

*What Is Realistic about Putnam's Internal Realism?*¹

David L. Anderson
Illinois State University

In 1976, Hilary Putnam, longtime champion of metaphysical realism, startled the philosophical community by abandoning metaphysical realism and offering his own alternative which he has been elaborating and defending ever since. Putnam makes an interesting claim about his new position: He insists that the appropriate description of the view is "internal realism" (or sometimes "pragmatic realism") and that it is a view that a person motivated by "the realistic spirit" might justifiably hold. Very few contemporary realists are sympathetic to Putnam's suggestion that his view is "realistic." Admittedly, Putnam's position does boast a rich ontology. Electrons exist every bit as much as chairs and tables do, and electrons can even help to *explain* the superficial properties of macro-objects. Few realists, however, are willing to count this as a sufficient condition for being a "realist." After all, Putnam insists that ontological commitment is always internal to a conceptual scheme; there is no scheme-independent fact of the matter about the ultimate furniture of the universe. Putnam, then, is no more a realist than is Kant—and for many contemporary philosophers, that is to be no realist at all.

Failure to recognize the "realistic" motivations for Putnam's rejection of metaphysical realism has led to a widely shared misunderstanding of Putnam's arguments against metaphysical realism. Realist critics of these arguments, convinced that they pose no serious threat to their views, typically

offer rebuttals that are either unapologetic appeals to realistic intuitions or demonstrations that the arguments in question are simply misdirected and thus *irrelevant* to the truth of metaphysical realism. Responses of this kind do not defeat Putnam's arguments; they fail to confront them. A principal shortcoming of these standard interpretations is their disregard for the *role* that these arguments play in the overall case that Putnam is building against metaphysical realism, a disregard which ultimately deflects their argumentative force. Simply put, Putnam's arguments are intended to show that *metaphysical realism itself is not sufficiently realistic*. If that claim can be substantiated then Putnam can go on to argue that his own view is, in relative terms, *more realistic* than metaphysical realism.²

Such widespread misunderstanding of Putnam's arguments requires more than a simple corrective—it requires an *explanation*. I will attempt to provide that explanation in the first half of this paper (Sections I–III). Attention will be focused upon several popular interpretations of Putnam's favorite arguments. These interpretations are *not* wildly irresponsible; they are, in fact, quite natural and reasonably motivated. Understanding where and why they miss the intended force of Putnam's arguments will reveal something of interest about the nature of the realism-antirealism dispute itself. Once we have understood why Putnam's arguments receive the interpretations that they do, we have the background against which the arguments can be properly understood—as *reductios*. The second half of the paper (Sections IV–VI) will make clear why Putnam thinks that metaphysical realism is not sufficiently “realistic” and why he judges his own internal realism to be more in keeping with the “realistic spirit.” This also provides an answer to a question that, on standard interpretations of Putnam's arguments, remains altogether mysterious: Why would Putnam—blessed with strong “realist” intuitions—have abandoned realism on the force of *these* arguments?

My interest in all of this is not to defend internal realism. Putnam's arguments notwithstanding, I have not myself abandoned metaphysical realism. However, he has forever changed my relationship to realism. First, because I think that he is absolutely right that many popular contemporary versions of realism violate the “realistic spirit” in ways too little appreciated. A version of realism that does not compromise this tradition is a robust kind, out of favor with many contemporary physicalist-realists and not easily defended. And second, because he is also right that realism can be alienating.³ We need a philosophical perspective that does not undervalue the “human perspective” *not* simply because our lives are more apt to flourish if we do, but because, as I have argued elsewhere,⁴ the failure to give proper place to the epistemic perspective of human beings inevitably leads to false *semantic* theories.

I. METAPHYSICAL REALISM

Putnam has argued that metaphysical realism, while initially appealing to the realistic spirit in us, can be shown to be irreconcilable with that spirit. I define “metaphysical realism” as a *conjunction* of commitments—ontological, semantic, and alethic. To be a metaphysical realist on my account is to be committed to the following three *metaphysical tenets of realism*:

(M1) *Correspondence Truth*. Truth is a relation of correspondence between pieces of language and the world (i.e., *ding-an-sich*-reality). A statement is correspondence-true *iff* it bears the (unique) relation “correspondence” to *ding-an-sich*-reality.

(M2) *Semantic Realism*. Statements that express an existential commitment to concrete objects (middle-sized and theoretical) will be true or false in virtue of the intrinsic nature of mind-independent reality, and thus in virtue of conditions the obtaining of which may be, in principle, inaccessible to human beings.

(M3) *Ontological Realism*. All (or most) of the objects (middle-sized and theoretical) countenanced by twentieth-century science and common sense exist independently of any mind.⁵

While Putnam’s definition of metaphysical realism is slightly different from this one,⁶ these are the tenets to which metaphysical realists typically have the fiercest loyalty, and Putnam offers reasons for rejecting all three.⁷ When he abandoned realism, Putnam embraced a view he calls “internal realism,” supplanting (M1) with what can broadly be considered a coherence theory of truth, deposing (M2) for an idealized verificationist semantics, and replacing (M3) with an “empirical” realism that is idealistic in the Kantian tradition.

There are a variety of arguments that brought him to his conversion. I shall consider three of the most famous: Dummett’s “language acquisition argument” (as I shall call it), Putnam’s “model-theoretic argument,” and Putnam’s “brains in a vat argument.” My aim here is not to provide a detailed reconstruction of these arguments—such reconstructions are available in abundance. Rather, I will provide a brief sketch of the basic intuition behind each of the arguments and a sampling of typical and predictable realist responses to them. The purpose of this survey is to show how remarkably easy it is for realists to resist these arguments, relying primarily on two distinctive strategies: The first, the “*Burden of Proof*” gambit (the subject of Section II) calls for realists to simply refuse to accept the burden of proof which the language acquisition argument and the model-theoretic argument are (purportedly) attempting to place upon the realist; the second, the “*Shielding Metaphysical Realism From Attack*” strategy (the subject of Section III) demonstrates how easily the force of both the model-theoretic

argument and the brains in a vat argument can be directed away from metaphysical realism, preserving it from harm. These two responses to Putnam's arguments are so widespread (coming from realist and antirealist commentators alike) that it is important to appreciate the very reasonable assumptions which ground them.

II. THE BURDEN OF PROOF

One of the arguments responsible for Putnam's rejection of metaphysical realism is Michael Dummett's *language acquisition argument*. For the past twenty-five years Dummett has forcefully argued that realism cannot ultimately be reconciled with any plausible theory of understanding. How is it that we come to learn a language, to learn what it means to assert that p ? On one reasonable account we learn to recognize the conditions under which competent speakers of a language are disposed to assert that p . That is, we acquire the skill of recognizing, from among the conditions that are epistemically accessible to us, those conditions which must obtain for p to be assertible. Now, if what it means to assert that p just is to assert that justification (or verification) conditions (possibly of an idealized sort) obtain, then we have at least the beginning of a theory of how a finite language-speaker could learn the meaning of p . If, however, these conditions of assertibility are at most *evidence* of the truth of p , where the truth of p consists in something entirely non-epistemic—say, for example, “bearing the relation of correspondence to mind-independent reality”—then, it would seem, we are nowhere near a theory which could account for how finite language-speakers could come to grasp realist truth-conditions and thus know what p means. If realist truth-conditions could be learned, it is reasonable to think that we would have some idea about how such learning takes place. Since we don't know how we could have the capacity, then we probably don't have it.

