification theory it is not by insisting that the first order sentences expressing the properties of the relation be true in every model but by insisting that "a=b" is true in a model if and only if *a* is assigned the element of the domain which is assigned to *b*. Indeed it is by conditions like this, expressed in the metalanguage, that the quantifiers and connectives which characterize a first-order language are introduced. Thus to formulate the languages which he claims to have distinct models, the non-realist must rely on concepts which are not defined by structural conditions but are introduced absolutely by a description in the metalanguage. If the non-realist does this he owes us an account of why it is only logical expressions and not any expression with a sense which can be so treated. If the non-realist does not do it he can't characterize a first-order language at all.

Pure holism depends on a solution to this dilemma and I'm not sanguine about its chances. The alternative is to abandon the metaphor of the map and think of language piecemeal as a structure in which meanings are fixed not by the global role terms play in the language but by non-linguistic relations between them and the world. This is to treat languages as more like photographs than like maps. Perhaps the causal theory of reference points the way.

Suppose that both the realist and the non-realist positions are defensible. Such a standoff would, I suggest, support the realist. Realists have no principled objection to the possibility that there may be facts we cannot discover, even facts about language. But what could a non-realist make of an irreconcilable disagreement? If the only constraints on meaning are given by the requirements of satisfaction in a model then the realist and the anti-realist pictures of language, being satisfiable in the same models, are really one picture. But then it becomes a mystery why they seem so different and one should not multiply mysteries beyond necessity.

Clark began the discussion by challenging Normore, "Try making 'There are uncountably many' a part of logic as a quantifier." Clark went on, "A realist claims to have a conception of reality as having in advance properties quite independent of what he can say. But how can reality in this sense be fixed in meaning?" Blackburn asked, "Do not interpretations of what people say require more constraints—from behavior—than just expressions that correspond to a model?" Clark retorted, "Give me written down your theory of reference. I add my theory of truth. Now I Skolemize this combination," which was to say, treat it as a method for getting from one proposition to another without being committed to any unique external model.

Normore asked, "How in your view can one say anything false?" Clark said, "That's a problem for the realist, not for me." Blackburn objected, "But if anything one gives as an explanation of referring can be Skolemized, must not the anti-realist face the same problem in respect to speaking truly?"

In reply to an objection from Tomkow, Clark asserted, "The realist-antirealist argument gets off the ground when we are dealing with infinite collections and our verification procedures break down."

Wright contended that methodological constraints operated independently of the Loewenheim-Skolem Theorem. Then, in a further exchange, Clark asked Normore, "How can you divide epistemology from metaphysics?" Normore replied, "The structure of what we know is a different matter from origin and justification."

This report of the 1983 conference has so far presented the papers and the discussions in the order of their occurrence. Here, however, instead of bringing up Blackburn's paper and the discussion that it provoked, we shall put them off to come just before Okruhlik's paper. Both Blackburn and Okruhlik tried to give a general view that summed up the discussions of the peak week. Matheson's paper, which at the conference was given after Blackburn's, was more specialized.

Carl Matheson
*Syracuse University**

The Irrelevance of Realism For Commensurability

1. *Introductory Polemics:*

Contemporary philosophers of science attribute great medicinal powers to the realist thesis, claiming that it is the *panacea* for all or

* Now at Iowa State University.

most of the outstanding maladies in the philosophy of science, be they methodological, epistemological, or linguistic. I do not share this faith in the curative powers of realism. Realism, properly construed, is a metaphysical thesis. Those who wish to make a tasty soup from a stone must add tasty ingredients to their gravelly broth. Those who wish to base their solutions to epistemological problems on realism must add many epistemological premises to their stock of assumptions. And, just as the savour of the soup owes nothing to the stone, the epistemological power of the resulting position owes nothing to realism. In this paper, I show that, contrary to popular belief, realism is of no great help in meeting the challenge of the argument for the incommensurability of scientific theories.

2. *Background:*

Feyerabend[1] was the first to propose the incommensurability thesis. In the course of his attacks on the positivist claim that science progresses via the reduction of earlier theories to later ones, Feyerabend devised a theory of meaning according to which the meaning of a scientific term is completely determined by its place in the inferential structure of the theory in which it occurs. Same location in the inferential structure, same meaning; different location, different meaning. Given that two theories will rarely have the same inferential structure, terms will always be assigned different meanings by different theories. Thus the claims of different theories can (nearly) never be compared. If the claims of different theories can never be compared, then theories cannot be evaluated with respect to each other. One theory is neither better or worse than another; rather, the two theories are incommensurable. The possibility of scientific progress is a myth. Shifts in theory are shifts in fashion.

Following Scheffler[2] many writers have claimed that invariance of meaning is not necessary for inter-theoretic comparison. Instead, these writers have demanded only that the terms of different theories be invariant with respect to reference (or extension). Using versions of the causal theory of reference, Putnam, Devitt, Kitcher[3], and many others have attempted to show that the comparison of theories is possible. Most of these writers seem convinced that their position is tenable *because* of the truth for realism. For example, Kitcher closes his paper with a paean to realism:

> Trivially, there are just the things there are. When we succeed in talking about anything at all, these entities are the things we talk about, even though our ways of talking about them may be radically different. However variable the connections we draw among its constituents, the world supplies a common content for our references.[4]

Here Kitcher seems to be saying that relativism is false, that what exists is not relative to conceptual scheme or theoretical framework, and that, because relativism is false, the referents of most theories should be specifiable in terms of most other theories. Since realism entails a denial of relativism[5], realism also entails the falsity of the incommensurability thesis.

3. *Realism is Insufficient for Commensurability*

If the above is indeed what Kitcher means, then he is clearly wrong. Realism is insufficient for commensurability on two counts:

(i) The mere fact that relativism is false does not entail that the terms of your theory will be coreferential with the terms of my theory. The following is possible: the terms of phlogiston theory refer and the terms of oxygen theory refer, but the terms of the two theories share none of their referents in common. This is an entirely realist picture, because the entities of phlogiston theory and oxygen theory exist *simpliciter*. They form islands in the sea of all the things there are. However, since there are no bridges between the islands, there is no way to specify the referents of the one theory in terms of the other.

(ii) Even if the terms of two theories refer to the same group of things, that does not guarantee that anyone can tell that they do. If, say, 'phlogiston' and 'oxygen' are co-referential, that is of no use to the oxygen theorist in evaluating the claims of phlogiston theory unless he can find out that the two terms are co-referential. This constraint reflects more than a mere sceptical possibility that is to be dutifully considered and then forgotten. When one notes the great difference that modern theories of reference allow between the actual criteria of application of a term and what a speaker may believe to be those criteria, one should realize that we may have relatively little insight into the intension of 'oxygen'. After all, we grant that some phlogiston theories referred to oxygen with tokens of 'dephlogisticated air', although many of the properties that they attributed to dephlogisticated air are properties that we would certainly deny of oxygen. We say that the phlogiston theorists simply had false beliefs here; the procedures they used to fix the reference of 'dephlogisticated air' were just wrong. However, we must also realize that future generations may have the same opinion of our reference fixing procedures for 'oxygen'. Maybe we're just wrong too. Faulty scientific theories lead to misleading reference fixing procedures—and to misleading procedures for determining co-reference between the terms of different theories. These problems stress the true nature of the incommensurability thesis.

Its challenge is not metaphysical; rather it is epistemological. It proclaims that we can't tell what other theories say. If scientific realism gives us semantic comparability in principle while denying it in practice, it offers small solace to those who wish to allow for scientifc rationality as a practical possibility.

4. *Realism is Unnecessary for Commensurability:*

But perhaps we should be more charitable in interpreting Kitcher's closing tribute to realism. Rather than claiming that realism itself is sufficient for commensurability, Kitcher may be claiming that realism is a necessary component of any theory that would disprove the incommensurabilty thesis. Perhaps he is offering this argument for realism: (a) scientific theories are commensurable; (b) the referentialist approach provides us with the only way of explaining the commensurability of scientific theories; (c) the referentialist approach presupposes realism; (d) therefore, by inference to the best explanation, realism is true.

Kitcher may justify (c), the claim that the referentialist approach presupposes realism, as follows: The referentialist route presupposes that there is some measure of referential invariance between theories, i.e. that terms in different theories may refer to the same thing. But this entails that what exists is prior to theory. What exists is not determined by one's theory; one's theory merely determines which entities one talks about. In other words, ontology is not relative to theory: relativism is false. And, for the purposes of the present discussion, we can equate the falsity of relativism with the truth of realism. Therefore, (c) is true.

This argument for (c) is unsuccessful. Briefly, the referentialist account does not presuppose realism, because a relativist is as entitled to the concept of reference as he is to any other concept. Suppose that I, an oxygen theorist, say that both 'oxygen' and 'phlogiston' refer to the same thing. The relativist can construe my statement as "With respect to the conceptual scheme surrounding oxygen theory, 'oxygen' in the oxygen theory refers to the same thing as 'phlogiston' in the phlogiston theory." Those who accept the claim that truth can be relativised to theory should not balk at the relativisation of ontology to theory. The oxygen theorist views the phlogiston theorist's utterances through the lens of oxygen theory and, in doing so, he will assign referents to them according to his ontology. Similarly, when the phlogiston theorist makes judgements of coreference, he does so through the lens of his theory. His referential assignment will be made according to the ontology of phlogiston theory. The relativist can make full use of the referentialist approach without giving up his analysis of truth.

5. *Conclusion*

This paper should not be taken as an argument for relativism. No doubt, relativism is incoherent even by its own eldritch standards. However, if refuted at all, it will be refuted *via* the traditional method of metaphysical dialectic. It will not fall as the result of its failure to provide an interpretation for certain episodes in the history of scientific language.

