HOW TO BE AN ANTI-REALIST*

Alvin Plantinga
Calvin College

The dispute between realism and anti-realism has been centerstage in Western philosophy ever since Kant's alleged Copernican Revolution. It is center stage now--both in Anglo-American philosophy, and on the continent, where anti-realism runs rife among the hermeneutical epigoni of Heidegger. It is therefore not surprising that the dispute has become a popular topic for APA presidential addresses. Lining up on the anti-realist side we have Hilary Putnam, with his 1976 Eastern Division Address entitled "Realism and Reason" and Richard Rorty, who delivered "Pragmatism, Relativism and Irrationalism" to the Eastern Division in 1979. On the realist side we have my illustrious immediate predecessor but two in this august chair, William P. Alston, whose 1979 presidential address was entitled "Yes, Virginia, there is a Real World". You will notice here a certain imbalance; so far it's been anti-realism two to one. You will also notice a certain flavor of interdivisional dissension: as you can see, it's been the Western Division--that sturdy and stalwart bastion of such traditional values as home, family and realism--against the more effete and epicene Eastern Division, with its old world tendency towards cynicism and world weariness. Of course the Pacific division is yet to be heard from; but it would be rash indeed to predict the behavior of anything containing southern California.

Now I hope to mediate the dispute. I shall argue that anti-realism in its presently popular forms is wholly unacceptable; unbridled realism, however is also unlovely; and I shall suggest what I take to be the right way to be an anti-realist.

I. Creative Anti-realism Characterized

First, I must say just which dispute it is I mean to mediate. Anti-realism comes in a bewildering variety of forms. I said the realism-anti-

---

realism dispute goes back at least to Kant; but of course there is an important use of the term 'realism' going back much further to the medieval dispute about universals. As this example suggests, one speaks of realism or anti-realism with respect to a given area or subject matter: universals, say, or the past, or other minds, or sets, or micro-entities in physics. And in one use of these terms, the realist is just a person who argues that there really are such things as universals, or other minds, or propositions. In this way of using the term, a realist with respect to inferred entities in science thinks there really are such things as the elementary particles–atoms, electrons, quarks and the like--endorsed by contemporary physics; he adds that they have pretty much the properties contemporary science says they have. And of course an anti-realist with respect to inferred entities denies these things. Call this sort of anti-realist an 'existential anti-realist'.

But there is another brand of anti-realism, one that is substantially a modern, post-Kantian phenomenon. The Kantian anti-realist doesn't deny the existence of an alleged range of objects; he holds instead that objects of the sort in question are not ontologically independent of persons and their ways of thinking and behaving. Kant didn't deny, of course, that there are such things as horses, houses, planets and stars; nor did he deny that these things are material objects. Instead his characteristic claim is that their existence and fundamental structure have been conferred upon them by the conceptual activity of persons. According to Kant, the whole phenomenal world receives its fundamental structure from the constituting activities of mind. Such structures as those of space and time, object and property, truth and falsehood--these are not to be found in the world as such, but are constituted by our own noetic activity. Were there no persons engaging in noetic activities, there would be nothing in space and time, nothing displaying object property structure, nothing that was true or false. We might think it impossible that the things we know--trees and mountains and animals--exist but fail to be in space-time and fail to display object-property structure; indeed, we may think it impossible that there be a thing of *any* sort that doesn't have properties. If so, then Kant's view implies that there would be nothing at all if it weren't for the creative structuring activity of persons. Of course I don't say Kant clearly *drew* this conclusion; indeed he may have obscurely drawn the opposite conclusion; that is part of his charm. But the conclusion in question does seem to follow. The fundamental thrust of Kant's Copernican Revolution is that the things in the world owe their fundamental structure and perhaps their very existence to the noetic activity of our minds. Or perhaps I should say not minds but *mind*–whether there is just one

transcendental ego or several is, of course, a vexed question for Kantian exegesis. The fact is, this question is considerably *more* than vexed; given Kant's view that quantity, number, is a human category imposed on the world, there is presumably no cardinal number n̲, finite or infinite, such that the answer to the question "how many of those transcendental egos are there?" is n̲.

Now we might call this Kantian kind of anti-realism "subjective anti-realism"; but perhaps that could be considered unduly pejorative. Instead, let's call it 'creative anti-realism.' So there are at least two kinds of anti-realism: creative and existential. Each, furthermore, can be restricted to a certain domain, or taken globally (although global existential anti-realism--the view that nothing whatever exists--has never been popular). With respect to a given domain, one can be either a creative anti-realist or an existential anti-realist, but not both. With respect to mathematical objects, for example, one can hold either that there really aren't any such things as numbers (although such sentences as 'There is a prime number greater than 17' are both true and useful) or that there are such things, but they owe their existence and their character, somehow, to our noetic activity. One might be an existential anti-realist with respect to unobservable entities such as quarks, but an existential realist with respect to the ordinary middlesized object of everyday life. Or one might be an existential realist with respect to the former and a creative anti-realist with respect to the latter. In some areas one brand of anti-realism may seem considerably less plausible than the other. In theology, for example, existential anti-realism--i.e., atheism--is fairly common; creative anti-realism, on the other hand--the view that there is such a person as God, all right, but he owes his existence to our noetic activity--seems at best a bit strained. Oddly enough, however, certain contemporary theologians do seem to adopt just such a position; Gordon Kaufman and John Hick, for example, apparently hold that the word 'God' denotes what they call an "imaginitive construct" or a "mental construction", so that while indeed there *is* such a person as God, he is, in an astonishing reversal of roles, something *we* have brought into existence.

The kind of anti-realism I mean to consider is creative anti-realism. Of course creative anti-realism with respect to *some* things is very attractive; your average house, or automobile, or B-1 bomber, for example, really does owe its existence and character to the noetic activities of persons. But the creative anti-realist is not ordinarily content to restrict his creative anti-realism to such things as houses, automobiles, and B-1 bombers. And it may not be initially obvious that the anti-realist I mean

to discuss are properly thought of as *creative* anti-realists. For their characteristic claim is not that human beings create or structure the world; instead they make a certain claim about *truth*. They claim that truth is *provability*, or *verifiability*, or perhaps *warranted assertability*. According to Rorty, for example, "the only sense in which we are constrained to truth is that, as Peirce suggested, we can make no sense of the notion that the view which can survive all objections might be false"; and his book, *Philosophy and the Mirror of Nature*, Rorty apparently suggests that truth just *is* "what our peers will let us get away with saying".