This is an argument which, according to Putnam, played a prominent role in his conversion to internal realism. At a crucial moment, when Putnam was struggling with the obstacles that a substantive correspondence theory of truth must ultimately overcome, he spent time (in Jerusalem) discussing these issues with Dummett and became convinced that Dummett was, in the main, correct. In his 1976 Presidential Address to the APA, where he announced his conversion, Putnam says:

The point is that Dummett and I *agree* that you can't treat understanding a sentence (in general) as knowing its truth conditions; because it then becomes unintelligible what *that* knowledge *in turn* consists in. We both *agree* that the theory of understanding has to be done in a verificationist way. . . . But now it looks as if . . . I have given Dummett all he needs to demolish metaphysical realism—a picture I was wedded to!⁸

While there are some who find Dummett's argument compelling, many a realist finds this little more than an attempt to shift the burden of proof, unjustifiably, to the realist. If realism is true, says the realist, then an empirical theory of understanding of a Dummettian kind is just what one would *not* expect to be available. Just as in other philosophical disciplines—ethics, epistemology, logic, etc.—the strategy is to refuse to concede the skeptic any ground.

For purposes of comparison, consider the skeptic about the external world and his stance on the skeptical conditional

(S) If I am deceived by an evil demon, then this is not a hand.

The skeptic insists that *modus ponens*, the closure principle, and our inability to know that the antecedent of (S) is false demonstrates that we do not know that this is a hand. Is it incumbent upon the realist to give an argument to show that the antecedent is false, or otherwise to abandon the claim that she knows that this is a hand? Many realists about the external world have insisted that the answer is, "No!" In the spirit of G. E. Moore, the realist is apt to say that *modus tollens*, the closure principle, and the known falsity of the consequent of (S) demonstrates that we know that we are not being deceived by an evil demon. This is an instance of the most common kind of philosophical standoff. (As Putnam himself is fond of saying, "One person's *modus ponens* is another person's *modus tollens*.")

The same strategy is available in the case of the Dummettian challenge over a theory of understanding. Coincidentally, the skeptical conditional (S) plays a role in this dispute as well. For realists, (S) expresses a coherent skeptical worry; for antirealists, since the antecedent is (by stipulation) not verifiable, the conditional is trivially false. Thus, the conditional around which the dispute over a theory of understanding revolves is as follows (note that (S) is mentioned in the consequent):

(U) If there exists no true theory specifying how finite language-speakers grasp realist truth-conditions, then human beings cannot grasp the meaning of (S) realistically interpreted.

The antirealist argues that the lack of a plausible theory of understanding makes untenable the realist's claim that she does in fact understand the meaning of (S) as a genuinely *skeptical* conditional. The realist, of course, insists that the undeniable fact that she understands the meaning of (S)—as a skeptical hypothesis and, thus, realistically interpreted—justifies the claim that there exists a true realist theory of meaning even if we are never able to articulate it to the verificationist's satisfaction. (This is the "*modus tollens*" strategy: Deny the consequent and infer the falsity of the antecedent.) Consider the following expression of this "*modus tollens*" strategy:

The "demon hypothesis" is not just a *noise* that happens to evoke some "pictures in the head"; it is a grammatical sentence in a

language; it is one we can offer free translations of; it is subject to linguistic transformations; we can deduce other statements from it and also say what other statements imply it; we can say whether it is linguistically appropriate or inappropriate in a given context, and whether a discourse containing it is linguistically regular or deviant. The verificationists would retort: "It doesn't follow it has *meaning*". But they would just be wrong, for this is just what meaning is: being meaningful is being subject to certain kinds of recursive transformations, and to certain kinds of regularities; we may not know much more about the matter than that today, but we know enough to know that what the verificationists were propounding was not an analysis of meaning but a persuasive redefinition.

Where does the burden of proof ultimately lie on the issue of the coherence of semantic realism? Must the antirealist give some knockdown argument (not yet in evidence) to shake the realist's strong intuition that she can entertain the "demon hypothesis"? Or is the burden of proof on the realist to produce some plausible theory to support the claim that human beings have the capacity to grasp realist truth-conditions?

The passage just quoted is of particular interest because Putnam himself is the author, *circa* 1971.⁹ At this point in time, Putnam knew of Dummett's arguments, but he did not believe that these placed the burden of proof onto the semantic realist. We may not know a great deal, Putnam argues here, "but we know enough . . ." The burden, he thought, lay with the antirealist. By 1976 he had changed his mind. Why? What convinced him that the burden of proof had shifted? Ultimately, we will attempt to answer that question. First, however, it is important to understand why Putnam's *own* arguments tend to receive a similar response—i.e., the burden of proof gambit—from present-day realists.

THE MODEL-THEORETIC ARGUMENT

Complementing Dummett's language acquisition argument is Putnam's *model-theoretic argument*.¹⁰ Putnam has contributed a whole family of arguments that advance what we might call the "problem of reference" objection to metaphysical realism. Each of these arguments is intended to show that literally *nothing* that twentieth-century philosophers believe in could fix the reference of terms of a human language to mind-independent objects. And if that is the case, we lack the capacity to either speak or think about mind-independent objects and the traditional realist picture collapses. Putnam has advanced various versions of this argument,¹¹ and there is considerable disagreement over the proper reconstruction of the details. It will suffice for this discussion, however, to offer the basic intuitive argument that requires little in the way of complicated machinery.

According to the correspondence theory of truth there is no logical relation between the truth of a statement and the justification of it. This generates a principled distinction between how the world is in itself (independent of the human perspective) and how the world appears (even in the long run) from the human perspective. If this distinction can be collapsed, then correspondence truth will be undermined and metaphysical realism with it. Putnam argues that a necessary condition for correspondence truth and the accompanying distinction is the existence of a determinate relation of reference between our words and mind-independent objects. Otherwise there simply could be nothing in virtue of which an epistemically ideal theory could come out false. The core of his argument is quite simple: Assume that we have a theory, *T*, which meets all theoretical and operational constraints. *T* is then epistemically ideal. Now, if truth is correspondence to mind-independent reality, it is possible that *T* is false. If it is false, then reality will not be as *T* says it is. For example, assume that *T* is committed to the existence of tables, electrons, and planets, among other things. Reality, let us assume, includes no such entities, but is populated instead by F's, G's, and H's, entities radically different from anything countenanced by *T*. On the face of it, it seems that the situation we have just described underwrites the metaphysical realist picture. Putnam demurs, and asks: How can we be sure that *T* does *not* countenance F's, G's, and H's? We know that there will be an interpretation of the sentences of *T* according to which terms of *T* (terms like 'table,' 'electron,' and 'planet') refer to objects in the real world, for example, F's, G's, and H's. We also know that there will be an interpretation of this kind such that *T* comes out *true*. If there is an interpretation of *T* which not only meets all theoretical and operational constraints but also comes out true, then in virtue of what could it fail to be the proper interpretation of *T*? The realist will be quick to insist that an interpretation according to which 'table' refers to F's is a deviant interpretation: it is not the one *intended* by the speakers. Putnam has a ready reply. He demands that an account be given of what there could be (under heaven and earth) that could conceivably *fix* the so-called "intended" interpretation. It will not do, in the late twentieth century, to appeal to "magical" theories of reference or occult properties.

The first move available to the realist is the "burden of proof" gambit which requires nothing in the way of philosophical argument. One simply refuses to admit that the burden of proof is on the realist. "Why," asks Carsten Hansen, "does Putnam think that it is incumbent upon realism to prove that our terms have determinate extensions?"¹² David Lewis states the matter more bluntly:

Since Putnam's paradoxical thesis is patently false, we can be confident that there is indeed some further constraint, *whether or not we can find out what it is*.¹³

The model-theoretic argument poses no serious threat, in Lewis's mind. Since he takes it to be an indisputable fact that our words have determinate reference, we are confronted only with a paradox that we may attempt to resolve or not as we wish.