NOTES

1. Paul Feyerabend, "Explanation, Reduction, and Empiricism", pp. 28-97 in *Minnesota Studies in the Philosophy of Science*, ed. H. Feigl and G. Maxwell; Minneapolis, 1962.
2. Israel Scheffler, *Science and Subjectivity*, Indianapolis, 1967.
3. Hilary Putnam, 'The Meaning of Meaning', pp. 215-271 in *Mind, Language and Reality*, Cambridge 1975.
   Michael Devitt, 'Other Terms', *Designation*, New York, 1981.
   Philip Kitcher, 'Theories, Theorists and Theoretical Change', pp. 519-47 *The Philosophical Review* vol. 87, no. 4, October 1978.
4. Kitcher, op. cit., 547.
5. I regard the realist as making two main claims. The claim stressed in this paper is that relativism is false. The realist claims that what is true is true simpliciter and that what exists exits simpliciter. Secondly, he claims that the concept of truth is independent of the concepts of knowability and assertibility; that is, he denies epistemological theories of truth. What is true is not true because it is warrantedly assertible in the ideal limit. The sentence 'There could be truths that we could never come to know' is *not* analytically false.
   On my view, the much discussed debate over the existence of so-called theoretical objects does not have to be a debate between realists and anti-realists. It could be a civil war between different realist factions over what happens to exist. Whether those who deny the existence of theoretical object are anti-realists or not depends on whether they choose to epistemologize truth.

Beginning the discussion of Matheson's paper, MacIntosh declared that there were two absurdities to be avoided by a theory of meaning: first, that what endows an expression with meaning is meaning; second, that we cannot tell what does the endowing. But a theory could avoid both these absurdities and still make it no surprise that realists and antirealists could agree on what counts as a good translation. Matheson maintained, in reply, that realists have always relied on sameness of meaning, paradigmatically, on the sameness of propositional meaning. He pointed out later that many have followed Scheffler in holding that we should argue for coreference, not for invariance in meaning.

Tomkow objected that there was not, as the paper presumed, even a *prima facie* commitment to realism in referentialism. Referentialism aims at turning the problem over to science; that is not *ipso facto* realism. Matheson replied that many metaphysicians insist on referentialism. Kitcher thinks that having a common ontology is a sufficient condition for realism. In a later exchange with Tomkow, Matheson contended that it's at the metatheoretical level that realism and antirealism separate.

Churchland commented that coreferentiality was not the only way to reach genuine epistemological issues. According to core realism, one theory may be preferred on the ground of coherence. Incommensurability can be limited in various ways that give us means of choosing between theories.

The conferees would have been hard put to give, at this point, impromptu, a consolidated account of what the papers and discussions had amounted to. As time went on, they became more and more sympathetic to demands for consolidation, or at least for a new general orientation. Okruhlik was assigned, from the beginning, the role of meeting such demands. Blackburn, continually, as the other papers succeeded one another, played the part in the discussions of trying to establish continuities, bringing up questions that pointed to a track through the profusion of contending ideas. In his own presentation, at the end of the conference, he gave a systematic account of his track. We give the account, not in synopsis, but complete.

Simon Blackburn
Oxford University

The Options In Debating Realism

I'll begin by saying the single most important thing I'm going to say which is if anyone picked up a red and white striped towel on the beach, it's mine.

I'll begin properly by telling a story which was told to me by the son of a friend who's doing a project on the history of Halifax. This struck me as a story with a moral for philosophers. Apparently the first winter when Halifax had just been founded there was a kind of stockade and the governor brought with him a gardener to help cultivate things, and he sent this gardener out on a foray to collect some local herbs and things, and the area around this stockade was teeming with Micmac Indians who'd been egged-on by French missionaries to slaughter the English, and when the Micmacs saw this gardener they scalped him, so he didn't get back. The governor had also brought with him a savant, a sort of secretary or, let's say, a philosopher, who didn't know about this episode and who very shortly afterwards decided he wanted to find out what the country was like. Now, he could have been an idealist, or one of these Edinburgh people, and just sat down and made up a maximally coherent story about what it was like, but he didn't, he went and looked. The Micmacs, being inductivists, saw another one coming and they perceived him as bounty and set on him. The philosopher, however, was wearing a wig; he did the only thing he could think of doing which was to throw his wig at them. Now, to the Indians, a self-scalping philosopher was an anomaly; they had never seen anyone do this, so they turned and fled in horror and the philosopher survived. Well, there is a number of morals to this story: one is, always wear a wig; but the one I like best is, I suppose a sort of Thurberian moral which is that philosophers often seem to scalp themselves but they always get up and walk away.

Well, that's the story; I'm not sure what it's got to do with the talk except that I think we've seen a lot of positions that appeared to scalp themselves and then get up and walk away. But I've been moved to try to generate some maps of this area, and see whether we can really define a debate. Some of this material is put in my book, *Spreading the Word* out shortly, moderately priced, be sure to order now. I think my principal problem, throughout this week and before, has been really gripping solidly the methodology of any debate there should be between the realist and the anti-realist. When one hears it proposed that the realist must say this or the anti-realist must say that, I find the motivations for these proposals quite obscure. It seems to me much harder than some speakers and some debates seem to have presupposed. It seems to me much harder than they presupposed actually to define the debate properly and make sure we've got a methodology, to make sure we've got a stable thing which the realist can say and the anti-realist can't or vice versa.

Well, the first map I've got here for us is a map of the options in the area. I don't think there's anything very controversial here, but let me just take you down the way I've drawn it up. I think a lot of the steam in the realist / anti-realist debate comes because we know what the debate is like or we think we do in local areas where we're not talking about the whole of our theory of the world or anything as grand as that; but we're talking something local. Well, morals is a very obvious example. Probably most people in this room would think of themselves as

DIAGRAM 1
Choice Points in the Realism-Antirealism Debate

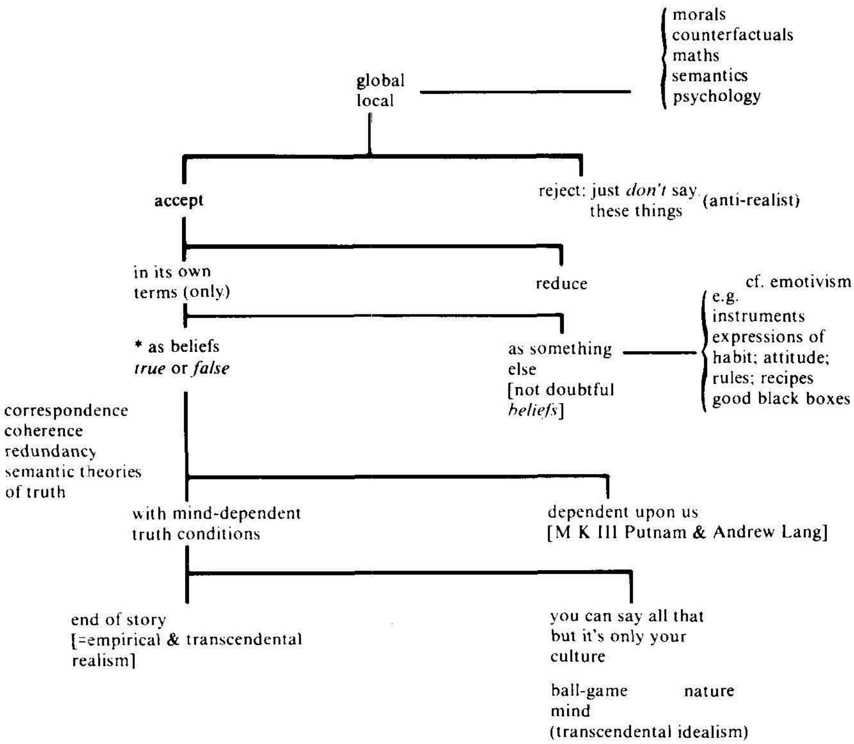

anti-realists of some sort about rights, duties, obligations and so on. Maybe possible worlds is another, counterfactuals and possible worlds, in spite of that famous paragraph in Lewis's book where he thinks that if you believe in ways things might have been that stamps you as a realist. Well, if you believe that and you don't reduce them, that stamps you as a realist. Some of us would like to have the apparatus of possible worlds, perhaps as a heuristic device or something without really being committed to what Lewis would call realism about them. Mathematical entities, perhaps certain of the entities of semantics, perhaps even folk psychology are areas where people would like at least to have the option of what they think of as anti-realism.

Now, I'm going to go through a local anti-realism shortly. I think it's a mistake, it's obviously a mistake as soon as it's pointed out, to suppose that because local anti-realisms and realisms can be debated, it follows that there's a global problem about the whole of our discourse or the whole of our science, because it may be, and in fact it seems to me the appearance is that it is so, that the local options get some of their steam precisely by contrast with other more homely scientific belief. For instance, when somebody denies the reality of possible worlds, he might have precisely in mind a contrast with chairs and tables. So you can't go and say, well, we know how to define realism versus anti-realism in the case of possible worlds so we just take over that methodology and apply it to the discussion about chairs and tables. That would seem to me an unwise course.