Now how is this a case of creative anti-realism? As follows. The core of creative anti-realism is the idea that objects in the world owe their fundamental structure--and, if they couldn't exist without displaying that structure, their existence--to our creative activity. The world as it is in itself, apart from this structuring activity, doesn't display any of these features. The idea is that if there were no persons (or if there were some and they didn't structure the world in the way in which we do in fact structure it) then there would be no objects in space or time, none displaying object property structure, no *number* of things of any sort, and the like. So the Kantian claim is essentially a modal or counterfactual claim: there is a sort of intellectual or conceptual or noetic activity we engage in, such that if we didn't engage in that activity (and no other creatures leapt into the breach) then things would not display the sorts of structure in question.

And now consider Rorty's suggestion that truth is what our peers will let us get away with saying. The idea is not, I take it, that our peers are both so splendidly informed and so fastidious that as a matter of fact they'll let us get away with saying something if and only if that thing is true. It is rather that truth just *is* what our peers will let us get away with saying; it is therefore *necessary* that a proposition is true just in case our peers will let us get away with saying it. Of course there are problems of interpretation here; but the idea seems to be that if our peers had *not* let us get away with saying what in fact they do let us get away with saying, then those things would not have been true. If our peers had let us get away with saying that there have never been any dinosaurs or planets, then it would be true that there have never been any dinosaurs or planets. But of course if it were true that there have been no dinosaurs, then there would have been no dinosaurs. Whether or not there were dinosaurs, therefore, depends upon the noetic activity of our peers--that is, upon us, since we are they.

Putnam's view is similar, if less than wholly clear. One thing he quite clearly holds is that an ideal scientific theory--one which appro-

priately measures up to our standards for excellence in a scientific theory–
couldn't possibly turn out to be false. "The most important consequence
of metaphysical realism," he says,

> is that *truth* is supposed to be *radically non-epistemic*–
> we might be "brains in a vat" and so the theory that is
> 'ideal' from the point of view of operational utility,
> inner beauty and elegance, "plausibility", "simplicity",
> "conservatism", etc., *might* be false. "Verified" in any
> operational sense does not imply "true" on the meta-
> physical realist picture, even in the ideal limit.[1]

Putnam, by contrast, holds that it is *not* possible that such a theory be
false; he holds that 'verified' in the operational sense, whatever exactly
that is, *does* imply 'true'.

How shall we understand this? It is of course possible that Putnam
is here displaying a powerful faith in our epistemic powers and proce-
dures; these are so good that we simply couldn't go wrong in the long
run. But the fact is that's probably *not* how he's to be understood;
what we have instead is a thesis about truth. Truth, he thinks, just *is*
verifiability, or assertibility–verifiability or assertibility according to
the standards we do in fact adopt. Now there are obvious problems
here: *whose* standards for verifiability or assertibility are we referring
to? Those of, e.g., a thirteenth century Frisian milkmaid? Or a twen-
tieth century Moslem? Some people think the way to determine the
age of the earth is to consult the Bible; for others, that is the sheerest
foolishness. There are deep disagreements as to what constitutes veri-
fication; must we say, joining what we had always thought of as a pecu-
liarly benighted sort of sophomore, that what is true for you might not
be true for me? And what about creatures much like us but much more
intelligent–i.e., much more intelligent by *our* standards? Even if there
aren't any such creatures, modesty demands that we concede there
could be some; shouldn't our account of truth respect the epistemic
standards and requirements such creatures have, if there are any, or would
have, if there were some?

These questions receive a partial answer in Putnam's Piercian refer-
ence to the 'ideal limit' of scientific inquiry. He thus seems to be sug-
gesting something like the following account or analysis of truth:

(1) p is true if and only if p ideally meets our epistemic requirements,

or perhaps a bit less vaguely

(2) p is true if and only if if there were an ideally rational inquirer
   in epistemically ideal conditions, she would accept p,

or, perhaps, in order to eliminate what may be seen as the improper
individualism of the last statement,

(3) p is true if and only if if there were an Ideally Rational Scienti-
   fic Community that had all the relevant evidence, it would accept
   p.

Now, how, exactly, does creative anti-realism enter in? As follows.
Suppose we ask under what conditions an assertion or statement is veri-
fied; what determines that, e.g., the assertion "there were dinosaurs
roaming the earth long before there were human beings" is verified?
Putnam doesn't give an explicit answer; he does say, however, that verifi-
cation is a process that goes on *within* a theory, within a way of think-
ing or speaking. And what he means, I think, is something like the fol-
lowing. In adopting a language we adopt a complicated interlocking
set of practices, including rules and procedures for verifying such asser-
tions. In coming to understand an assertion such as "there are three
cows in that meadow" we learn, among other things, under what condi-
tions that statement is assertible, i.e., verified. And what determines
that this statement *is* verified, under certain conditions C, is just the
fact that we have adopted practices and procedures according to which
it is thus assertible under those conditions. What determines then, that
a statement S is assertible or verified under conditions C is our having
adopted certain practices and modes of behavior. So whether the state-
ment *Dinosaurs once roamed the earth* is verifiable, depends upon our
ways of thinking and behavior--upon what practices and rules for veri-
fying that statement we have adopted. To use a word Putnam employs
elsewhere, it depends upon the conventions we adopt. But of course
dinosaurs once roamed the earth entails and is entailed by *it is true
that dinosaurs once roamed the earth*; since according to Putnam the
latter owes its truth to us and our behavior, the same goes for the former.
On Putnam's view, therefore, whether dinosaurs once roamed the earth
depends upon us and our linguistic activities.