Realist objectors to the model-theoretic argument rarely leave the matter here, however. Most, including Lewis, are prepared to meet Putnam's demand by offering some account of what the reference-fixer might actually be. Putnam separates these theories into two broad categories: (a) those candidates that he considers to be scientifically respectable but incapable of actually fixing reference, and (b) those candidates that might have the metaphysical clout to fix reference but which require—precisely because of their metaphysical extravagance—commitment to entities, properties, or powers of dubious scientific merit. The first category includes the most popular candidate for reference-fixer, "causal connections of the appropriate type." The second category includes solutions which, to Putnam, smack of "medieval essentialism"¹⁴ or "magical" powers.¹⁵ Two realist responses that apparently fall into the latter category are (i) David Lewis's suggestion that there exist "elite properties" (properties that cut reality at the joints) which provide an additional constraint on reference,¹⁶ and (ii) Alvin Plantinga's simply stated claim that we have the power to "grasp properties."¹⁷ Putnam gives little attention to these unabashed pronouncements that there exists potent metaphysical machinery which can fix reference. Instead he focuses upon the one candidate which (in his opinion) has, because of its modesty, the potential of being part of a plausible empirical theory with scientific respectability and genuine explanatory power. This is the view that *causal connections* provide the additional constraint that will enable reference to be fixed.

Putnam himself used to assume that causal connections were the most likely answer to the reference-fixing problem. Given the model-theoretic argument, however, he became convinced that the causal theory lacks the necessary resources to meet the challenge. How, exactly, is a causal connection supposed to accomplish the task? Certainly we can add the sentence " 'Cat' refers to mind-independent cats by virtue of a causal connection of the appropriate type" to our theory, *T*. But this is just "more theory" according to Putnam. 'Causal connection' is just another linguistic symbol which itself will lack determinate reference without an acceptable reference-fixer. It simply begs the question to assume that your language has a determinate reference while you are in the process of showing how that reference is fixed.

Or does it beg the question? The realist will insist that it is not the linguistic symbol 'causal connection' that fixes reference, it is causal connections themselves that do the work. To assume that sentences which the realist uses to express her theory of reference-fixing do not have determinate reference is itself to beg the question against the realist. Michael Devitt expresses it this way:

Putnam, in effect, accuses the metaphysical realist of begging the question in appealing to a theory to determine reference for a theory. I have accused him of begging the question in claiming that the reference of 'causally related' is not determinate.¹⁸

Again the question arises, where does the burden of proof lie? Is producing a sentence about "causal connections" sufficient? Even if causal connections are considered to be possible reference-fixers, has anything really been explained? Aren't causal connections at best brute facts with little explanatory value? Isn't the burden of proof on the realist? No, argues Anthony Brueckner,

We maintain that the use of language together with non-linguistic facts about the world (e.g., causal relations between the world and the use of language) do fix the intended interpretation of the language. Unless Putnam can discredit the foregoing claim, he has given us no reason to suppose that every attempt at a linguistic specification of a theory's intended interpretation must fail.¹⁹

Brueckner implies that Putnam has given no substantive argument to discredit the causal theory and so he has given the realist no reason to abandon his position.

The purpose of this inquiry is not to arrive at a judgment on the merits of either Putnam's argument or this particular realistic response to it. Rather, it is to make explicit how the realist perceives her position and why she considers the "burden of proof" response a natural one. After all, it seemed natural to Putnam himself in 1971. If the realists' initial convictions are reasonable (as they obviously assume they are), then any argument that is interpreted as nothing more than a demand for a more sophisticated theory will hardly be received as a demonstration that metaphysical realism is incoherent. If one comes to the table assuming that the burden of proof is on the anti-realist to give a knockdown refutation of realism, it is little wonder that the language acquisition argument and the model-theoretic argument are perceived to have so little force.

III. SHIELDING METAPHYSICAL REALISM FROM ATTACK

The "burden of proof" gambit is basically a justification for ignoring an argument, denying that it provides an objection substantial enough to warrant rebuttal. Realist critics of Putnam are, however, rarely satisfied with a bald dismissal of his arguments. They are more than happy to provide a careful analysis of the details of each argument. Such an analysis frequently produces the following results: It is found that the arguments, even if sound, turn out *not* to entail the falsity of metaphysical realism. Since the arguments are

not strictly *about* metaphysical realism, why would they give Putnam or anyone else reason to abandon that view? Employing this second strategy, one simply protects metaphysical realism whatever the cost. David Lewis argues that *even if Putnam is correct*, even if reference is *not* fixed by anything and thus is indeterminate, realism still survives. The model-theoretic argument is simply about the wrong thing to be successful against metaphysical realism. Putnam has greatly overstated the consequences that would follow even if it were shown that reference is radically indeterminate. He explains it this way:

My point is rather that even if the model-theoretic argument worked, it would not blow away the whole of the realist's picture of the world and its relation to theory. Something vital would be destroyed, but a lot would be left standing. There would still be a world, and it would not be a figment of our imagination. It would still have many parts, and these parts would fall into classes and relations—too many for comfort, perhaps, but too many is scarcely the same as none. There would still be interpretations, assignments of reference, intended and otherwise. Truth of a theory on a given interpretation would still make sense, and in a non-epistemic way. Truth on all intended interpretations would still make sense. Despite Putnam's talk of the 'collapse' of an 'incoherent picture', he has given us no reason to reject any of these parts of the picture. The only trouble he offers is that there are too many intended interpretations, so that truth on the intended interpretations is too easily achieved. That is trouble, sure enough. But is it *anti-realist* trouble, except by tendentious definition?

... The *metaphysics* of realism survives unscathed. What does suffer, if Putnam has his way, is realist semantics and epistemology.²⁰

Lewis's basic idea, I take it, is that if the heart of realism is a commitment to correspondence truth and to a mind-independent reality to which true propositions correspond, then acceptance of the model-theoretic argument does not force one to abandon realism. Our language will lack a unique, intended interpretation and so each sentence of our language will express a whole host of propositions, but this is not meaning skepticism in the sense that we simply don't know which unique proposition is being expressed; this is meaning pluralism in the sense that each sentence expresses as many propositions as there are interpretations of *T* which meet theoretical and operational constraints. This picture does not *abandon* correspondence truth and mind-independent reality; it literally *presupposes* it. The result, of course, will be a raging semantic pluralism that hardly leaves traditional realist semantics in place and wreaks considerable havoc with epistemology. Still, "The *metaphysics* of realism survives unscathed," according to Lewis, and since it is intended to be an argument against metaphysical realism, the

model-theoretic argument misses its mark. Admittedly, Lewis would have to pay a fairly high price to continue to love metaphysical realism in the face of this rather unseemly and promiscuous semantics. Yet love can make a person do remarkable things, and a true love of metaphysical realism can be sure to inspire tremendous feats of philosophical courage. The model-theoretic argument clearly lacks the necessary clout to *force* one to abandon metaphysical realism.