O.K. We start off where we've got a global or local question. What's the next option? Well, I suppose the most radical thing—and you could be an anti-realist it seems to me by just rejecting the whole of some discourse. And I might just say, look, possible worlds, yuck; this is a Quineian attitude, you just don't talk about them, you don't use that term, it's a bad theory, it involves errors of various kinds, you just don't do that. You reject it. Rejection doesn't seem to bring many problems, although there are things which shade into it. One of the problems I'm going to try to say something about is the extent to which the debate in science is defined by one side thinking they can have more confidence in theories than the other. Somehow the realist really believes whereas the anti-realist somehow pays lip service and that comes close to rejecting. (We saw something of that during Paul's talk yesterday, the exent to which the anti-realist defines himself by just saying, look, I don't really believe it.) But I'm actually going to discuss real belief versus some other kind of acceptance later, farther down. So I'll pass on through rejection, noticing it as an option. It's not really a very attractive option in the case of science. I mean, what are we doing sitting here saying I reject, you know, the Copernican theory of the solar system. We're just not competent to have an opinion on it, most of us. So it certainly shouldn't be touted around as a very attractive option in many cases.

The next option, which I suppose is familiar, is reduction. We might say, it's O.K. I'll talk about e.g. possible worlds but when I say it, you could paraphrase everything I say in terms of, e.g., maximally consistent sets of sentences or something and you offer an analysis which is supposed to remove the problem of metaphysics or ontology or whatever it is that's worried us about the theory or discourse in question. I think there are a lot of things to say about reduction, as we all know. I'm not going to concentrate on it in this talk partly because it hasn't figured very highly in the agenda; it doesn't nowadays figure as an option when we come to scientific realism. Most people who try to define an opposition to scientific realism these days or assume themselves as trying to do that don't peddle reductive analyses; they do something else. Well, it's the next choice point which I find really the exciting one. Here's somebody who says, "I'm going to accept this theory; I'm going to accept it in its own terms. I don't think you can, for instance, paraphrase thought about possible worlds into talk about something else; or paraphrase moral commitments into other terms. The whole point, you might say, about this kind of discourse is that it's its own animal, it's not going to be reduced down to something else." And then one side says, "Furthermore, when we voice these utterances, when we voice these commitments which we accept, we are expressing belief, judgements, propositions, things that truth conditions, beliefs are going to be true or false or something else if we go in for truth value gaps and the like, but at any rate they're beliefs, they're judgements, they're telling us how the world is." And the other side says "No. When we accept these things, we're not voicing belief, we're not making genuine judgements of truth conditions; we're doing something else." The most familiar option of that kind, I suppose, is what used to be called emotivism in ethics. An emotivist, if he knows his business, doesn't have to give a reductivist analysis, he doesn't say, "When you say you've got a right to free speech what you mean is that you yourself have certain sentiments," or anything like that. That's a naive subjectivism, reductive theory, and that's crazy. What the emotivist who knows his business says is, "You say that sort of thing, you say you've got a right to free speech, don't try to paraphrase it into a sort of non-moral way of moralising, but when you voice that sort of thing you're not expressing a belief with a truth condition, you're expressing an attitude, or announcing a policy, or something. You're evincing ways you'd like the world to go. So, you're expressing something other than judgement." Now, I class instrumentalism in its more plausible form as a version of this choice point; that is, if you say, "I accept the theses of some mature science, I think they're the best we've got, or I really like them, I assert them, I use them in explanations, and so on, but don't construe me as believing something with a truth condition. I'm putting this forward as an instrument of control or an instrument of generating empirically adequate belief."

O.K., so that's the third choice point. I've put "not doubtful" here because I confess that it seems to me a mistake and one that's not infrequently made, I think, to confuse that position with a position which says, "well, they're beliefs but they're really kind of doubtful so when I accept them I'm really very insecure, I don't give them a high probability." That should go up there. A person who says this in my map isn't saying they're doubtful. He's not saying, "Oh, I'm not really sure whether there's a right to free speech", or, "When I say there's a right to free speech I'm frightfully conscious of other ethical theories in which there's not one, or when I say that the solar system has a certain shape I'm frightfully conscious of how doubtful science is and I have problems of induction which make me feel all jittery about it." He's not saying that. He's saying, "When I say it, with whatever confidence I've got in it, I'm not expressing belief, I'm doing something else, and that's a quite different thing from the issue of confidence."

O.K., well, each of these, of course, is the anti-realist side, the realist is down here. A friend saw me draw this map, he suggested that I ought to have done it the other way around because realism is the sort of more right wing position; I don't know why it has that air.

Well, now, there's a problem of methodology, I think a very severe and interesting one about this choice point. The problem comes because it seems to presuppose or acquire a theory of what it is for an utterance to voice a belief with a truth condition, and there are various different things to say about truth, we all know. In particular, I think a beast that hasn't had a very good outing this week, and perhaps deserves a better one, is the redundancy theory of truth. I think it's not as easy as some people assume to give substantive theories of truth. I mean, we all say, "Perhaps truth is acceptability in the ideal limit theory, or something," as though we know what we're saying when we say that. One of the things I want to get onto later on is the difficulty in saying that kind of thing and the plausibility of a redundancy theory of truth, and just to give a brief forewarning of what's going to come; it seems to me quite obvious that if you held a redundancy theory of truth this choice point becomes very much harder to debate and perhaps collapses, because after all, if I hold a redundancy theory and then every time I say, e.g., if I've got a right to free speech, I can just add without any extra import of theoretical commitment: "Furthermore it's true that I've got a right to free speech because I've just said the same." It doesn't import an extra set of concepts. And if I'm going around saying that it's true, then it sounds as though I've dipped over to this side of the fence, or, in other words, as though the choice point wasn't a real one and that's something we're going to need to investigate.

Well, I've listed four words that you hear in connection with truth; there are probably others, correspondence, coherence, redundancy, semantic; and it will obviously be far too big a task to discuss the impact of all of those. But one thing I do want to try to stress and that is that I don't think you get into these choice points just because you're wedded to something called the correspondence theory of truth. I think this is a mistake that Hilary Putnam is responsible for, and to some extent Richard Rorty. People think that serious epistemology, a kind of thing that leads you to interest in anti-realisms is really just the other side of the coin from a correspondence theory of truth, and I argue that that's not so.

Well, if we're plugging down this side (of Diagram 1), we haven't been tempted or lured into any of the anti-realist options so far. Is there anything left to do? Well, we could say look, I accept these things, I don't reduce them, I think the acceptance is genuine belief of genuine truth condition, but it's mind-dependent. It's us. It's we who create these things, and that's obviously opposed by the realist. I want to argue that that's a bad choice point, that you shouldn't go around saying that sort of thing, or, at any rate, you shouldn't say it in general. I think that as Jim showed us, there may be a good scientific reason sometimes for saying that our minds influence things. Perhaps there exist cases of the mind influencing the world. Some truths, it may be, in this sense [are] dependent upon our minds, but that can't be something that an anti-realist or realist has to dispute about. The anti-realist and realist aren't disputing about the truth of particular causal judgements, particular scientific theories. They're trying to dispute about the status of any such belief. Now, the question of whether there's a global option of saying that everything's mind-dependent is, I think, very moot, very doubtful.

I draw then one more which I want to distinguish from that choice point. That is empirical realism coupled with transcendental idealism, Kant's patch. At this point you go around saying look, it's just crazy to say that whether there's a projector in this room depends upon my mind. Obviously it's just a false causal theory about dependencies in the world. Putnam says that whether cats are conscious depends on what we think about it. O.K. Do you all remember that? That just seems to me to be a false causal theory of the world. We know enough about the way the world works to know that whether other animals are conscious does not vary with what we think. It may vary with whether we shoot them, or kick them, or give them various injections, but it certainly doesn't vary with what we think. At least, we've got not the slightest reason to believe it, and anybody saying it does is actually putting forward a false or completely absurd causal picture of the world, the way the world works. So that choice point (empirical idealism) is inferior to this one (transcendental) which tries to protect what I've just said. It says yes, that's right, you mustn't go around thinking your mind does more than it does. However, that's empirical realism and you could have that together with a sort of background God's eye position; most people say, "Of course we say all these things, we can't dip out anywhere along here, we've got our theories, we've got realism winning all along the line, *but*, in the end, this is all filtered through our concepts, minds, categories, etc." That is the final choice point, where Kant resides.

I'm very doubtful about whether there it is genuine. People fall for transcendental idealism far too quickly. It seems to me the interesting choice is the third one, actually.

O.K., well, that's the first map I wanted to draw for you. Now, I now want to make one or two remarks about the methodology of this choice point in particular cases. A particular case I worked on recently is morals, although it may not be intellectually the most interesting case. I think the more interesting case which Crispin is beginning to work on is that of necessary truth. A lot of philosophers in Britain anyhow are now unhappy with Quine's blanket merging of necessary and high grade empirical truth and believe that necessary truths do have a curious and strange status, and one of the options is not that we've got a necessity-detecting faculty but when we describe a truth as necessary we somehow project an expression of an attitude towards it or a ahabit of reliance on it or something about the way it enters into our intellectual life. We project this by describing it as necessary. So there are interesting local cases here and I think I'll spend a quarter of an hour talking about one of them. Because this gives us some sense of what the anti-realist says here.