We can approach the same point by different routes. According to
Putnam whether the term 'cow' (to move to his favorite bovine example)
denotes, and whether, indeed, it denotes *cows*, depends upon us and
our behavior--the network of practices, and rules and conventions we
adopt. Putnam puts it as follows:

> What I am saying is that, in a certain "contextual" sense, it is an *a priori* truth that 'cow' refers to a determinate class of things. . . . Adopting "cow talk" is adopting a "version", in Nelson Goodman's phrase, from within which it is *a priori* that the word 'cow' refers (and, indeed, that it refers to cows) (p. 495).

The details of this suggestion are not easy to make out, but the main thrust is fairly clear. For those who adopt cow-talk, it is an *a priori* truth that the word 'cow' has a denotation, and an *a priori* truth that it denotes *cows*. But then it is an *a priori* truth, for those who adopt cow talk, that there exist such things as cows; for most of us will find it but the work of a moment to infer *there exist such things as cows* from *'cow' refers to cows*. And what *makes* this an *a priori* truth (for those who adopt cow talk) is just the fact that they *do* adopt cow talk. If there are those who do not adopt cow talk, then for them this is not an *a priori* truth; and if we had not adopted cow talk then for us it would not have been an *a priori* truth. So whether or not there are cows, depends upon us--upon the categories, rules and strategies we adopt for verification, upon the linguistic practices and procedures we employ. Or rather perhaps we should say, whether or not there are cows for us, cows *für uns* as opposed to cows *an sich*, depends upon us and our linguistic proclivities, thus accommodating the puzzling suggestion that *there are cows* could be an *a priori* truth for some people–those who adopt cow talk–but not for others. But in any event, if it is a truth for us that there are cows, then it is indeed a truth that there are cows; on the suggestion in question, therefore, whether or not there are cows depends upon us and our noetic activity. The reasoning involved, furthermore, seems not to be restricted to the bovine, but to be perfectly general. The astronomical moral to be drawn, then, is that whether or not there are stars and black holes will also depend upon us and our noetic activity; and the theological moral is that whether or not there is such a person as God depends in the same way upon how we think and talk. Putnam's suggestion is less than wholly clear; as far as I can make it out, however, he seems to endorse a creative anti-realism of the global variety. In fact he himself goes on to draw the Kantian moral:

> Let me close with a last philosophical metaphor. Kant's image was of knowledge as a "representation"–a kind of play. The author is me. But the author also appears as a character in the play (like a Pirandello play). The "real" author is the "transcendental me".

I would modify Kant's image in two ways. The authors
(in the plural--my image of knowledge is social) don't
write just *one* story: they write many versions. And the
authors *in* the stories are the *real* authors. This would
be "crazy" if these stories were *fictions*. A fictitious
character can't also be a real author. But these are true
stories (p. 496).

II. Argument For Anti-Realism

Now there is something initially off-putting, offensive and disturb-
ing about creative anti-realism, as even its partisans admit. It is not,
initially, at all plausible to suppose that whether dinosaurs once roamed
the earth or whether there are black holes depends upon us and what
we say and think; and creative theological anti-realism seems at best
a piece of laughable bravado. So the next natural question is: what is
to be said in favor of creative anti-realism? What is there about it that
might constrain us to accept it? What leads its protagonists to adopt
it? Of course a bewildering variety of arguments and quasi-arguments
have been offered for creative anti-realism. In order to narrow the field,
I shall confine myself to arguments offered in APA presidential addresses.
I want to examine a couple of these arguments, although, in the style
appropriate to an after dinner address, I shall not Chisholm away at them
with the relentless patience and endurance their importance warrants.
Putnam's presidential argument goes as follows. The realist, he says,
holds that even our best theories might be false; even if a theory meets
all the constraints--short of truth--we impose, it could nonetheless be
false. There is no guarantee that the theory arrived at by our best lights
and efforts is true. God could, if he wished, systematically deceive us;
or he could allow us to systematically deceive ourselves. It is logically
possible, for example, that the world has come into existence just 15
minutes ago, complete with all its apparent memories and dusty books
and other alleged traces of the past. This is possible, and compatible
with our having an epistemically ideal theory according to which Colum-
bus discovered America some 490 years ago and Aquinas died some
705 years ago.
So the realist claims that even our best theories could be false; and
Putnam argues for the denial of this claim. Let T be an epistemically
ideal theory; the argument then proceeds as follows:

I assume THE WORLD has (or can be broken into) infinite-
ly many pieces. I also assume T_1 *says* there are infinitely

many things (so in this respect T$_1$ is "objectively right" about THE WORLD). Now T$_1$ is *consistent* (by hypothesis) and has (only) infinite models. So by the completeness theorem (in its model theoretic form), T$_1$ has a model of every infinite cardinality. Pick a model M of the same cardinality as THE WORLD. Map the individuals of M one-to-one into the pieces of THE WORLD, and use the mapping to define the relations of M directly in THE WORLD. The result is a satisfaction relation SAT–a "correspondence" between the terms of L and sets of pieces of THE WORLD--such that the theory T$_1$ comes out *true*--true of THE WORLD--provided we just interpret 'true' as TRUE(SAT). So what becomes of the claim that even the *ideal* theory T$_1$ might *really* be false?

Well, it might be claimed that SAT is not the *intended* correspondence between L and THE WORLD. What does 'intended' come to here?

T$_1$ has the property of meeting all *operational* constraints
. . . . But the interpretation of "reference" as SAT also certainly meets all theoretical constraints on reference–
it makes the ideal theory, T$_1$, come out *true*.

So what *further* constraints on reference are there that could single out some other interpretation as (uniquely) "intended", and SAT as an "unintended" interpretation (in the model-theoretic sense of "interpretation")? The supposition that even an "ideal" theory (from a pragmatic point of view) might really be false appears to collapse into *unintelligibility* (pp. 485-486).