We have arrived at an important juncture. Having seen that the “burden of proof” gambit is altogether reasonable from the realist’s perspective and having seen that metaphysical realism can be preserved come what may, the following seems unavoidable: *If the success of Putnam’s arguments depends upon either (a) shifting the “burden of proof” by simply demanding more theory from the realist, or upon (b) forcing the realist to abandon metaphysical realism on pain of contradiction, then his arguments must fail.* Metaphysical realism can be maintained without logical contradiction even if reference is maddeningly pluralistic; the burden of proof cannot be shifted onto the realist simply because the antirealist would like a bit more theory. While this may seem to put Putnam at a distinct disadvantage vis-à-vis his realist opponents, such is not the case. Not only are these concessions not fatal to Putnam’s arguments, but recognizing that Putnam is happy to grant these concessions is a *prerequisite* for understanding his arguments. Putnam has been there before. He knows all too well that these dialectical moves are available to the realist. (He has, after all, used them himself!) The case he has built against metaphysical realism is not merely a demand for more theory, *nor* is it the bald claim that metaphysical realism entails a logical contradiction. The main thrust of his argument lies elsewhere.

To better appreciate the spirit of Putnam’s arguments, consider Lewis’s suggested strategy against a successful model-theoretic argument. Lewis says that one ought simply to marshal all available defenses and protect metaphysical realism from attack. Instead of allowing metaphysical realism to take the “hit” from the argument, sacrifice instead one’s semantic and epistemological commitments. There is no logical constraint to prohibit this strategy, and all *true* realists will recognize their responsibility to preserve metaphysical realism *come what may*. First, it is important to notice the very different *thresholds* of philosophical suffering that individuals are willing to endure before crying uncle and abandoning their basic presuppositions. What does the model-theoretic argument show if successful? Does it show (A) that we may have the inconvenience of “too many” intended interpretations, or (B) that the notion of “intended interpretation” collapses into incoherency on the metaphysical realist picture? Does it show (A) that metaphysical realism may reasonably be preserved as one’s defining paradigm so long as certain semantic and epistemological corrections are made, or (B) that the altogether bizarre and counterintuitive consequences in semantics

and epistemology make it an effective *reductio* of metaphysical realism itself? We know that philosophers can stick to their guns come what may, so long as they are willing to live with the consequences. Lewis, as we know, is unlikely to be deterred by our incredulous stares.

Lewis's strategy to protect metaphysical realism come what may provides a convenient foil to show off Putnam's arguments to their greatest advantage, revealing the *role* that these arguments ultimately play in Putnam's case against the classic realist picture. Putnam is well aware of the fact that the "burden of proof" gambit carries little conviction by itself. If a person has strong intuitions in favor of metaphysical realism and if that person detects no serious liabilities that cause unavoidable conflicts among other of her convictions, then Putnam's arguments will hold little conviction. Putnam assumes, however, that there are in fact deep and irreconcilable tensions within the traditional realist perspective—evidence that there is something disturbingly "anti-realistic" about the view. If this can be established then he may successfully shift the burden of proof by showing that metaphysical realism has absurd and intolerable consequences (even if it cannot be proven that it is logically inconsistent).

Consider a parallel situation. In normal science there is often a lack of a plausible theory—with no immediate hope of a breakthrough—right where scientists would most like to have one. This need not jeopardize fruitful research, nor pose any threat to accepted theories. Historians of science have helped us to see that normal science continues, that the overall theoretical assumptions that have been successful in the past continue to be employed without question and without (additional) justification in spite of data that do not fit, and in spite of theories that never materialize. Bona fide crises do not arise merely for the lack of answers to theoretical questions; crises arise only when there is some tension within the accepted theory which is intolerable, some tension which is increasingly detrimental to the fruitful practice of normal science.

I want to suggest that this is an apt model for the present case. It is true that the realist may not have as thorough and sophisticated a theory about realist semantics and correspondence truth as she might like. However, too little theory, alone, rarely precipitates a crisis; it often does no more than motivate more "normal philosophy." Crises arise *not* when there is too little theory but when deep conflicts are detected *within* the theory, conflicts that seem irreconcilable without a radical adjustment in the defining paradigm.

The force of Putnam's case against realism is not simply a demand for more theory from the realist. The arguments gain their persuasive force against a background of unanticipated anomalies that Putnam believes makes a "revolution" the only reasonable course. It is commonplace for realists

today to think that the causal theory will take care of the problem of reference and that traditional realist commitments will be preserved. Putnam thinks otherwise. There is a crisis within the realist paradigm, he argues, and there is simply no way to preserve the traditional picture intact. It may be possible to maintain a commitment to metaphysical realism without falling into logical contradiction, but one will be forced to abandon many strongly held convictions and, if Putnam is correct, protecting metaphysical realism is neither a rational nor a “realistic” thing to do.

Most philosophers will agree, I think, that realism in the modern period was not simply a narrow, isolated commitment to metaphysical realism, i.e., (M1)–(M3). Traditional realists have always taken for granted that one of the primary virtues of correspondence truth is that it seems to guarantee a particular commonsense view of semantics and epistemology. If it turns out that correspondence truth does no such thing, few realists, I think, would possess Lewis’s steely nerve in the face of adversity and be willing to abandon these strongly held convictions. If the model-theoretic argument *did* obviously succeed, Putnam *would* have a compelling objection to metaphysical realism—not because metaphysical realism together with radical indeterminacy constitute a logical contradiction, but because they are incompatible with strongly held convictions—semantic and epistemological—that are an inseparable part of the traditional realist picture.

Lewis’s argument, then, is a perfect example of the kind of pressure that Putnam seeks to place on traditional realism. It is only necessary that he show that metaphysical realism has serious, possibly intolerable, consequences in some part of the traditional realist’s philosophical picture. Even if we assume that this analysis is accurate, little has been done to actually further Putnam’s case. No realist that I am aware of (including Lewis) is willing to grant that the model-theoretic argument is successful in the least. As we have already seen, the realist assumes that there is no good reason to doubt referential determinacy and that Putnam’s argument is at best the unreasonable demand that the realist say something *more* about how causal relations fix reference and at worst the impossible demand that the realist say it in a way that doesn’t presuppose that reference is determinate.

All of this notwithstanding, Putnam believes that his arguments do reveal a crisis of irreconcilable tensions and anomalies within the traditional realist picture. The realist, for her part, is having considerable difficulty understanding how the model-theoretic argument does anything but beg the question. Putnam anticipates this difficulty, however, and so offers another argument, the brains in a vat argument, as a practical demonstration of the unexpected consequences that correspondence truth has in the areas of semantics and epistemology.

THE BRAINS IN A VAT ARGUMENT

The brains in a vat argument, if it is effective, will make manifest certain irreconcilable tensions that lie at the heart of the standard realist picture. It will help to show something that the model-theoretic argument alone could not: That even if we assume that “causal relations of the appropriate type” *do* fix reference, there are consequences that follow in the area of semantics and epistemology that are *not* reconcilable with strongly held convictions of most traditional realists. Realists have had great difficulty interpreting the brains in a vat argument as such an argument.

First a brief summary of the argument. One crucial premise of the argument is that we postulate (what Putnam considers to be) a plausible theory of reference. Appealing to insights that he helped to popularize in his work on the new theory of reference, Putnam insists that there is a condition that must be met for successful reference—the *causal requirement*:

One cannot refer to certain kinds of things, e.g. trees, if one has no causal interaction at all with them, or with things in terms of which they can be described.²¹

Invoking this requirement, Putnam argues that a brain in a vat—whose experiences are the result of computer-generated electrical stimulation and who bear no causal relation to real trees—cannot *refer* to real trees regardless of what mental images the brain produces and regardless of how many tokens of the symbol-string ‘t-r-e-e’ it produces in its computer-generated environment. Thus, an envatted brain asserting “I am a brain in a vat” does not refer to the *real* vat in which it is suspended nor to the *real* brain that it is, since it does not bear a causal relation of the appropriate type to either one. What does it refer to? Well, possibly to a feature of the computer program, or to the electrical impulses that cause the brain’s experiences, or to the phenomenal images themselves—“tree-images”, “vat-images”, and the like.