DIAGRAM 2
One's Mind and the World

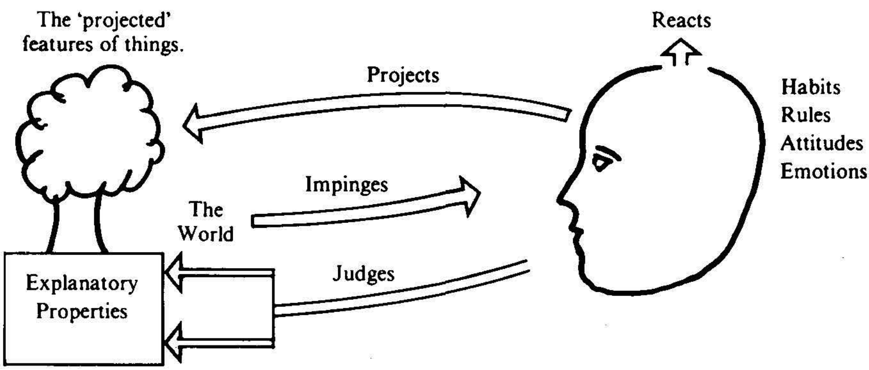

Now in this picture, we've got the world which is imprinting upon the subject, and perhaps there's a class of things the subject can do which just judges the way the world is. He describes the world as possessing properties he's observed, like colours or shapes of things, or whatever, or perhaps he invents explanatory properties, or theories, because at this point we're not presupposing that this anti-realist option can work right across the board. But he can do other things: he reacts, he forms habits, rules, attitudes, whatever it is, which he then spreads, projects on the world. And by that I just mean that he describes the world as though, as if, there were facts which in some sense answer to those habits, emotions, rules, etc. The classic theory of this kind is, of course, Hume on cause. Hume didn't give a reductive analysis of causation. It wasn't a reductive analysis and the identity of meaning wasn't a concept that he worked with. What he tried to do was explain why we describe things as causing each other, what the genesis of that belief is, the genesis of that concept, and the theories of the world which consists as far as we know, as far as anything we can see goes, just of regularities in the succession of events, that it impinges upon us, our minds then have a certain propensity, that is, they form habits which we can do nothing about, we spread those habits when we describe the world as containing causal connections between events. Or, similarly, perhaps the world impinges upon us, it causes us pleasures and pains to take the easier case, we react to that by forming policies, emotions, sentiments, attitudes. We project those on the world when we describe the world as containing duties or obligations or values or humor. So, that's the projective mechanism.

Now, idealism is the second one. The spread world just comes down and leaves the world which is really doing the causing, it's entirely noumenal, it's something you can't know anything about. There is a real world but it's completely blotted out by the conceptual categories that this mechanism causes us to fall into.

Well, now, I've been very interested in the resources of this picture. And to this end I've invented a character who's a kind of aide-de-camp for the projectivist. You don't have to be a projectivist to like this character. He's different; he's the sort of methodological side of it. And I call this character a quasi-realist, or, again, as some of my colleagues call him, a queasy-realist. The quasi- or queasy-realist takes a different stand in certain debates which tend to go on about this picture. Let me just remind you about what tends to happen in first-year moral philosophy courses. People introduce emotivism and then they say, "Of course, on an emotive theory there's no such thing as truth in ethics. You're simply saying boo or hurray to various things." And then of course you notice immediately that the actual shape that our ethical views take doesn't seem to square with that at all, because we say things like "I may be wrong". We allow expressions of fallibility. After all, maybe even a better ethical theory than mine would be wrong! Perhaps even the best of all possible ethical theories that I could form would still exhibit various defects. So you start talking as if (we *do* talk as if) our moral judgements were open to correction and refinement and improvement just as much as our judgements about the natural world. And then there are two possible reactions: one is John Mackie's which sticks with the basically subective picture, this picture, and says, "Well, that just shows that we're in the grip of a false metaphysics, there's an error imbedded in the way we use our moral vocabulary", so he adopts an error theory. That's kind of heroic because it's not obvious how we'd do any better if we gave up these practices which he doesn't like and he never says why we should. The other line is to say well, no, the practice is perfectly in order and refutes emotivism. That's quite a common line in England at present where people are driven back to some kind of weird realism about ethics because they think that the projective picture just isn't adequate to these kinds of things we say about fallibility, about improvement, about truth, about moral truth. Well, my quasi-realist, the aide-de-camp to the projectivist, says, "It's O.K., you can have this picture *and* you can have your ordinary practice, we can explain why we talk and think as if, e.g., there was one truth in ethics, or as if our opinions might be false or as if moral utterances voiced judgements. So we can explain why even in the case of ethics there's a whole bunch of appearances which lead peole to think that, you see." That's the intention of the quasi-realist. So he's heavily involved in the methodology of that debate. He's trying to say, what would you expect if this were the right picture? Well, let me give a simple example of the kind of thing which, it seems to me, he can explain. Peter Geach charged that a projective theory or emotive theory of ethics which regarded a moral utterance as voicing something other than the proposition or judgement couldn't explain the fact that moral utterances go into the antecedents of conditionals with exactly the same meaning they have when they're outside them. There's a context in which they're not asserted. Take his example: "If telling lies is wrong, getting your little brother to tell lies is wrong." Most of us would say, "Yes, that's right." I think we're committed to that. It doesn't matter whether we believe it or not; what does matter is that in that conditional you've got the utterance sentence, "Telling lies is wrong" coming into the antecedent of the conditional. And whatever else you're doing when you voice that you're not committing yourself to an attitude to telling lies. It's unasserted. You don't say telling lies is wrong, you say if it is then something else. Notice that you couldn't say, "If boo are lying, then boo to something else." It's syntactically ill-formed. And the charge would be that the projectivist just can't explain why in this little way the moral utterance looks as if it voices a judgement. And my quasi-realist response is, "Yes we can explain it." I won't go through the whole story but roughly you can divide it into two parts. First of all, can we see what they're up to when they utter those conditionals? Yes, we can. The conditional actually voices another moral commitment although what you might think of as a second order one; it's one about the structure of attitudes that we are endorsing. When I utter that conditional, I'm endorsing only moral sensibilities, which, if they reject telling lies, if they're against telling lies, are also against getting little brothers to tell lies. I set myself, I announce myself, in other words, or express the attitude of being against a potential structure in morality which would say you mustn't tell lies but it's quite all right to get other people to do so. You can imagine a morality which has that as a thesis, e.g., some sort of very heavy personal kind of Bernard Williams style morality, but that's not mine, I'm against it, and I express myself as being against it by voicing that conditional. So, you can explain what the conditional is doing, right? That's what we're up to when we voice it. The next problem will be can you explain why we've got the conditional form to do it, and that, of course, would depend upon a background theory of conditionals, again that's a longer story that I'm not going to go into. My answer would be that you can.

I'll give one more example because that's more relevant to the scientific case. Suppose I say, "I may be wrong. I believe that telling lies is wrong but I may be wrong." Now, that sounds really bad, that sounds as though I've got moral truth. What can the projectivist say, because there's no reality there which his attitude is to match up to. That's the whole point. It's not a belief which corresponds to anything at all. So how can he understand us when we go around puzzling about whether we're right? We go in for moral self-reflection, and self-doubt. How can the projectivist make any sense of that at all? Same two problems: Can he explain what we're up to? Can he explain why we've got that vocabulary to do it? Well, explain what we're up to, yes. The image here is Neurath's boat. Any morality has various theses which interact with each other; some give us standards whereby others can seem doubtful. We get the idea of flaws and failures to which systems are prone. So of course I can doubt whether your attitudes are any good. Is it more surprising that I can worry whether *mine* are? Not a bit more, because the standards for acceptance or rejection may lie inside my own theory, my own moral set of attitudes. To take a simple example, maybe I'm too quick to get angry with people, and so I'm severe with people in various ways. I can worry whether that's so. I know that it's potentially a flaw and therefore I can worry whether it's a flaw which I have on exhibit, I can worry whether it's a flaw which might infect my own moral system and that kind of worry I can express by saying maybe I'm wrong. It's all an internal thing; the trick is that you turn what seemed like an external reflection on the adequacy of the whole coherent system into a worry which can be quite well understood just when you remember that the system will contain elements which can be used to criticise others. O.K., so that explains what we do when we talk about fallibility and the like. Does it explain why we can do it by saying, "I may be wrong"?

Well, this brings us to the point of all this, which you may have lost, rather, which is whether the quasi-realist can construct a notion of truth because if he can construct a notion of truth on behalf of morality then he ought to be able to do it with science, you see. Well, now I think the root worry here is divergence and convergence. You might say look, "You've given us the ingredients from which we can picture things going like this. You start off with the null morality and then life goes on and in comes various inputs and you start forming attitudes." O.K., now you might say, "Well let's just simply define truth as the ideal set of attitudes, membership of the ideal set of moral attitudes. And so when I'm worrying whether it's really true that I've got a right to free speech, I'm worrying whether that goes into the ideal set of attitudes." And here the trouble of course is that I've got absolutely no right to any such notion as we've seen. And indeed in morality you might easily start thinking in terms of a divergent tree structure; there'd be various points where improvement could be got by going either of two ways. And so the thing just collapses outwards, and you might say, "Well, whatever else you should be talking about you shouldn't talk about single moral truth." I think that's a very worrying picture, and it's the one that's obviously in people's minds even in the case of science.

I don't worry about this as much as other people. The reason is I think it's not so easy as people suppose to believe in these nodes—these choice points where improvement could equally be got by going in either of two directions. The classic discussion of this is Hume. Hume discusses it not in the Inquiry or the Treatise but in a very great paper called "Of the Standard of Taste," and the example he takes is this. He says, "How can we say that one, one aesthetic judgement is true, another one wrong?" And he considers the literary dispute between two critics, one of whom is a young man of warm and amorous passions who likes reading Ovid. O.K., since it's the twentieth century, let's say *Playboy*. His literary system has that as a paradigm, good thing to be reading. And the other one is a man of calm, mature judgement who likes reading Tacitus, or *The Economist*. Now, imagine each of these refining their literary systems and their set of judgements as much as the ingredients permit. They roll along growing more and more ideal, but this fundamental divergence of constitutional sensibility or taste is going lead to them to different and apparently conflicting judgements. One's going to say, "*Playboy* beats *The Economist*," and the other's going to say, "*The Economist* beats *Playboy*." How on earth can we, looking at that, believe that truth attaches to just one side? Because, as Hume himself says, the difference arose from such a difference of internal constitution as was blameless on both sides. It's not that one of them's exhibiting a defect from the word "Go"; and I think that perhaps when Paul talks about different kinds of epistemic engines, I think what we've got in mind is a difference which is blameless on both sides. Neither of them exhibits a defect because then it would be easy, you see, to barge a notion of truth through. The fact that some one is defective and isn't going to agree with you, he's going to get a different system, that doesn't matter, so long as you've got a title to regarding him as defective, you're still in business. You've got a unique truth. But here we've got differences leading to divergent judgements, blameless on both sides; how can you recover a notion of truth?