Now how shall we understand the argument implied here? Suppose we concede that THE WORLD has a cardinality, and suppose we pass over the question just how we are to think of its alleged pieces. Let T$_1$ be our ideal theory; T$_1$ says Putnam, meets three conditions. First, it correctly predicts all true observation statements. Secondly, it meets whatever "operational" constraints there are. While Putnam doesn't here explain what operational constraints might be, what he says elsewhere suggests the following. Let OP be a set of statements giving, for each measurable magnitude M̲–gravitational potential, for example–and each rational space-time point p̲ the value (to rational approxima-

tion as close as you please) of M at p. Then to say that T_1 meets all operational constraints is to say that it is consistent with OP. And the third condition is this: since T_1 is consistent, it has a model; we can use that model to define a satisfaction relation SAT–an interpretation function that assigns extensions from the set of pieces of THE WORLD to the terms of T_1 in such a way that its sentences come out true. In other words, T_1 has a model in THE WORLD. Putnam then asks the following question: "so what becomes of the claim that even the *ideal* theory T_1 might *really* be false?"

But the suggestion embodied in the question isn't at all compelling. True, T_1 meets these conditions; but so do a bewildering bevy of theories inconsistent with it. In fact any theory that contains all true observation statements--statement of the form 'S is appeared to F-ly at t'--and is consistent with OP has a model in THE WORLD, thus meeting these three conditions. Pythagoreanism, for example, a theory that contains all true observation statements together with the affirmation that there exists nothing but numbers, meets these conditions. So does T_2, a theory that contains all true observation statements and states that God created the world just 15 minutes ago, complete with all its apparent memories and dusty books and other traces of the past. The same goes for T_3, a theory that contains all true observation statements together with the assertion that there aren't any human beings. But surely we have here no reason at all for concluding that these three theories are true.

Now we can put essentially the same point as follows. This interpretation function SAT assigns extensions to the terms of T_1 in such a way that its statements come out true. Of course there are *many* interpretation functions that meet this condition. Suppose T_1 contains the claim that, say, some cows weigh more than a ton. There will be an interpretation like SAT in which 'cow' is assigned as its extension a set S of prime numbers; if the bovine population of the world is, say, half a billion, S might contain the first half billion primes, the other assignments being such that "weighs more than a ton" has as *its* extension a set that overlaps S. But again, why should the fact that there is a model, in this sense, of T_1 incline us to think that this claim T_1 makes about cows is true? It hardly seems so much as relevant. Of course there is *one* interpretation that is, we might say, privileged: the *intended* interpretation A* that assigns to each term of T_1 the extension that term actually has. Thus A* assigns to 'cow' the set of cows rather than a set of numbers; and 'weighs more than a ton', under A*, will have in its extension not numbers, but just those things that weigh more than a ton. If under *this* interpretation all the sentences of T_1 came out

true, then perhaps the question whether T_1 was really true would indeed collapse into unintelligibility. But the mere fact that there is *some assignment or other* under which T_1 comes out true has no tendency at all to show that this question thus collapses. T_1 could still be false, we want to say, and would be false if it came out false under A*.

Now here Putnam asks the following question: how *can* we sensibly deny that SAT is the intended interpretation? How can we sensibly deny that the extensions it assigns our terms are the extensions they actually have? "The interpretation of 'reference' in L as SAT", he says, "certainly meets all *operational* constraints on reference. But the interpretation of 'reference' as SAT certainly meets all *theoretical* constraints on reference—it makes the ideal theory, T_1, come out *true*. So what *further* constraints on reference are there that could single out some other interpretation as (uniquely) 'intended' and SAT as an 'unintended' interpretation . . .?" The suggested argument, then, goes as follows: SAT meets three conditions: first its domain is the set of pieces of the world. Second, the claim that SAT assigns to our terms the extensions they do in fact have is consistent with OP. And third, SAT is a model of our ideal theory; that is, under SAT our ideal theory T_1 comes out true. But then we haven't any grounds at all for stigmatising SAT as unintended; and hence "The supposition that even an 'ideal' theory (from a pragmatic point of view) might *really* be false appears to collapse into *unintelligibility*" (486).

Once again, it is initially hard to see this argument as compelling. True enough, SAT does meet these conditions; but then so do any number of interpretations. If the pieces of THE WORLD include the natural numbers, an interpretation according to which all of our terms denote numbers will meet this condition. If among the pieces of THE WORLD there are people, an interpretation that assigns to 'shark' a set of people– shyster lawyers and used car salesmen, perhaps--also meets these conditions. Are we to conclude that we can't sensibly ask whether such interpretations are unintended? It is certainly hard to see why we should think so.

Now Putnam has *another* presidential address. This one is entitled "Models and Reality" and was delivered to the Association for Symbolic Logic.[2] And here Putnam returns to and expands the above argument. He points out first that by the Lowenheim-Skolem Theorem, any first order theory in a countable language that is satisfiable in an infinite domain has a countable model. Hence if your favorite first order theory says that there are uncountably many objects of some sort–real numbers for example–then that theory will have models, in the set theoretic sense, in any countably infinite domain–the positive

integers, for example. And even if we add to the theory in question a formalization of all of science (either our present science or some appropriate idealization of it) and also add OP, the set of statements giving, for each measurable magnitude M and rational space time point $\underline{p}$ the value (to a rational approximation as close as you please) the value of M at $\underline{p}$, the resulting theory will *still* have countable models. Now this, of course, is well-known and widely appreciated, if somewhat startling. It can lead, however, to puzzled questions: if this theory, the theory of real numbers, has countable models--models in which the domain of interpretation is countable--then can it really be true that there are uncountably many real numbers? A common answer is to distinguish intended from unintended models; the intended models won't be countable. And this is the point at which Putnam demurs. All of the models in question, he points out, meet both theoretical and operational constraints; as we have seen, they are models both of all of science and of OP. How then can we make these invidious distinctions among them, stigmatising some as "unintended"? If they all meet both operational and theoretical constraints, what could possibly be meant by claiming that some of them are "unintended"? The conclusion to be drawn, Putnam thinks, is that if we are realists, then either we must hold that we have "non-natural" powers of apprehending the property of being a set or we must concede that the term 'set' doesn't have a determinate extension.