Putnam’s enticing conclusion is that we can know, therefore, that we are not BIVs. If I exist in a normal world consisting of real trees, real brains, and real vats, then the English sentence, “I am a brain in a vat,” does express a claim about real brains and real vats, a claim which is, however, *false*. If, on the other hand, I actually *am* a brain in a vat, bearing no (appropriate) causal relation to real vats, then my utterance of “I am a brain in a vat” is ultimately a claim about electrical impulses or vat images. In that circumstance, the statement that I would express by such an utterance would be something like, “I am presently having experiences as of being a brain in a vat.” But of course if *this* is what it means, it is straightforwardly false. With this constructive dilemma, we see that the skeptical threat is disarmed. I can be assured that the hypothesis expressed by “I am a brain in a vat” is false and thus I know that I am *not* a brain in a vat.

If this argument is to prove effective against the realist, it must reveal some consequence of metaphysical realism that the realist herself finds unacceptable. There has been a very common and predictable response to this argument, however, according to which it is no threat to realism whatsoever. While there are numerous articles that make exactly the same claim,²² I shall consider Anthony Brueckner's version because I find his overall approach to the argument especially instructive. Brueckner admits that there is some question about the force with which the argument is to be taken. The conclusion of Putnam's argument seems to be: *That I am a brain in a vat* is not a logically possible proposition. This conclusion does not follow from the premises, says Brueckner, and Putnam is wrong when he says

In short, if we are brains in a vat, then 'We are brains in a vat' is false. So it is (necessarily) false. (15)

Brueckner takes Putnam's last sentence to express the claim that there exists a proposition (i.e., the proposition that is expressed when a person in the real world utters the sentence "I am a brain in a vat") such that *that* proposition is necessarily false. How could that be true? Brueckner asks. If it were, then there would be *no possible world* in which I am a brain in a vat. But the brain in a vat scenario offered by Putnam himself *just is* such a logically possible world. No, Brueckner insists, Putnam has not shown that the relevant *proposition* is necessarily false, but only that a *sentence*—a sentence that expresses different propositions in different contexts—will express a false proposition in all possible worlds in which the sentence is betokened.

All I can claim is the metalinguistic knowledge that a certain sentence expresses a false proposition, rather than the object-language knowledge that I am not a brain in a vat. Since the latter knowledge was required in order to refute the skeptical argument in the envisaged manner, the present anti-skeptical strategy fails.²³

Brueckner rightly determines that one of the premises of the BIV argument *presupposes* metaphysical realism, just as Descartes's evil demon scenario does. He concludes, then (somewhat hesitantly), that the argument purports to show metaphysical realists that there is a "tricky" way to solve Descartes's skeptical dilemma.²⁴ The basic strategy of the BIV argument, if Brueckner is correct, is something like this:

1. Assume that metaphysical realism is true.
2. Describe a brains in vat scenario.
3. Show that no one can truly believe that she is a brain in vat.
4. Claim to have provided a clever way of undermining Cartesian skepticism.

5. Conclude, therefore, that even if metaphysical realism is true radical skepticism is not a real threat.

First it is important to recognize the parallel between this response to the BIV argument and Lewis's claim that the model-theoretic argument, even if successful, does not count against correspondence truth. Both responses refuse to countenance an interpretation of the argument that makes it a *reductio*. In Lewis's case the refusal is simply a bravura "in your face" stance which he (gleefully) takes against Putnam. In Brueckner's case there is some reason to suspect that he doesn't actually recognize that it is intended as a *reductio*. Throughout the paper, he admits his uncertainty about this interpretation and in a footnote to the first paragraph, he says:

The argument of chapter 1 ["Brains in a Vat"] indeed depends upon causal-theoretic assumptions about reference which Putnam explicitly rejects in chapters 2 and 3.

At no point, however, does he explicitly entertain the possibility that it is a *reductio* and thus he makes no attempt whatever to consider what force it might have, so interpreted. It is quite common, and not a wholly disreputable business, for analytic philosophers to consider a philosophical argument as an isolated text, as if it washed up on the beach somewhere with no clues as to its intended argumentative *force*. However, in this case, where the argument clearly *presupposes* metaphysical realism, and where the book, of which it is the opening chapter, is clearly committed to demonstrating the bankruptcy of metaphysical realism, it is difficult to miss the role that the argument is meant to play. Putnam makes it quite clear when he asks:

Why is it surprising that the Brain in a Vat hypothesis turns out to be incoherent?²⁵

Brueckner's interpretation is puzzling. How could it seem reasonable to interpret this argument as a straightforward rebuttal of Cartesian skepticism, presupposing the truth of metaphysical realism? The answer, I suspect, is that many commentators simply cannot figure out how the argument could plausibly have any force as a *reductio* against someone strongly committed to metaphysical realism. To substantiate this claim, consider Peter van Inwagen's interpretation of the argument. He rightly concludes that a proper interpretation of the argument must not ignore the fact that the BIV argument is advanced as an argument against metaphysical realism. He characterizes the argument this way:

(1) If truth is radically nonepistemic, then it is intelligible to suppose that we are all wrong about almost everything.

(2) It is not intelligible to suppose that we are all wrong about almost everything.

hence, Truth is not radically nonepistemic.²⁶

While conceding that Putnam never makes this argument explicit, he says “But if he does not mean to be defending this argument, or some argument very much like it, . . . then I do not see the point of that chapter.” This is a good beginning, if only because the conclusion is consistent with the stated purpose of the book. Van Inwagen is, then, asking the right question: *Why* would Putnam (or anyone else) think that the BIV argument raises any difficulty for the metaphysical realist?

Van Inwagen offers a reasonable answer to this question. The argument will have force against the metaphysical realist because realists have traditionally believed that (1) is true and now Putnam has given us reason to believe that (2) is true (because, given the causal requirement, even brains in vats will not be wrong about everything). Van Inwagen’s response to the argument is predictable. He argues that even if one grants Putnam premise (2), for the sake of argument, one still need not concede the conclusion. Premise (1) is still vulnerable and, if one is really committed to non-epistemic truth, it will be sacrificed. Realists may have *thought* that (1) is true—and even with some justification. But as Putnam points out, realists also have tended to believe that “*what goes on in our heads* must determine what we mean and what our words refer to.”²⁷ If that belief proves unwarranted, then premise (1) will lose much of its intuitive force. The obvious strategy for realists, van Inwagen argues, is to give up (1) and hold on to correspondence truth. What could it hurt?

Why shouldn’t the “copyist,” when he has seen Putnam’s argument for the incoherency of supposing that we are brains in vats simply say, “Oh, now I see,” and continue to be a copyist? Why shouldn’t the proponent of the thesis that truth is radically nonepistemic react in that way, too? . . .

From the realist’s point of view, Putnam’s argument has nothing to do with the nature of truth. It has nothing to do with the nature of the relation that must hold between a proposition and the world in order for the latter to *confer truth* on the former. It has, rather, to do with the nature of the relation that must hold between a proposition and a subject in order for the latter to *accept* the former. The realist who accepts Putnam’s argument will describe its import like this: Putnam has shown that certain of the necessary conditions for accepting the proposition that everyone is a brain in a vat—or for accepting any proposition about brains and vats—could not be satisfied by anyone if everyone were a brain in a vat. Of course (the realist might add), if everyone were a brain in a vat, this proposition would be *true*; it is just that no one would be able to accept it or grasp it.²⁸

Notice that we have now come full circle. The last sentence quoted above reflects Brueckner’s reconstruction of the argument. Van Inwagen is well aware that this is *not* how Putnam wants the argument to go, but he sees no

way that Putnam could prevent the realist from taking this way out. After all, there is no strict *logical* relation between correspondence truth and *any* epistemological position. One can hold on to correspondence truth come what may. It is only necessary that the realist admit that a BIV could not think that it was a BIV and that our beliefs about the external world could not be globally false in the way of Cartesian skepticism. No problem. What has any of this to do with realism?