I think that's almost pure image, actually. I don't think there is a problem there. Why not? Well, you could ask me, "What do I think about *Playboy* versus *The Economist*?" And the important point is for this case to work they've got to be equally O.K. It's actually not got to be my judgement that one is much better than the other, because if it is my judgement that one is much better than the other, then the case is one where one of these was defective, I'll say, and I won't draw this picture. If that's my judgement, then both these systems are flawed because they've claimed a strict inequality when in fact there's an equality. Now, that sounds like a trick, doesn't it? It sounds as though I have said, "There's going to be one truth about this issue" when, as it were, I've stacked it up, so you can see with the ingredients I've provided myself there can't be just one truth about it. So, it sounds like a trick—one bound the philosopher was free and you all think I've scalped myself.

It seems to me very important, though, that whenever you think of one of these choice points, these alleged different ways in which the systems can roll off in different directions, you've got to think about it from your own standpoint. You've got to make your own judgement about the case, and I think people suppose it's much much easier than it is to just go and say well, you could imagine in an equally good science as ours, an equally good theory which would tell us that the solar sytem has nothing like the shape it has. A difference which is blameless on both sides. Western Europe goes thataway and describes the solar system with a familiar pattern. Another culture with an equally good science goes thisaway and makes the solar system something else. And people get this image; they think, "Yes, that's obviously possible. It could be equally good." And I want to say,"Why do you think you can say that? You know as well as I do the shape of the solar system. Any constuction, any theory, or any epistemic nature which led people to say it was anything other than heliocentric, Mercury and Venus are nearer the sun than we are, and so on, would have to be wrong. I don't care what other virtues these people have. There's one thing they're not good at and that's detecting shapes of solar systems." Similarly in this case: I set up a *prima facie* case in which people of different sensibilities make different judgements and the difference is blameless on both sides, then I've only really got three options: I can say that *Playboy* is as good as *The Economist* in which case each of them's wrong; I can agree with one of them; I can agree with the other of them. Perhaps I might say that in some sense comparisons shouldn't enter really, you shouldn't be in the business of making this kind of judgement at all. This is what people say when they reject terms of discussion. Is Beethoven a better composer than Brahms, you know, just don't talk like that, O.K.? And of course I can say that but in that case a literary system which is in the business of making comparisons is defective. You get the same result. You haven't got a candidate for a node or choice point.

I think in some respect the most important thing I want to say is that the image of divergence needs filling out if you're to use it, divergence which is blameless on both sides, as an argument against allowing yourself a single notion of truth. Quite what notion of truth we are allowed remains to be seen, but at any rate that's not an argument against allowing ourselves one.

If I had time here, I'd give you a little proof why travel broadens the mind. I'll just do that very quickly, it's quite nice actually. You often think that things can go this way, O is obligatory, right, P is permitted. You often think that, you know, moralities as they evolve can go in the direction of obliging certain things or permitting you not to do them. It's obligatory, you know, to be chaste before marriage, it's permissible not to be, you see. Blameless cultures, relativism, they go different ways. Why does travel broaden the mind? Because if you've got a *prima facie* case of this you think you're at a node, you think you've got it permissible to obligate P and it's permissible to permit not P. So it looks, this is the *prima facie* case, you know, you come to Canada and they say it's permissible not to be chaste before marriage; you go to Britain and they say oh no, it's obligatory. And here's a traveller who comes across a *prima facie* case where that's permissible and that's obligatory. Now, he's got to make up his own mind. He's got to decide whether it's permissible to be chaste before marriage. And the interesting thing is that once he's got a *prima facie* case of this and that, he only needs the modal system S_4, the reduction principle to get from PP to P. It's permissible to permit something; you only need S_4 to get down to it's permitted, whereas from permissible to oblige something you need the stronger modal logic S_5 to get down. It's not nearly as plausible to say that the fact it's *prima facie* permissible to oblige something yields the judgement is obligatory. And that's why travel broadens the mind.

Where was I? Now, you're still going to say, "Look, Blackburn, it's absurd to think that there's, that you've justified our believing that there's a single, say, ideal moral system, and the argument may show that we shouldn't glibly assume that there are choice points all over the place. We shouldn't just announce that it's obvious that there are not, but surely this isn't anything like a sufficient justification of saying what there is." And I'd agree, if we're talking ideal limits. I don't think there is one. But here the plot thickens; it may be thick enough already, but it's thickening. Because what I want to ask is whether it's the kind of thing I've been saying is sufficient to justify us in practising and thinking as if there were an instrumental view of the ideal limit, you see. And I would argue that it is; that is, on the basis I've given you, you can see why it's intellectually reputable, permissible, to tackle, say, a case of a moral dilemma as if truth was going to attach to just one side, just one. So you go around saying things like, "There's going to be an answer here. I don't know what it is but one of these sides is going to be better than the other, and maybe I'm wrong when I come to an answer." In other words, to have the habits that we associate with the use of the concept of truth, I think it's possible if you marinade and digest what I've been giving you for you to see that that will become the intellectually reputable thing to do. And that would justify us in the use of a concept of moral truth which has—I know Peter was cross about this yesterday—but it has what I dare call a regulatory or regulative status. We use it to regulate our process of forming opinion. We behave as if there's a unique truth and in that sense you use the concept.

Well, that's what the quasi-realist does on behalf of people. He goes around saying you might expect us to practice differently if we had this genesis of our beliefs and you might think that the way we practice stamps us as realists but in fact it doesn't. You can earn all the concepts that we use from the anti-real position. Well, he does look then like fouling up the choice point. Because it's going to make it very difficult indeed to say what would be different about the something else. The quasi-realist, in other words, is in the business of making that choice point much less hard to debate than people think.

Let's get back to science. I want to go through very quickly some of the ways that people have been trying to draw the realist anti-realist distinction, and see which of them, if any, look plausible. I just drew this up last night and I may have missed some, but at least this much has been going on this week and there's probably more. Some people think that a real realist must believe in an observation-theory distinction. He needs it because he wants observation to dig piles into truth. I don't think that's true. Some people think they're real anti-realists, although they believe in the observation theory of distinction, as van Fraassen needs to. It seems to me you can believe or disbelieve in that distinction and not thereby stamp yourself as a realist or an anti-realist. I shan't argue this now; I was going to argue that but I've just not got time. Some people say the real realist believes in the convergence of scientific theories and uses his real realism to explain it (Putnam, Boyd). I agree with van Fraassen about that. I don't think the real realist alone has an explanation of that. Furthermore, I don't see why he should be more or less optimistic about the progress of science than anyone else. So I don't think that works either. The use of logical indicators such as bivalence is tricky. I said that I'd given you materials from which to construct a notion of moral truth. I think it's a notion which would permit bivalence as well. I would advocate, in other words, using that as a logical principle to govern moralising. As I say, I don't really believe in an ideal limit, but I do believe in doing our moralising properly, and I'd say that using bivalence is something we can purchase on quasi-realist techniques. I don't think bivalence is going to do very much work here. Sorry, in deference to Crispin and Peter, I ought to say I don't think this is generally so. I mean in mathematics you get quite a different construction I think.

The fifth one we had that we ran last night, so I won't say very much more about it, that's the Loewenheim-Skolem result. It seems to me the problem with that argument and Putnam's in the scientific case was properly aired last night. It't just not true that the interpretation of other people proceeds simply by seeing them as accepting a set of sentences which you've then got the methodological right to interpret as you wish in any domain that you please to do. So, you can interpret them in any domains of given cardinalities. The mistake is thinking of other people as coming to us simply with a portmanteau of sets of sentences. They don't. In any case, I think it's very noticeable in Putnam that although he points the result of that proof—the proof that from knowing the truth value that some of these sentences would have, not just in this world but in any possible world, you can't get a determinate interpretation—at this character the metaphysical realist, the problem had nothing to do with metaphysical realism. It just arose in a world in which there are people and words and things, and everybody believes in that. It's not a philosophical doctrine. So, the anti-realist and the real realist are not separated by that. We've had the anti-realist thinking that whether bats are conscious depends on us, and we're going to come back to mind-dependence to finish off, so I'll leave that. A final suggestion is that the real realist allows the global skeptical position and the anti-realist doesn't; that is, the real realist alone can say things like, "Maybe even the best idealisations of our current scientific theories are in doubt, improved as far as we can, maybe even all that would be false. Brains in vats." The anti-realist, who's supposed to be seeing truth as some kind of idealisation of ways of forming opinion, is supposed not to be able to say that. I just don't think that's so. It seems to me that the anti-realist who knows his business can say that with just as much of a puzzle on his face as any real realist. Again, taking it in the moral case, I think one should be allowed to say, you know "We are terribly corrupt animals, maybe even the best moral position that any of us is going to articulate or would articulate will still not be very good." I think you can allow yourself a skeptical possibility even on an anti-realist basis. I think I'd better leave that to the discussion, but I think you can.