Having softened us up in this set theoretical fashion, Putnam goes on to generalize his point. *Everything*, he says, can be Skolemized. Consider T, the totality of our beliefs, or T_1, an appropriate idealization of it: T_1, if consistent, will have a bewildering variety of models. Since it has models in the positive integers, it has models in which the extention of 'dog' is a set of numbers. If theology is included in T_1, there will be models in which the Pope is an inaccessible cardinal. There will be models in which the extension of 'dog' is a set of cats and that of 'cat' a set of dogs. And of course these models won't be so much as isomorphic; by the Lowenheim-Skolem Theorem any consistent theory with an infinite model has models of every infinite cardinality. And this much is undeniable; T_1 will indeed have models of this bewildering variety. But what philosophical conclusion shall we draw? Why should this fact strike terror or even mild alarm into the soul of the realist? Putnam's claim here isn't entirely explicit; what he seems to hold, however, is that if the realist does not assume that we have what he calls "non-natural" powers of apprehending properties, he will have to concede that we cannot discriminate among these models, dismissing some of them as "unintended" or "non-standard". And this is the point of contact with the previous argument for the conclusion that an ideal theory

couldn't fail to be true. The response to that argument, you recall, was that if T_1 were true under A*, the intended interpretation, the one which assigns to the terms of T_1 the extensions they do in fact have, then indeed T_1 would be true; but the mere fact that there is some interpretation or other under which T_1 comes out true of THE WORLD shows only that T_1 is consistent. It does not so much as suggest that it is, in addition, *true*. Putnam's response, as we now see, is a question: what makes you think there *is* any such thing as "the extension the terms of T_1 do in fact have"? In model M_0 the word 'dog' has a set of cats as its extension; in model M_1 a set of natural numbers; in model M_2 a set of angels; how then can we sensibly speak of its extension *simpliciter*? It seems rather just to have different extensions in different models. How can we sensibly pick out one of these models and one of these interpretations as somehow favored? "In short, one can Skolemize absolutely everything. It seems absolutely impossible to fix a determinate reference (without appeal to non-natural powers) for any term at all" (476). The conclusion is that from a realist view point, either we have non-natural powers of property apprehension, or there is no such thing as the assignment function that assigns to each of our terms the extension that term actually has. But then we must acquiesce in the anti-realist claim.

I find this argument tenuous at best. A few years back, Paul Benacerraf pointed out that various *different* set-theoretical objects can 'play the role' of the natural numbers; we can identify O with the null set, for example, and each succeeding number with the unit set of its predecessor or with the set of all its predecessors, or with still other things. And of course there is no reason to stick thus unimaginatively to sets; any countably infinite set will serve as the natural numbers, even if some of its members are propositions or elephants or planets. Another way to put the fact Benacerraf points to is that the theory of natural numbers has many different models; in some of these models '7' will be assigned a unit set, in others a proposition, and in still others a small elephant. Now one conclusion Benacerraf considered was that the terms of natural number theory--such terms as '1', '2', '3', etc.--lack a determinate extention. A sensible response would be: indeed number theory does have models of these different sorts, but so what? The same goes for, say, biology. We can find models or our total biological theory in which the term 'elephant' is assigned a set of numbers, others in which it is assigned a set of horses, and maybe some in which it is assigned a set of philosophers. Should that cast us into doubt as to whether the term has a determinate extension? Should

59

that suggest, perhaps, that we must give up realism with respect to elephants? And now to return to Putnam; his reply seems to be: yes it should--or rather it should push us either into giving up realism with respect to elephants or else into postulating non-natural powers of grasping properties; and the latter course has little, he thinks, to recommend it. "The Platonist will reply that what this really shows is that we have some mysterious faculty of 'grasping concepts' (or "intuiting mathematical objects") and it is *this* that enables us to fix a model as *the* model . . .; but this appeal to mysterious faculties seems both unhelpful as epistemology and unplausible as science."

But suppose we step back a moment to try to assess the force of this argument. Its premiss is that a first order formalization of our actual or ideal set of beliefs will have models in which our terms get assigned many different sorts of objects. It will have models in which the extension of 'cat' and 'dog' are permutted and models in which the extension of the term 'human being' is a set of numbers; indeed, for any obect $\underline{x}$ you please and any non-empty-term $\underline{T}$, there will be a model of our beliefs in which $\underline{x}$ is in the extension of $\underline{T}$. This is surely true; but what follows? As logicians have taught us, if we formalize a consistent theory, the models of that theory will reflect a certain structure the theory says holds; but of course the models of the theory won't in general preserve the extension of the theory's terms. But why concede that we can't discriminate among these models? Why, for example, can't we point out that there is an assignment function $A*$ that assigns to each 1 place term the extension it has in fact, so that under $A*$ the extension of 'dog' is a set of dogs, not some set of numbers? Exactly what is the problem? It is hard to see here anything that should cause the realist to take alarm or for that matter to resort to non-natural properties. The models of a formalized first order theory don't determine the extension of the terms of the theory, even up to isomorphism. What this shows is that the process of formalization is severely limited in a certain dimension; but how does it cast doubt on the view that our terms have determinate extensions?

The facts Putnam points to would be presently significant, I think, only if we had some reason to think that the terms of our language get their meanings or extensions, somehow, by virtue of the set theoretical models of first order formalizations of the body of our beliefs. If we had reason to think, for example, that our terms get their meanings and hence their extensions by virtue of a vast network of *implicit definitions*, then indeed these facts would be significant. But is there even the slightest reason to think that implicit definition is the process whereby the terms of our language do acquire their meaning? Is there reason to

think they *could* be given meaning in this way? Suppose we wrote down all that we know or believe; suppose we then replaced each referring term by its mirror image; and suppose we then proposed to define these new terms by declaring the entire structure an implicit definition of the terms in question. Is there even the slightest reason to think these terms would thus acquire either the meanings or the extensions of their English counterparts?