As a matter of fact, I think that Cartesian skepticism has a great deal to do with traditional realism. In the remainder of this paper I hope to show that the realist pays a much higher price for these concessions than van Inwagen might lead us to believe.

IV. THE REALISTIC SPIRIT

Putnam insists that he has abandoned metaphysical realism for reasons that are genuinely “realistic” in character. He says:

Recognizing such facts as these is part of what might be called “rejecting ‘realism’ in the name of the realistic spirit.” It is my view that reviving and revitalizing the realistic spirit is the important task for a philosopher at this time.²⁹

Most realists are not convinced by this talk. They tend to think that there is nothing very “realistic” about Putnam’s position and that his arguments will be compelling only to those who already have antirealist leanings. Michael Devitt holds this view. He sees nothing “realistic” about internal realism, which he considers to be a version of “radical idealism” and which he assumes will have appeal only to those with “anti-realist” intuitions.

In a vivid and comprehensive way Putnam has captured most of the intuitions that motivate anti-realism. These intuitions have always had some appeal in philosophy and have recently become popular.³⁰

On what basis, then, does Putnam claim that his view is motivated by the *realistic spirit*? It hardly need be said that if the “realistic spirit” is nothing more nor less than a disposition to embrace the (three) metaphysical tenets of realism, then there is nothing realistic about Putnam’s position. But is realism, traditionally understood, nothing more nor less than the conjunction of these three doctrines? Admittedly, there is some justification for identifying “realism” with the three metaphysical tenets of realism in recent literature because of the semantic focus of the realism-antirealism dispute on the contemporary scene. Nonetheless, realism has a long and rich history in the modern period, and fixation on the metaphysical tenets is a relatively

recent phenomenon. I shall argue, further, that even today most self-proclaimed realists are committed to realism more broadly construed.

It is not uncommon for “realism” to be considered as much a philosophical attitude, or even a temperament, as it is a doctrine. The “realistic spirit” reflects a comprehensive intellectual perspective which is often typified by a commitment to certain familiar doctrines, but is not necessarily exhausted by those doctrines and could, conceivably, require a rejection of some of them. Crispin Wright has captured at least part of the perspective I have in mind in the following passage:

The realist in us wants to hold to a certain sort of very general view about our place in the world, a view that . . . mixes modesty with presumption. On the one hand, it is supposed, modestly, that how matters stand in the world, what opinions about it are true, is settled independently of whatever germane beliefs are held by actual people. On the other, we presume to think that we are capable of arriving at the right concepts with which to capture at least a substantial part of the truth, and that our cognitive capacities can and do very often put us in position to know the truth, or at least to believe it with ample justification. The unique attraction of realism is the nice balance of feasibility and dignity that it offers to our quest for knowledge. . . . We want the mountain to be climbable, but we also want it to be a real mountain, not some sort of reification of aspects of ourselves.³¹

This broad “realistic” perspective has found expression in a commitment to a variety of distinct philosophical doctrines reflecting the particular interests of different historical periods. Today, realism is assumed to be a theory about ontology and/or semantics and truth. Not very many years ago, however, the most common use of the term “realism” was to refer to a doctrine (or group of doctrines) that reflects the *epistemological* orientation of the spirit of realism. While there are any number of beliefs that “epistemological realists” have typically held, I will focus attention on what I shall call the *four epistemological tenets of realism*, beginning with the core doctrine:

(E1) *Epistemological Realism*. Belief in the mind-independent existence of all (or most) of the objects of common sense and of science is epistemically justified.

A wealth of different realistic epistemological theories have been advanced over the years, each offering a separate account of *how* it is that our beliefs about the external world are justified.³² As a result of its commitment to mind-independent objects, therefore, epistemological realism takes skepticism quite seriously. It is logically possible that the ultimate nature of reality is, in principle, inaccessible to the cognitive and sensible faculties of human beings and thus it is possible, as Putnam says, that

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*.³³

To take skepticism seriously, in this sense, is to hold that skepticism is *coherent*. Thus, realists of this stripe have been committed to

(E2) *The Coherence of Skepticism*. The skeptical conditional
(S): If we are all brains in a vat, then this is not a hand.
is both intelligible and true.

This doctrine reflects the *modesty* of which Wright speaks. The mountain to be climbed is no mole hill and success is no mere consequence of having justified beliefs. This is an expression of an attitude central to the realist spirit, viz., that the human pursuit of knowledge is a risky affair and that the strides that have been made in our knowledge of the world represent genuine victories won in the face of very real obstacles. Thus, skepticism is intelligible, but false; skepticism is taken seriously, but not too seriously. Not only is the mountain real but we *really* are climbing it. According to the realist, then, (S) is true, but both the antecedent and the consequent are false. (E1) and (E2), then, lead naturally to:

(E3) *Anti-Skepticism*. Most of our beliefs about common sense and scientific objects are *known* to be true and the brains in a vat hypothesis is *known* to be false.

Finally, since epistemological realism presupposes that the beliefs which result from a careful exercise of our faculties, both cognitive and sensible, possess (to a greater or lesser degree) significant epistemic warrant, this is a perspective that has historically resisted any philosophically motivated *revision* of the deliverances of common sense or of scientific inquiry. Phenomenalism, for example, was immediately resisted by epistemological realists—even before the internal difficulties of the view were made manifest—simply because it constitutes too radical a revision of what is assumed to be the common sense (pre-philosophical) interpretation of our discourse. Thus, the fourth epistemological tenet:

(E4) *Anti-revisionism*. Since we are justified in believing that our cognitive and sensible faculties give us reliable access to reality, there is *prima facie* justification to resist any radical, philosophically motivated revision of the deliverances of common sense and of scientific inquiry.

While *anti-revisionism* is rarely picked out as a distinct doctrine, it is a regulative principle that has shaped the epistemological attitude of a great many realists through the years.³⁴ Recapping the taxonomy of *traditional realism*, thus far presented, we have

SEVEN TENETS OF TRADITIONAL REALISM

Three Metaphysical Tenets of Realism

- (M1) Correspondence Truth
- (M2) Semantic Realism
- (M3) Ontological Realism

Four Epistemological Tenets of Realism

- (E1) Epistemological Realism
- (E2) The Coherence of Skepticism
- (E3) Anti-skepticism
- (E4) Anti-revisionism

Today, it seems that realism has been largely disassociated from epistemology, and many philosophers are untiring in their insistence that realism is *not* an epistemological doctrine.³⁵ The four epistemological tenets of realism, though still *believed* by most realist-minded philosophers, are no longer considered to be *necessary* for being a “realist.” The metaphysical tenets—especially (M1) and (M2)—are considered to be the ineliminable core of “realism,” both necessary and sufficient for being a realist.