The anti-realist, or the real realist, I think shouldn't take issue on the things that Jim was talking about. I've already said that. But, finally, you're left with what tends to happen when people call themselves realists which is that they flap—they wave their hands in front of them—they say, "I believe in the world out there." Right? And the other people say, "I don't." but what I've been trying to say is the flapping doesn't define the debate; it might define an image or metaphor but if you want to debate you've got to find a methodology for it. I believe in a world where my hand is. So, I think many ways of trying to define what the one lot can say and the other lot can't don't work. However I've just given you a local option; I'm an anti-realist about morals and I've tried to explain why I can go on and practice with a very realist cast to the things I say; that's the quasi-realist construction. So, it seems to arise whether the same thing will do in science. After all, if I'm a projectivist about morals and I can make a realist-sounding moral theory up (and I can say realist-sounding things about morals, in spite of this philosophical anti-realism), is it a possible position as far as science goes? And I confess that my answer is "No." I don't think that you can be an anti-realist about science. And I don't think this for the same reason that Bill was getting on to yesterday.

This is because of the role of explanation. Projectivism is about what explains our practice, in this case, of moralising. And the explanation is not that we react to moral fact; the world is kept thin; it's not moral fact, it's not your perception or awareness or the world in which people have rights which explains why you say that they've got rights. And that's a possible thing to say about moralising, in fact it's the best thing to say about it in my view. I don't think it is a possible thing to say about science; I just don't think there is an option. In other words, I don't think there is a philosophical option of saying, "Aha, the reason why you talk about the shape of the solar system as you do is . . . " and then say mention something other than the shape of the solar system. I just don't think there's an option there. In other words, science itself, or our scientific world view, explains why we hold the views we do, and it explains it by citing the facts we believe in. It's because the solar system has the shape it does that you believe, that, for instance, Venus is closer to the sun than us. There's simply no option of saying anything else, because if you say anything else you're not taking up a philosophical stance, you're just abandoning the science. It's because of the fact that Venus is nearer the sun than us that we believe it is. It's an astronomical discovery. Why did we believe it, then? Well, you might come up with a story; look, these telescopes are all bent and stuff, but that's science, that's not philosophy. You'd be actually contributing to first-order astronomical theory if you'd said that. You wouldn't be standing outside and taking up an armchair *a priori* position about the status of all this. So, I just don't think that it's a serious option of saying any of the projective things when it comes to straightforward science of that kind.

Let me try to sum up what I've done so far; I think I'm just coming up to the hour. It may seem like hours, but it's just one so far. I insist on the debate being defined; I insist on people saying not just, "I announce myself as taking one stand or another of these choice points." I want to say, "Now, what is it to take one stance or another? What makes the difference? How do I tell whether you've got a belief or a truth condition or say an expression of attitude or an instrument?" Now, you could say, "Ah, I'll practice differently." You practice differently. One side would say things, the other side can't, or something. One side would accept bivalence, the other wouldn't. I think one of the great virtues of Dummett's work on this was precisely that he saw the need to say, "Look, don't just flap, don't just say mathematical objects are out there or they're not. Show us how a realist or an anti-realist practices differently," and, of course, he identified it in the acceptance of intuitionism. Now, you'll say well, it's very, very hard to see anything that the projectivist or in this case the instrumentalist need do differently. We can see how he can earn the practices which you'd think were characteristic of realism. So, he'll end up saying the things that the realist thought were his private property. So, how can one go on to be a projectivist in the case of morals if there's no different practice and if I've insisted on some difference? You can do it in the local cases by giving a difference of explanation, and that's the final thing, that's the only thing you've got here that can mark you down as on that side. You can say look, I can be on this side say locally about morals or counterfactuals or whatever it might be because I don't believe that the way I talk about rights or duties or possible worlds or whatever it is is explained by exposure to that. It comes to an explanatory thing, the projectivist offers that picture and the other person doesn't. But the explanatory thing won't generalise; it may be all right to explain why we talk about possible worlds or rights or duties or whatever it is by saying, "Well, we are reacting to a world that doesn't really contain those things; the explanation is this, you recite sentiments or policies or habits or whatever it might be." But you can't do that to science, because that's rejecting the science, at least you can't without leaving this room and going over to the other building and start in talking about why their telescopes were bent or something. Science has its own involvement in first order epistemology. That's the difference.

I'd like to just finish by saying two things: one is relatively local but it explains something I said earlier: I said that I thought this was a better choice point than these. And one reason I think might be very apparent now: my quasi-realist in the case of morals, how does he fare on the issue of whether morals is mind-dependent? Supposing somebody says, "Ah, look, I'm a real realist about moral values because construction morality is mind-dependent, and according to me it isn't." What does that mean? Well, he says, "Look, you'd think that if we had had different sentiments, it wouldn't have been wrong to kick cats, but N, he continues, "Kicking cats doesn't depend on our sentiments it depends on cats." This charge is just wrong, because the way the projectivist coped with indirect contexts shows that he's got no more reason to say that than anyone else. It's just not true. If we'd had different sentiments, we mightn't have appreciated that it was wrong to kick cats—that is the right thing to say. And the projectivist can say it because he'll give an account of this conditional which goes something like this, "Look, if you said that, you'd be endorsing a certain kind of morality, namely one which you might call the bourgeois morality, the one which can't tell whether things are good or bad until it tells how people are reacting to it. But that is not a sensibility to endorse. The structure to endorse is one which needs no information about what people are like, before delivering hostility towards kicking cats."

To use Kant, Carnap, or Putnam, we might say that this is a defence of 'internal' realism—as part of our theory of the dependencies of things in the world, we simply stand by the normal view, that mostly how things are is quite independent of how we take them to be. Traditionally, this drives idealism back to some large, external or transcendental position, and this gives the last choice point on my diagram. Science is here an expression of our perspectives—our quality space, or filtering devices, or categories, which so 'overlay' the world that—well either you just forget it somehow (lose THE WORLD, as Putnam and Rorty have it), or you lapse into scepticism (tellingly, it is at least as easy to see these modern idealists as sceptics, as not actually succeeding in abandoning any concept of THE WORLD. This is because the arguments they use—Goodman's arguments, usually, really force us into scepticism, not into a genuinely anti-metaphysical positivism).

I cannot refute transcendental idealism in the minus five minutes remaining. I don't think it can be refuted at all. I think its permanent appeal lies more in a kind of image, than in any habitable doctrine about knowledge or realism. The image is a very powerful one, one which arises when we distance ourselves from our own perspectives and normal habits, and try to ask general questions about their adequacy. And it's not that these general questions are improper, or in any way out of place. Idealism, particularly about space and time, becomes very appealing—it becomes a ghost that will permanently haunt anyone with a philosophical or even a religious temperament. So I'll finish with a poem by Andrew Lang about it, and about cricket, which ought just about to complete your education:

> If the wild bowler thinks he bowls
> Or the wild batsman thinks he's bowled
> They know not, poor misguided souls
> They too shall perish, unconsoled
> *I* am the batsman and the bat
> *I* am the bowler and the ball
> The umpire, the pavillion cat,
> The roller, pitch and stumps and all.

In an exchange with Wright, which began the discussion of Blackburn's paper, Blackburn maintained that if he accepted the medical story about the cause of measles, he also had to accept that it's the performance of that virus which is responsible for my belief. "It's true that when I say *that*, you can explain my saying it in terms of my having made certain observations; but then we go through the cycle again, don't we, and I say, 'Yes, but what explains the results which those observations gave was this virus,' and that's just a first-order thing, that's the important point, that's just part of medicine, not philosophy . . . . The crucial point for my purposes, the point I'm relying on, [is] that it is not part of first-order moral theory or first-order mathematics or first-order theory of counterfactuals—it's not part of those theories that you explain our beliefs in those things by citing the numbers or the rights or the duties or the possible worlds or whatever. That's metatheory, that's us, that's philosophers. So there's an option of doing it the other way around. It *is* part of first-order science that the reason why I believe 'There is a sun there' is because there is a sun there, and if you don't accept that then you stop believing."

At several points in the discussion, questions arose about the bearing of what Blackburn said on the relativist position of Putnam. Davies, for example, claimed that Putnam's position was that the category "cat" and the category "consciousness" were categories that we have "because we are creatures that have a certain biological constitution, a sort of interest, etc. Given that we have those categories, however, it's then I take it an empirical matter, a function of these inputs that he says we can't get away from, it's a function of those that we come then to a judgement, given that we're constituted the way that we are . . . . It's only in that sense I think that [Putnam] is saying that whether cats are conscious depends upon the human mind." Blackburn retorted, "If Putnam is just saying that then it's shocking to put it as he put it . . . . It's a very trivial thing to say that it's because of contingencies associated with our biology or our language or our cultures or our semantic spaces or God knows what . . . that we're capable of making a judgement. That's one thing. It's quite a different thing to say that it's because of those contingencies that galaxies exist or cats are conscious." The first thing, according to Blackburn, was a platitude. "But the question is whether you can get something which leads in an anti-realist or . . . relativist direction out of that platitude and you certainly can't do it in one swoop by saying, 'Therefore, the way the world is depends on us'."

Later, however, in answer to a question from someone else, Blackburn conceded that Putnam's argument would go through if the issues about realism and relativism simply turned on the posibility of giving any set of sentences multiple interpretations. However, Blackburn held, there was more to the issue than this. Both in the common sense world and in science people offered not just sets of sentences, but sets of sentences in behavioral contexts, and these contexts enabled us to arrive at determinate interpretations.