What Putnam points out, therefore, is that implicit definition will not suffice to confer a determinate extension upon our terms. But why does he suppose that if this is true, then our terms *have* a determinate extension, in the realist sense, only if we have non-natural powers of grasping forms or properties? The idea seems to be that implicit definition is the *natural* way for our terms to get meaning, so that if that doesn't work, then they have a determinate extension, in the realist sense, only if we have non-natural powers. But why think that? So far as I can make out, the only answer here is a question: "what neural process, after all, could be described as the perception of a mathematical object? Why of one mathematical object rather than another?" (471). Perhaps Putnam would say the same thing about properties: what neural process, after all, could be described as the grasping of a property? Why of *one* property rather than another? And perhaps the idea is that if it is implausible to suppose that some neural process could be described as the grasping of a property, then if we *do* have the power of grasping properties, that power must be, somehow, "non-natural". But this is surely dubious. What neural process, after all, can be described as thinking about Paul Q. Zwiers? More poignantly, what neural process could be described as believing that if realism is true and our terms can't get meaning by way of implicit definition, then we must have non-natural powers of property apprehension?

So perhaps Putnam's argument is best seen as a question directed to the realist: if, as you say, our terms do have determinate extensions, how do they get them? How does it happen that the term 'shark', for example, denotes sharks all over the world, including a lot of sharks no one has ever seen? This is a good question, and one at least some realists would answer in terms of grasping properties. But why should the power of grasping properties be thought of as non-natural? I don't myself have even the slightest objection to non-natural powers; and I'd be delighted if we could show that we could grasp properties only if, say, supernaturalistic theism were true; that would be a much stronger theistic argument than the cosmological, teleological, or even ontological arguments. But it isn't easy to see how such an argument would go.

Putnam's argument, therefore, is really a question addressed to the

realist. And the anti-realist is likely to ask still another question. Consider such set-theoretical statements as the Axiom of Choice and the Continuum Hypothesis. What is interesting about these statements is that they are independent of ordinary set theory–Zermelo-Frankel set theory, for example. For each of these statements, there are models of ZF in which it is true and models in which it is false. And Putnam's underlying question is something like this: how could there just *be* a truth about these matters if our best methods of theory construction and investigation–e.g., ZF set theory–don't enable us to *reach* that truth? How can there be truths independent of our best epistemic efforts? Perhaps Putnam's fundamental difficulty with the realist lies at this very point: how could it be that what is certified, even ideally certified, by our best methods–is nonetheless false? Isn't that in some way just unthinkable? At a fundamental level a cardinal anti-realist intuition is that truth, whatever it is, is something that can be known; if the best efforts of mind can't settle the question whether a proposition is true, then there's no truth there to be known.

Now if there is a strong suggestion of this line of thought in Putnam's addresses, it is even more explicit in Rorty's. Although Rorty disclaims any argument for anti-realism; his piece contains, nonetheless, a subterranean current of argument for that position. Perhaps we can approach the matter as follows. A striking feature of the intellectual situation is persistent *disagreement* about such matters of deep human concern as religion, morality and, for that matter, philosophy. Kant was appalled by the fact that after centuries of effort metaphysics had not yet attained the secure path of science; he therefore proposed that his predecessors had been confused and that what was needed was a Copernican Revolution. Indeed, ever since Descartes, modern thought has witnessed one alleged new beginning after another, each innovator declaring his predecessors utterly misguided. And this sort of disagreement is a source of wonder. It is also a source of philosophy. One sort of response it presently provokes--especially in continental thought-- is anti-realism. If disagreement--about the existence of God, or human freedom, or the nature of substance--persists century after century despite our best efforts, perhaps the conclusion to draw is that there isn't any real question of truth in these areas. And of course this impetus to anti-realism is connected with the one I mentioned above: if we can't come to agreement on these matters despite centuries of effort, then presumably we cannot by our methods grasp the relevant truth– in which case, according to the first impulse, there is no truth to be grasped there. And *this* impulse to anti-realism is suggested though not explicitly endorsed in Rorty's address. The effort to find a method

that will enable us finally to *settle* these disagreements, he says is misguided:

> According to this Platonic myth, the life of reason is not
> the life of Socratic conversation, but an illuminated state
> of consciousness in which one never has to ask if one has
> exhausted the possible descriptions of, or explanations for,
> the situation. One simply arrives at true beliefs by obey-
> ing mechanical procedures.[3]

> The pragmatist tells us that it is useless to hope that ob-
> jects will constrain us to believe the truth about them,
> if only they are approached with an unclouded mental
> eye, or a rigorous method, or a perspicuous language.
> (p. 724).

In his book *Philosophy and the Mirror of Nature* Rorty excoriates the entire program of what he calls "epistemology": the attempt to find or develop methods by which, in the fashion of Leibniz' *Characteristica Universalis*, one can settle or decide all important issues. Rorty argues that in fact there is no such method and that the search for it is utterly misguided; and he concludes that realism is mistaken. The pragmatist, he says,

> wants us to give up the notion that God, or evolution or
> some other underwriter of the present world picture, has
> programmed us as machines for accurate verbal picturing,
> and that philosophy brings self-knowledge by letting us
> read our own program. The only sense in which we are
> constrained to truth is that, as Peirce suggested, we can
> make no sense of the notion that the view which can sur-
> vive all objections might be false.

But notice: Rorty apparently agrees with the epistemologists, or Methodists, as they might better be called, that if, in the disputed area, there *were* such a thing as truth in the realist sense, then there would be a sure method for arriving at it. Since there seems to be no such method, truth must be thought of, as he puts it in *Philosophy and the Mirror of Nature*, as what our peers will let us get away with saying. There is surface disagreement, then, with the Methodist, but deep concurrence: truth is what our methods obtain; if in a given area truth can't be attained by those methods, then in that area there is no truth.

One characteristic anti-realist thought, then, is that in fact there is no methodical way to settle all important disagreements. There is no set of premisses all rational persons are compelled to accept, from which by argument forms or procedures required by rationality, we can infer the truth about God, Freedom and Immortality–or for that matter, the truth about numbers and material objects. Here he is right. He concludes, however, that there is no truth here to be had--or better, that whatever truth there is here is of our own making. And that is a giant step indeed.