Perhaps a reasonable case can be made for making (M1) and (M2) the heart of “realism.” Note, first, that one can be committed to (M1) and (M2) without being either an ontological realist or an epistemological realist, but one cannot be either of the latter two without being committed to (M1) and (M2).³⁶ On this account, then, a skeptic, even a radical skeptic, can be a realist. And that is precisely as it should be. The skeptic and the Cartesian both share a common vision of the philosophical landscape and thus of the epistemological enterprise. They share a common understanding of our cognitive aspirations, of the meaning of our discourse, and of the goal of epistemology. They agree that our linguistic capacities allow us to make claims about mind-independent reality and that our claims to knowledge are (for the most part) claims about the intrinsic nature of reality. They agree that we can engage in speculative metaphysics and thus consider the possibility that reality might be radically different from what it appears to be. Both are equally animated by the “realistic spirit” in this sense, with the former simply a bit more pessimistic about the extent to which our beliefs are justified and (in some cases) a bit more pessimistic about the truth of our present theories. Surely this difference—a disagreement about the truth of (E1)—is not what separates realists and antirealists. The skeptic who embraces (M1) and (M2) but rejects (E1) is misleadingly described as an “antirealist.” If realism-antirealism is a dichotomy, there is certainly justification for placing the traditional skeptic on the realist side of the divide.

All this is well and good. Let us assume that this picture of the

philosophical enterprise does in fact capture something of the “spirit of realism.” And, on this picture, the skeptic and the traditional realist do share a commitment to (M1) and (M2). This, however, should give little comfort to the realist who insists that Putnam’s arguments have nothing to do with realism. It is true that the realist and the skeptic hold (M1) and (M2) in common but that is hardly the *only* thing they share, nor is it the most important. The conflict between epistemological realism and skepticism in the modern period does not begin with an agreement about theories of truth and meaning. Rather, it begins with a consensus about the *coherence of skepticism*, (E2), with presumptions about semantics and truth being largely implicit and little developed. In fact, we typically attribute a commitment to (M1) and (M2) to a philosopher on the basis of her acceptance of (E2) because we assume that (E2) *entails* (M1) and (M2).

The philosophical vision which the ontological realist and the skeptic supposedly share is only coherent if Cartesian skepticism itself is coherent. But this is precisely why Putnam’s arguments have force against the classic realist picture. Putnam argues that correspondence truth and semantic realism when constrained by a plausible view of language (i.e., the causal requirement for reference) not only cannot guarantee the coherence of Cartesian skepticism but also are in fact incompatible with it. It may be true that radical skepticism presupposes correspondence truth and semantic realism; but if Putnam is correct, the reverse is not true. Correspondence truth and semantic realism, when joined with the causal requirement, are incompatible with radical skepticism.

What does all of this mean? As we shall see in the following section, Putnam’s brains in a vat argument reveals that introducing the causal requirement for reference into the traditional realist picture produces troubling tensions and anomalies which cannot be casually dismissed. Putnam challenges the assumption that a metaphysical realism which accepts the causal requirement (let’s call it *causal realism*) preserves most of the things that realists typically care about.

V. WHAT IS “REALISTIC” ABOUT METAPHYSICAL REALISM?

The brains in a vat (BIV) argument is intended to manifest the counterintuitive consequences of *causal realism* (i.e., metaphysical realism plus the causal requirement). While the BIV scenario presupposes metaphysical realism, the rather surprising implications for both semantics and epistemology ultimately give the argument the force of a *reductio*. This is not to suggest that metaphysical realism leads to logical contradiction. As we have seen, the realist can certainly avoid that. However, if Putnam is correct, merely avoid-

ing contradiction will not necessarily provide much comfort to the realist. Recall that there was a time when a person could continue to believe in the existence of the aether even in the face of the negative results of the Michelson-Morley experiment. The only price one had to pay was to admit that bodies “shrink” in size as their velocity with respect to the aether increases. While this concession saved the aether theory from outright contradiction, it seemed an intolerable consequence for most observers and today seems an effective *reductio* of the theory. Putnam’s *reductio* of metaphysical realism is analogous. To be successful he need not show that causal realism is logically inconsistent, he need only show that it requires unanticipated and ultimately unacceptable revisions somewhere within the realist’s philosophical perspective. It may be that the causal realist can muster the courage to hold on to correspondence truth. What Putnam will not concede is that this is a reasonable thing to do nor that the result is reflective of the realistic spirit.

Most philosophers who assess Putnam’s BIV argument begin by granting him the *causal requirement*. This is not necessarily because they believe that he is right that there *is* such a requirement, but because they believe that realism will not be threatened *even if* he is right. Having granted the causal requirement, they also grant that a brain in a vat who says, “That rhododendron is six feet high,” is not making a false assertion about a physical, mind-independent shrubbery, but is instead making a true assertion about electrical impulses, a computer program, or “rhododendron” images. It follows, then, that a BIV who asserts

(S1) If we are all brains in a vat, then there exists no six-foot-high rhododendron

is *not* expressing the same kind of skeptical worry that Descartes expressed. A necessary condition for a skeptical conditional to be of a Cartesian sort is that the antecedent express some non-standard theory about ultimate reality. The antecedent of a BIV’s utterance of (S1) expresses no such thing. At best it asserts a (philosophically uninteresting) falsehood about electrical impulses, a computer program, or “vat” experiences; at worst, it is incoherent. Since BIVs have no cognitive access to genuine Cartesian conditionals, they lack the capacity to do philosophy in anything like the way that we take ourselves to be doing it.

It is most telling that realists who assess Putnam’s BIV argument rarely voice any dissatisfaction with this revisionist interpretation of the BIV’s discourse. We might have initially thought that a BIV would be capable of doing philosophy, just as we do it. Once it is conceded, however, that the cognitive range of the BIV’s language is seriously limited because of the limits of its causal reach, it must be admitted that BIVs will be serious under-achievers as philosophers. This is not to say, of course, that a BIV could not utter the following words:

I believe that we are brains in a vat. If so, then in normal discourse our term 'chair' may refer to electrical impulses or to "chair"-images. However, when I say "We are brains in a vat" I do not mean to say that electrical impulses are causing us to have *experiences* as of being brains in a vat, *that*, of course, is clearly false. Rather I intend my words to reach beyond the immediate causal web within which we are trapped and to make a claim about a real computer-independent vat.

There is no reason to doubt that a BIV, having experiences phenomenologically indistinguishable from our own and having dispositions similar to our own, might make such a speech. If Putnam is correct, however, this speech is plagued by incoherency. The BIV simply cannot be saying what we would take ourselves to be saying if these were our words.

Note the concession that the realist makes in denying that the BIV has the capacity to entertain a Cartesian conditional. On the classic (Cartesian) picture, the BIV's utterances are interpreted quite differently. A BIV's everyday "rhododendron" statements would come out false, but the BIV's philosophical speech would come out both coherent and true. Part of the perspective shared by the traditional realist and the skeptic is that philosophical discourse of this kind not only makes sense but also reflects an amazing and ennobling human capacity. It is assumed that human beings would possess this capacity *even if they were brains in a vat*. It is this latter assumption, of course, which the realist must abandon if Putnam is correct. BIVs do *not* possess the capacity to entertain radically skeptical hypotheses. BIVs are *not* capable of doing philosophy in this way. The foregoing philosophical speech is not simply misguided, it is *nonsense*. We must conclude, therefore, that if we were BIVs we would lack these capacities as well. If realists find this outcome at all distressing, relatively few of them express that distress in print.