The last paper of all at the conference, Kathleen Okruhlik's, presented, by means of a diagram, a taxonomy of the positions that *philosphers take on the chief issues discussed under the heading of "scientific realism."*

Kathleen Okruhlik
The University of Western Ontario

A Taxonomy of Scientific Realism

My assignment this afternoon (as last speaker on the programme) is to pull together some of the strands of discussion which have figured prominently during the last few days and to provide a sort of overview or summing up of the issues that have exercised us. I'll start with three broad questions that still seem to require answers, three questions to which we have been promising the seminar students answers for a long time. About five o'clock this morning, I hit on a diagram that seemed helpful in answering all three questions; so I'll structure the talk around development of that diagram.

The first of the three unanswered questions concerns how the different varieties of scientific realism relate to one another and to the different varieties of anti-realism. Can we develop a workable taxonomy which illustrates similarities and differences in a helpful way? Though we may lump together van Fraassen, Feyerabend, the instrumentalists, and the conventionalists as anti-realists, we all know that there are important differences among their views. The question is whether we can exhibit the sources of agreement and disagreement within a systematic schema. Second, and this is a question which has already received considerable attention today: Does the answer to the question of *metaphysical* realism or anti-realism have any bearing whatsoever on the *scientific* realism debate? Third, why have there been so many conflicting interpretations of Putnam's recent work at this conference? If we look at Paul Churchland's interpretation and Dave Davies' interpretation and the view we got from Peter Clark last night, we find that there is virtually no overlap among them. A reasonable explanation is that there are certain tensions present in Putnam's work and that each of the conflicting strands of thought suggests a different interpretation. I want to ask whether these tensions can be resolved without fundamentally altering Putnam's position or whether something will have to give.

To begin, I'll try to develop a taxonomy of the scientific realism debate and then use that taxonomy to cash out its relationship to metaphysical realism or anti-realism. You'll recognise this description of scientific realism as being somewhat similar to Jim Brown's earlier account. We have a semantic component and an epistemic component. The semantic component says that theoretical statements in mature science are approximately true independently of us and the theoretical terms in mature science typically refer. The epistemic component says that we can have *good reasons* for believing that certain theories are true and that their terms refer, where the interpretation of truth and reference are left open. Thus, four possible positions in our taxonomy are determined by the possibility of affirming or denying each of the two components of full-blown scientific realism. Questions? Please feel free to interrupt at any point.

THE TAXONOMY

	Semantic	**Non Semantic**
Epistemic	I Sellars Putnam$_1$ Boyd Glymour Horwich	III (early) Positivists Newton-Smith Dummett (math) Quine Davidson Kant Wright/Clark Other verificationists, old and new
Non Epistemic	II van Fraassen Popper Laudan (?)	IV Classic Instrumentalists Duhem, Poincaré & other conventionalists Kuhn$_1$ Feyerabend Laudan (?) Goodman Rorty

Now, I'm somewhat hesitant about doing this taxonomy because for every single person that I've put on the map there will be somebody in the audience who thinks that that person is misplaced. In order, however, to make this thing work at all, a certain amount of squeezing and accommodating to the categories must be done. In the first group

I've put those people who affirm both the semantic and epistemic ingredients of scientific realism: Sellars, the early Putnam, Boyd, Glymour, Paul Horwich, just as examples. These are people who believe that there is truth independent of our epistemic activities and who also believe that we can have good reasons for believing that certain of our theoretical statements are true and that the terms mentioned in those theories refer. Now, in the second group we find those people who affirm the semantic ingredient but deny the epistemic ingredient of scientific realism. Here we have van Fraassen who thinks that our theoretical statements do in fact have truth values independently of us but who is pessimistic about our ever being able to know which of our theories are true and which are false. Popper at least in his classic period (and probably throughout) would fall into the second group. I'm not sure where to place Laudan because he's been cautiously and systematically agnostic about the semantic thesis, so I've put him in both the second and fourth quadrants. What's absolutely certain about Laudan is that he denies the epistemic component of scientific realism and he's been very careful not to take a position on the semantic component, so I'll put him in both groups with a question mark. I think what he has done so far is probably compatible with both. The third group will include those who affirm the epistemic component but deny the semantic component of scientific realism. Certainly the early positivists and Bill Newton-Smith must go here. Newton-Smith would have preferred to be in the first group, but we saw in earlier discussion that to accommodate cases of underdetermination (such as his example of periodic but open time versus closed time) he had to eventually go a verificationist route; in order then to preserve the epistemic ingredient he had to sacrifice the semantic ingredient. So, he ended up (rather against his will) being in this third group of philosophers. By analogy, although he doesn't have a position on scientific realism per se, Dummett and other constructivists in math would belong here as would Quine, Davidson, Kant, Wright and Clark at this conference. Any sort of verificationist will fall into the third quadrant.

In the fourth group, finally, we get those philosophers of science who can't be called realists on either criterion because they deny both the semantic and epistemic components of scientific realism. Here we get the classical instrumentalists, the whole tradition of saving the phenomena in astronomy, Poincaré and the other conventionalists, Kuhn (the way he was widely interpreted in the first edition of the *Structure of Scientific Revolution*), Feyerabend in *Against Method*, possibly Laudan, again depending on what position he takes on the semantic thesis, and Goodman and Rorty. Now, if we combine the four quadrants we get a map which eventually will help us to answer all the questions that we asked at the beginning of the session. I take it that Paul Churchland would probably like to be in the first group but he's really right off the map with his more recent views.

Paul Churchland: That's a charming interpretation. I accept it.
Okruhlik: I did say "map".
Churchland: You didn't say "off the wall"?
Okruhlik: No.

O.K., so in the first quadrant we get scientific realism in its sort of robust, old-fashioned sense, both components affirmed. In the fourth quadrant we get something that probably can't be called scientific realism on anybody's terms. In the second and third quadrants we have mixed cases because we have one component of traditional scientific realism affirmed and the other denied. We see that the later Putnam isn't on the map anywhere yet, and I hope eventually to say what's going on here.

Now, if we accept this as a reasonable taxonomy of the scientific realism debate (recognising that some people have been squeezed a little to fit the categories), then what can we say about the relationship of this debate to the metaphysical realism debate? Recall that in the metaphysical realism debate the issues are generally couched around the possibility that even in the ideal epistemic limit we could be wrong. Anti-realists deny that possibility, realists affirm it; and what it seems to come to is that metaphysical realists have a notion of truth which is radically, irreducibly non-epistemic, whereas for the anti-realists the question of truth is always going to be epistemic, irreducibly epistemic. So, what are the answers to our initial three questions? In terms of a workable taxonomy, I think the map is a good first approximation, a way of understanding why all these anti-realists are included as anti-realists while delineating the differences among them in terms of which of the components of scientific realism they affirm or deny. The second question was: What's the relationship between scientific realism and the metaphysical debate? When I first started working on this I shared the sentiments of Carl Matheson and Simon Blackburn this morning that the metaphysical question really had very little, probably nothing, to do with the scientific realism debate; and I haven't modified that opinion very much. I think that there is some impact but not much, and the reason that it's limited is exactly the same reason that Peter

Clark pointed out last night, that Simon pointed out this morning and Carl this afternoon: The metaphysical question isn't going to settle any epistemological issue and the epistemological issue is one that has been central to the scientific realism debate. Now, Peter told us last night that the idea of the transcendent truth which we probably can't get to and wouldn't recognize anyway is just useless epistemologically. I think that the Putnam-surrogate of warranted assertibility under ideal epistemic conditions is probably just about as useless as the transcendent notion of truth, partly because on some interpretations it's not going to be very different from the transcendent notion of truth.

In particular, I think the ideal limit of warranted assertibility is not going to help us establish reference for two reasons: first of all, even if we could use it to define reference by saying, "We successfully refer to those entities which will be part of the theory in the ideal epistemic limit", that wouldn't help us at all in knowing now which of our theories successfully refer and which don't. This is because in Putnam's view everything is up for grabs all the time, and we have no way of knowing what things are going to be like in the epistemic ideal limit. In this sense I think warranted assertibility is just as useless as the transcendent notion simply because it's just about as transcendent. Second, it's not clear to me that we can in fact define reference in terms of the theory which we would hold at the ideal limit since that theory itself could be Skolemized and we'd have the same problems with reference. Now, what Peter suggested quite reasonably last night was that Putnam could avoid the Skolemization problem by going a constructivist route, and in the part of the paper that he didn't actually read aloud to us last night he suggested that in this way we could in fact be realists about scientific entities, about theoretical terms. He also suggested that we might be realists in thinking that whatever our best scientific theories now affirm to exist really do exist, and we would do this on analogy with intuitionists in mathematics who are realists about the domain of constructable proofs.

So, warranted assertibility, in Peter Clark's sense, would be the analogue of proof in intuitionistic mathematics and we would have constructivist programmes in both domains; we could be realists in so far as a construction could be carried out. Now, what he noted quite properly in the written version of the paper is that there is an important disanalogy between proof in intuitionist mathematics and warranted assertibility in science, *viz.*, that proof in intuitionistic mathematics is a once-and-for-all affair, there's no flip-flopping whereas we've seen people try unsuccessfully to come up with an analogue for scientific theories which would yield the sort of stability of reference that we obtain in constructivist mathematics programmes.