III. Objections to Anti-Realism

Such are the impulses to anti-realism. Now the problem with anti-realism, in addition to its intuitive unloveliness, has always been a tendency towards self-referential incoherence. Think, for example, of Kant's alleged Copernican Revolution. Enormously ingenious and deeply attractive, it founders on the fact that if things are the way the scheme says they are, then we cannot so much as think the scheme; for then there is, (or there are–one doesn't know which to say) a world of things in themselves that, unthinkably, aren't in space or time, display no object-property structure, and are neither singular nor plural.

And similar strictures apply to the present versions of anti-realism. Rorty suggests that truth is what our peers will let us get away with saying. But this suggestion immediately and obviously falls prey to self-referential difficulties. For neither his peers nor mine will let either him or me get away with saying any such thing. If it is true, therefore, it isn't true; so if it is true it both is and isn't true, in which case it isn't true. Perhaps it doesn't follow, on Rorty's view, that it is *false*; for our peers, being an irascible lot, might not let us get away with saying either this or its denial. But in any event it isn't true.

This sort of self-referential argument sometimes raises eyebrows; it is thought to be cheap, too easy, a mere dialectical trick, somehow a bit unfair. But surely it isn't. If a view is such that, together with obvious truths, it implies falsehood, then clearly it isn't true. But a view that implies its own untruth does imply falsehood, since it also implies its own truth.

Now I believe that a similar self-referential difficulty afflicts Putnam's development of anti-realism. Time is short however, and I shall turn instead to a related difficulty. Putnam's suggested analysis or definition of truth, as we have seen, is something like

(2) $\underline{p}$ is true if and only if if there were an ideally rational inquirer

in ideal epistemic conditions, she would accept p̱,

or perhaps

(3) p̱ is true if and only if if there were an Ideally Rational Scientific Community (IRS), it would accept p̱.

Both of these lead to grief; at any rate each has consequences most anti-realists won't be pleased to accept. Consider (3) first: if (3) is an *analysis* or *definition* of truth, or anything *like* an analysis or definition of truth, it must be necessary:

(4) Necessarily, p̱ is true if and only if if there were an IRS, it would accept p̱.

Now Let A be *there is an IRS.* If

(5) There is an IRS and it does not accept A
is *possible*, then the same goes for
(6) A is true and there is an IRS and it does not accept A

by the principle that a proposition p is equivalent, in the broadly logical sense, to the proposition that p is true.
(6) however, entails

(7) A is true, and it's false that if there were an IRS, it would accept A;
so if (5) is possible, so is (7). But the possibility of (7) obviously contradicts (4); hence (4) entails that (5) is not possible. Accordingly, (4) entails
(8) Necessarily, if there is an IRS, it accepts A.

By the unimpeachable principle that what is necessary is necessarily necessary, (8) is equivalent to

(9) Necessarily (8),

so that (4) entails (9). But (8) entails

(10) If there were an IRS, it would accept A.

Hence (4) entails both (10) and its necessity:

(11) Necessarily, if there were an IRS, it would accept A.

Of course (4) also entails

(12) Necessarily (if, if there were an IRS it would accept A, then A is true).

But (10) and (12) together entail

(13) There is an IRS;

and (11) and (12) together with Modal Modus Ponens, the principle that if P is necessary and entails Q, then so is Q, entail

(14) Necessarily, there is an IRS,

certainly a dismal conclusion if there ever was one.[4]

So (3) entails the necessary existence of an ideally rational community of scientists who have all the relevant evidence; and (2) entails, in the same way, the necessary existence of an ideally rational inquirer. Now I have no objection to the necessary existence of an ideally rational inquirer (or rather an ideal *knower*, since God does not inquire); but at least some anti-realists may have reservations. And presumably no one will welcome the necessary existence of the IRS.

Secondly, the anti-realist views under consideration encounter a serious problem about *disagreement*. Such views begin by taking fundamental disagreement seriously; ironically enough they end by denying its possibility. According to Putnam, to understand a statement is to know what it is for that statement to be verified: "All the 'paradox' shows is that our understanding of 'The real numbers are nondenumberable' consists in our knowing *what it is for this to be proved*, not in our 'grasp' of a 'model'". The idea is that my understanding of a statement *consists in* my knowing how to verify it. I understand a statement when I know how to verify it, that is, know how to verify it according to our standards of proof and methods of inquiry. But the fact is there isn't any such thing as *our*--our human or even our Western--methods of inquiry or standards of proof. Putnam imagines extra-Terrestrials who reject the Axiom of Choice; but we need not go nearly as far afield. Members of the Creation Research Society hold that the way to determine the age of the earth is to consult the *Bible*; Carl Sagan and his friends disagree. Methodist and classical foundationalists believe that serious philosophical inquiry must appeal to premisses acceptable

to all or nearly all careful and reflective human beings; others demur. Suppose, therefore, that you and I disagree as to the proper method for determining whether it's true that there is such a person as God. Although anti-realists are commendably modest when it comes to specifying just what difference in use constitutes difference in meaning, surely, on their view, difference in methods of verification would. But then you and I don't use such sentences as "there is such a person as God" in the same way; and therefore don't use them with the same meaning. And this means that if I affirm theism and you deny it, we aren't contradicting each other. Similarly, if we disagree about the proper methodology in philosophy, we find ourselves, on the view in question, unable to disagree about philosophical conclusions. This way of looking at the matter has the consequence that people cannot disagree at profound levels; it begins by noting and taking seriously our deep human disagreements and ends by legislating them out of existence.

IV. How To Be An Anti-realist

There is thus much to be said against anti-realism. The arguments for it—at least the arguments we've considered, and I know no better—are frail reeds indeed. The versions considered, furthermore, suffer from self-referential and other sorts of difficulties; and, fundamentally, it just seems incredible that whether there were dinosaurs, for example, depends in any way upon how we think or speak. Is anti-realism, then, a mare's nest of confusion, at best a mere galimatias? If so, why have so many sensible people accepted it?

The answer is that anti-realism is not at all a mere confusion; there is strong intuitive support for it or something like it. This intuition is often dressed up in fancy argumentative clothes of one sort or another, partly because, these days, one feels the need of argument for respectability. But the arguments aren't successful; and what there is of substance here is just this intuition, this impulse within us towards anti-realism. How could there be truths totally independent of minds or persons? Truths are the sort of things persons know; and the idea that there are or could be truths quite beyond the best methods of apprehension seems peculiar and *outre* and somehow outrageous. What would account for such truths? How would they get there? Where would they come from? How could the things that are in fact true or false—propositions, let's say—exist in serene and majestic independence of persons and their means of apprehension? How could there be propositions no one has ever so much as grasped or thought of? It can seem just crazy to suppose that propositions could exist quite independent

of minds or persons or judging beings. That there should just *be* these truths, independent of persons and their noetic activities can, in certain moods and from certain perspectives, seem wildly counterintuitive. How could there be truths, or for that matter, falsehoods, if there weren't any person to think or believe or judge them?

'Platonism' is often used to name the view that among the furniture of the universe are such abstract objects as propositions, possible worlds, numbers, and properties. Your true Platonist, however--Plato, for example--doesn't hold merely that these things exist; she holds that they exist *independently* of everything else. Hence they exist independently of minds and their noetic activity; they aren't in any way dependent upon mind. This is realism run amok; and it is this that the impulse towards anti-realism is an impulse *against*. It is worth noting that Platonism properly so-called has been a rare bird in our philosophical tradition. Plato, as I say, was in at least some moods a Platonist. Bertrand Russell was too, at least for a while, and so was the young Husserl, although he outgrew it. No medieval philosopher was, I think, a Platonist, and neither was any modern philosopher before Frege, if indeed Frege was a Platonist.

So what we really have here is a sort of antinomy. On the one hand there is a deep impulse towards anti-realism; there can't really be truths independent of noetic activity. On the other hand there is the disquieting fact that anti-realism, at least of the sorts we have been considering, seems incoherent and otherwise objectionable. We have here a paradox seeking resolution, a thesis and antithesis seeking synthesis. And what is by my lights the correct synthesis, was suggested long before Hegel. This synthesis was suggested by Augustine, endorsed by most of the theistic tradition, and given succinct statement by Thomas Aquinas:

"Even if there were no human intellects, there could be truths because of their relation to the divine intellect. But if, *per impossible*, there were no intellects at all, but things continued to exist, then there would be no such reality as truth." (*De Veritate* Q. 1, A.6 Respondeo). The thesis, then, is that truth cannot be independent of noetic activity on the part of persons. The antithesis is that it must be independent of *our* noetic activity. And the synthesis is that truth is independent of our intellectual activity but not of God's.

The suggestion I mean to endorse can be put as follows: truth is not independent of mind; it is necessary that for any proposition p, p is true only if it is believed, and if and only if it is believed by God. This is truth *de dicto*; but it is also true, *de re*, that every proposition has essentially the property of being true only if believed, and if and

only if believed by God. In the same way propositions themselves, the things that are true or false, are not independent of mind. It is necessary that a proposition p exists only if it is conceived or thought of or the object of some other propositional attitude, for it is necessary that every proposition is conceived of by God; furthermore, every proposition has essentially the property of being conceived or thought of, for every proposition has essentially the property of being conceived by God.

You will no doubt be happy to learn that I don't propose to develop this view here; life (and time) is short and philosophy is long. But I wish to make one clarificatory remark by way of conclusion. First, we need some more distinctions. Creative anti-realism, I said, is the claim that truth is not independent of mind; and divine creative anti-realism is the view that truth is not independent of God's noetic activity. Now divine creative anti-realism comes in several varieties. In the first place there is the view that truth, the property of being true, *just is* the property of being believed by God, or perhaps that of being the object of some other noetic activity on the part of God. This is the theistic analogue of the claims made by Putnam and Rorty–the claims that truth just is verifiability by our standards or what our peers will let us get away with saying. Second, there is the view that truth and *being believed by God* are distinct but necessarily coextensive properties; but the latter is in some way *prior* to the former. This priority can be spelled out in a variety of ways. It might be said, for example, that God's believing p is what *makes* p true, or alternatively that p is true *because* God believes it, or that God's believing p is an explanation of p's being true. From this point of view the fact, for example, that 7 + 5 = 12 is to be explained, somehow, in terms of God's believing this proposition.

I mean to reject both of these views. Truth is not the very same property as being believed by God, even though the former is necessarily coextensive with the latter. But neither is it the case, in general, at any rate, that God's believing p is prior, in some important sense, to p's being true. God's believing p is not, in general, an explanation of p's being true, or what makes p true, or the reason for p's being true. In particular, truths about the free actions of persons other than God are not true because God believes them; on the contrary, God believes them because they are true. Here what is needed, of course, is an explanation of the sense of 'because'. I say each of *it is true that there are cows* and *God believes that there are cows* entails the other; but then what sense, it may be asked, does it make to affirm that God believes that there are cows because it is true that there are, while denying that it is true that there are cows because God believes it? This is a good question: it

is also, as one says when one doesn't have a very good answer, one we shall have to defer to another time.

It is thus not the case that a proposition is true because God believes it. On the other hand it is the case, I think, that a proposition *exists* because God thinks or conceives it. For propositions, as I see it, are best thought of the thoughts of God. You might think this idea compromises the necessary existence of propositions; but not so. For God is a necessary being who has essentially the property of thinking just the thoughts he does think; these thoughts, then are conceived or thought by God in every possible world and hence exist necessarily. As we know, serious difficulties attend the claim that propositions just are *our* thoughts; these difficulties fall away for the claim that propositions are *God's* thoughts. Accordingly, while God believes a proposition because it is true, a proposition exists because God thinks it.

By way of conclusion then: the fundamental anti-realist intuition--that truth is not independent of mind--is indeed correct. This intuition is best accommodated by the theistic claim that necessarily, propositions have two properties essentially: *being conceived by God* and *being true if and only if believed by God.* So how can we sensibly be anti-realists? Easily enough: by being theists.

NOTES

1. *Proceedings and Addresses of the American Philosophical Association,* August, 1977, p. 485.
2. September 1980, p. 464ff.
3. *Proceedings and Addresses of the American Philosophical Association,* September, 1979, p. 724.
4. I owe the idea for this argument to Ernest Sosa, to whom I express my thanks.