Peter didn't really argue for the possibility, the workability, of such a programme except to say that the only alternative was a truth conditional analysis which was pretty hopeless and there's nobody who has been able to carry it out so far. Now, I'm not quite sure why he's any more optimistic about the verificationist programme for scientific theoretical entities. It seems to me that just as many people have worked on such a programme and it's been just about as hopeless (although I know that Peter is attaching a lot of confidence to Crispin Wright's endeavours in that area). I think, though, that Peter is certainly right that this is one way Putnam wants to go. We might want to say this: that in any event the anti-realist is no worse off than the metaphysical realist, since again the scientific realism debate (as we have presented it here) has this very strong epistemological component; and it doesn't appear that either side in the metaphysical debate can get any mileage, any epistemological mileage, out of its position.

Whether this is exactly true or not I think will depend upon whether the semantic component of scientific realism is separable from metaphysical realism. In an article called "Three Forms of Realism", Paul Horwich argues that they are separable in one direction, that you can deny metaphysical realism while still maintaining semantic realism. He argues in that article that the reason that Putnam and Dummett have gone the route they have is because they share with the metaphysical realists the mistaken belief that you need more than a redundancy theory of truth, that there is some sort of surplus meaning to truth. Horwich suggests that we could drive a wedge between semantical realism and metaphysical realism by adopting a redundancy theory of truth and a use theory of meaning. Thus, he himself claims to be a semantic realist and a metaphysical anti-realist. So my answer to the relationship between the metaphysical debate and the scientific debate will turn partly on whether Horwich is right or not. If he's right, then you can give up metaphysical realism while maintaining semantic realism, and if we refer back to our diagram then any of the four quadrants would be compatible with metaphysical anti-realism. So, there wouldn't be any constraint in this direction from the metaphysical debate.

If, on the other hand, Horwich is wrong, if metaphysical realism just is semantic realism, then it's obvious that a metaphysical anti-realist will have to occupy either the third or the fourth quadrant, in our taxonomy. So, we could either go the constructrivist route or we could put ourselves in the fourth quadrant with those people who deny not only the semantic component but also the epistemic component.

Now, what if this debate were resolved to everybody's satisfaction and metaphysical realism turned out to be the correct view; what would be the implications here? I don't see any way you could fail to be a semantic realist if you are a metaphysical realist, so I think that consideration would mean that you have to stay on the left hand side of the diagram and occupy either the first or the second quadrant. In either case, you're going to be a semantic realist just because that would reduce to some form of metaphysical realism. The only question open would be whether you were epistemologically optimistic or pessimistic. If you were optimistic, then you could be a full-fledged scientific realist of the early Putnam-Boyd type. If you were pessimistic, if you were skeptical about our ability to know whether our theories actually did latch onto the truth, then you'd be down in the second quadrant with van Fraassen. So, just to sum up the answer to the second question, I think that there are some ramifications for scientific realism from the metaphysical debate. But none of these affects the epistemic component, so resolving the metaphysical question is still going to leave the scientific question open no matter what route you go. The epistemic question's going to be open no matter what. Whether the semantic component is fixed or not will depend upon whether we're forced to identify semantic realism with metaphysical realism. Yes Paul?

> Paul Churchland: I think it leaves it open a bit more than that, because I want to be a metaphysical realist, for example, but I don't want to be put in either One or Two. It's quite possible that there's a single world out there being the way it is quite independently of us, but that reference and extension and truth are simply not the relations we bear to that world and so one could deny the standard semantic thesis about how our cognition relates to the world and still be a metaphysical realist. But this is the only case where I think the taxonomy breaks down and it could be repaired quite quickly. Otherwise it's succeeding very nicely.
>
> Okruhlik: Well, there still would always be the implication in one direction that if you were a metaphysical realist the semantic component in the scientific debate would get hit just because those two would reduce to one another. If it turns out that metaphysical realism is the answer, then we're going to have to locate ourselves in the first or second quadrant—barring your exception for just a minute. I have to talk to you in the question period about the sense in which you could still have the metaphysical component. I thought maybe we could get into that a little later?

So, that's my answer to the second question. No matter what we do with the metaphysical issue, the epistemic ingredient's going to be completely unaffected; so the scientific debate will still be up for grabs. It's possible that the semantic component will be fixed by the resolution of the metaphysical debate.

Third question: why do we get so many divergent interpretations of Putnam at this conference and everywhere that Putnam's recent work is discussed? And assuming that the reasonable answer is that those tensions and perhaps contradictory strands are present in Putnam himself, why? How do we explain it? How do we locate Putnam on this map? I think the basic problem is that Putnam wants to locate himself in both the third and fourth quadrants and that this can't happen, that it's just logically impossible to hold both positions simultaneously. When he's doing his Dummett line, when he's taking the sort of constructivist view of *Models and Reality*—that version of Putnam that we got from Peter Clark last night—he's trying to locate himself in the third quadrant. When he's doing his Goodmanian thing, talking about pluralism and about the indispensability of frameworks and so on, he's trying to locate himself in the fourth quadrant. And since one of these entails the acceptance of the epistemic ingredient and the other the denial of the epistemic ingredient, it doesn't seem that he can have it both ways. If Putnam is located in the fourth quadrant then it's not clear that he's entitled to any notion of the ideal limit of warranted assertibility because there is no guarantee that even the modest kind of convergence that he needs is available to him. There's certainly no argument in the book that would make it available to him. Probably a notion of warranted assertibility makes perfect good sense in the third quadrant. It's just an analogy with constructivist mathematics, it's what would be in the ideal theory; that is, if all the proofs were in, then that's what would be true and the terms of those theories would refer. But even in a case like that where we can define an ideal limit without involving ourselves in any contradictions it's not clear again that it can give us any epistemic guidance at all. And it's not clear that Putnam could in fact maintain the sort of pluralism, the anti-scientism, the anti-reductionism which he's so keen on in *Reason, Truth and History*, if he really goes consistently to the third quadrant. (This is in addition to the problem I alluded to earlier about whether you can make a constructivist programme work in science because there does seem to be this rather blatant disanalogy with the area where it has been successful, in intuitionistic mathematics.)

So, I think that if we say that Putnam is trying to straddle this divide, trying sometimes to locate himself in the third quadrant and sometimes in the fourth quadrant, this explains a lot of the tensions that we've seen in the debate. Paul Churchland always refers to Putnam as a relativist because he concentrates on passages that put him down in the fourth quadrant. Everything that Peter Clark said last night seemed to come out of the earlier models of reality theory where he's trying to be in the third quadrant univocally. Dave Davies was, I think, trying to juggle both sides and get a coherent story out. I'm not quite as optimistic as Dave (even when Dave expressed reservations) about the workability of this programme. I suppose I'm even less optimistic that Putnam can have it both ways. And I think that probably explains those instances in the book where he seems to be working with a stronger theory of reference than he's entitled to, reference across paradigm shifts, through stability of concept even when conceptions vary. It explains that sort of tension that arises because he seems to want something analogous to Kant's transcendental nature, some ideal of human flourishing that will give us the constraints that will keep us in the third quadrant, and at the same time he wants to be a pluralist. It would also explain the sort of silly move where all of a sudden he drags Davidson in to solve the problem of radical incommensurability and non-translatability after he has made very much of his earlier case depend on the presence of just such elements. So, that's my answer to the third question: Why have there been so many conflicting interpretations at this conference? Why does Putnam seem to involve himself in the different position that he does in the book? I think it's because he's basically torn in two directions. He likes the work that Goodman has done recently, he wants to be that kind of pluralist; and on the other hand he's trying to carry out the constructivist programme, he's trying to be the new Kant but without the transcendental nature. I don't think he can have it. And *that's* the whole truth.

In the discussion of Okruhlik's paper, the main division lay between those who thought the taxonomy useful and illuminating, and those who were dissatisfied with it in one aspect or another, or (in some cases) disinclined to believe that any taxonomy was needed. Blackburn wondered whether when notions like "mind-dependence" were properly sorted out there wouldn't be a "mass exodus" to the top left-hand corner, all parties finding themselves subscribing after all to both semantic realism and epistemic realism. MacIntosh held that the most interesting issues were lost from the taxonomy; for example, what was interesting was what is the theory of truth and what is the correct theory of meaning?

Others defended the taxonomy. Okruhlik herself, replying to MacIntosh's doubts about the need for a taxonomy, said, "One thing . . . that is very important about doing a taxonomy like this is . . . that it is very hard at the beginning even to read the literature because in a scientific realism debate the epistemic and the metaphysical go together whereas in the metaphysical debate they're apart, that's the very definition of realism . . . . It took me a long time to realize that very often people who are calling each other names and seem to be addressing the same issue were in fact talking at cross purposes." Brown interjected an example: "We might have thought that van Fraassen and Popper were poles apart. With the taxonomy, we can see that . . . one calls himself a realist and the other calls himself an anti-realist, and they are in fact holding exactly the same position."

The most vigorous defense of the taxonomy came from Tomkow. It was, according to him, a clear thesis that there are propositions that are true or false even though we have no way and will never have and could never have any way of finding out if they are true or false. "That's a clear thesis and whether or not it's true it's occupied an awful lot of philosophy." Other questions that led to talk of realism, "Do our theories have integrity" "Are mature theories, whatever they are, true in the ideal limit, whatever that is?" can generate a lot of debate, but are these different issues or are they non-issues? . . . "I take a strong line. There aren't any issues outside of here . . . . If you can't locate yourself on the picture, maybe you ought to wonder, not about the taxonomy, but whether or not you've got a philosophic topic to wrestle with."

MacIntosh remained unsatisfied. "We come away asking exactly the same questions that we began with." So many philosophical debates do end; but not every participant felt that way about this one.

(This report was compiled and edited at the Dalhousie Department of Philosophy by David Braybrooke and Thomas Vinci, with the indispensable assistance of Margaret Odell.)