There are several reasons, I think, why realists rarely consider it a major concession to grant that a BIV's cognitive repertoire may be seriously limited. In the first place, it is easy for us to shrug our shoulders and say: "Well, it is no wonder that there are odd consequences given the bizarre nature of the BIV scenario." But there is a second reason that is far more important. There is a strong temptation to assume that it is only brains in vats whose cognitive powers are seriously constrained by the causal requirement. It is frequently assumed that none of the foregoing argument gives *us*—those of us who are *not* BIVs—any reason to worry. *We* can certainly do philosophy just the way we suppose ourselves to do it because our epistemic situation is crucially different from the BIVs. BIVs cannot speculate about wildly different ultimate ontologies; their language goes awry when they contemplate statements like (S1). Our language, however, the language of non-BIVs living in a world of mind-independent objects, will be plagued by no such

difficulties. We retain the resources necessary to do philosophy, to speculate about radically non-standard theories about reality, and to worry (unnecessarily) that some such theory might be true. The difference, we are told, is that speakers in the actual world bear the proper causal relation to the *kind* of objects that make speculation about the ultimate nature of reality possible. Thus, for example, while the BIV cannot think that it is a brain in a vat, we can think that we are because our term 'vat' refers to the very kind of physical vat that would house our brains were we brains in a vat. Thus, causal realism limits the cognitive range of the BIV but not of us. The only concession we need make is that *if* we were brains in a vat, *then* we could not think that we were. Given the fact that we aren't BIVs, this is a concession of no consequence.

But this is not a proper assessment of the situation. Causal realism of the kind under consideration is far more revisionist than this picture indicates. It does *not* leave speakers in the actual world with sufficient resources for philosophical reflection of the Cartesian type. Consider the skeptical conditional:

(S1) If we are all brains in a vat, then there exists no six-foot-high rhododendron.

There is a substantial constraint on the kind of reflection that is possible on the causal realist picture, *even in the actual world*. While I can reflect upon a counterfactual possible world in which I am a brain in a vat and can speculate about the truth value in that world of 'The rhododendron is six feet high,' there is one thing that I cannot do. I cannot (consistently) entertain the thesis that the *actual* world is a world in which I am a BIV. I cannot consistently think that my present reflection on the skeptical conditional is the reflection of a brain in a vat. The limitation arises because the two cases require different interpretations of the antecedent of the conditional. In the case where I assume I am in the actual world imagining how things are in some counterfactual possible world, (S1) will be in English, and the antecedent will assert that my physical brain is in a physical vat. So far so good. If, on the other hand, I am asked to consider the genuinely *skeptical* hypothesis that the actual world is a world in which I am a BIV and thus that my present reflection of (S1) is the reflection of a BIV—that I cannot do. To accomplish that I must take (S1) first to be a sentence in vat-English, and second to be the occasion of my entertaining the possibility that I am a brain in a vat. I cannot do both. If I do the first, then (S1) is a sentence that a BIV could possibly entertain, but it does not express the hypothesis that I am a brain in a vat—and, thus, I fail to do the second. If I do the second, I am presupposing that I am not a brain in a vat and thus I fail to do the first. To entertain the possibility that the actual world is a BIV world I must be able coherently to believe both that some sentence I entertain expresses the claim

that I am a BIV and that no sentence I could possibly entertain expresses the claim that I am a BIV. Attempts to entertain the possibility that the actual world is a BIV world will get no further than do attempts to entertain the possibility that my present cognitive state is the state of someone who is unconscious (and cognitively inactive).

As a consequence of this, the causal realist must compromise two epistemological tenets of realism. Skepticism, at least of the traditional Cartesian sort, is *not* coherent, and thus (E2), "The Coherence of Skepticism," must be rejected. Further, because human beings are disposed to believe that skepticism is coherent, at least some of what they say will ultimately be incoherent. Causal realism must be *revisionist*, then, just as verificationism is revisionist—determining that many utterances of otherwise competent speakers have (unbeknownst to them) some fatal semantic defect. A causal realist must, therefore, reject (E4), "Anti-revisionism."

One of Putnam's most profound insights is his recognition that causal realism shares many of the apparent weaknesses of verificationism while lacking many of its strengths. For many years Putnam resisted the temptations of verificationism because it was revisionist: it told people that they couldn't say many things that, as otherwise competent language-speakers, they were prone to say. Realism, it seemed, put no such constraints on our linguistic capacities, providing semantic resources that verificationism lacked. As Putnam explored the implications of causal realism, however, he made a startling discovery: causal realism does *not* provide the resources that it promises. Cartesian skepticism is no more coherent on the causal realist picture than it is with verificationism. While verificationism limits our cognitive access to that which is (at least in principle) verifiable, causal realism limits our cognitive access to that which is within our causal sphere. The limitations of causal realism, however, are far more insidious. Realism was originally embraced because it promised to give us the semantic resources to express all those thoughts that we recognized (from our own epistemic perspective) as being cognitively meaningful. Causal realism is not merely revisionist, it does not merely legislate that certain types of discourse are incoherent, but it makes us strangers to our own thoughts and utterances. The privileged access that we always assumed that we had to the content of our own thoughts is now a myth. John Heil, one of the few commentators on the BIV argument who addresses this issue, expresses the situation well:

Let us suppose for the moment, however, that Putnam is right: meanings are, if they are anywhere at all, someplace other than in the head; content is externally determined. A startling corollary of such a view seems to be that agents have no special access to the content of their own thoughts. I am in no privileged position to appreciate what my thoughts are about simply in virtue of their being *mine*.³⁷

The meaning of our utterances, on the causal realist picture, is an external matter. The cognitive content of my thoughts and utterances is, in a very real way, not only beyond my control but also beyond my epistemic access. The ties between the epistemic perspective of the speaker and the meaning of his utterances is severed. The *significance* of my own cognitive activity is removed altogether from the realm of the epistemic. Putnam rejects this view of language and insists that it is not consistent with the spirit of realism because it alienates speakers from their own language.

In the earlier book [*Reason, Truth and History*] I described current views of truth as 'alienated' views, views which cause one to lose one or another part of one's self and the world; in these lectures I have tried to elaborate on this remark, and on the connection between a non-alienated view of truth and a non-alienated view of human flourishing.³⁸

It may reasonably be argued that Putnam himself is responsible for this view of language. He is the one who helped create this revolution in the philosophy of language by arguing that "meanings just ain't in the head," and by insisting upon the causal requirement. Putnam, of course, stands by these claims. The extension of the term 'gold' is fixed by (possibly unknown) facts about objects to which I bear a certain causal relation. Putnam's whole point, however, is that this very reasonable theory of reference, when joined to a verificationist semantics, has none of the consequences just described. It is only when the causal requirement is joined with metaphysical realism that the speaker is alienated from her own thoughts. Internal realism is, for Putnam, a more "realistic" philosophical perspective than is causal realism.

On a similar note, Crispin Wright spoke of the "dignity that [realism] offers to our quest for knowledge." How does causal realism fare on this measure? The familiar rap against verificationism is that knowledge is too easily achieved. Since there is no room for radical skepticism, there is no risk of radical failure. The epistemological enterprise is thought, then, to be diminished in value, to lack the nobility that comes from facing a challenge where one risks complete failure. On the causal realist picture, however, the risk of radical failure seems to be missing as well. The reason that the brains in a vat are not engaged in the traditional epistemological enterprise which realists have found so ennobling is that the brains are caught in a *causal web*. Their words cannot escape the limits of that web and thus they cannot even think the kind of thoughts necessary for philosophical speculation of the traditional realist sort. They also cannot *fail* to be right in their beliefs about the world, assuming that those beliefs meet sufficient theoretical and operational constraints. Regardless of what goes on in their heads, their beliefs about "tables" and "chairs" *just will* be true beliefs about whatever reality is ultimately causing them to have table- and chair-experiences—whether that reality is a computer program or an evil demon.

The meaning of our utterances, on the causal realist picture, is an external matter. The cognitive content of my thoughts and utterances is, in a very real way, not only beyond my control but also beyond my epistemic access. The ties between the epistemic perspective of the speaker and the meaning of his utterances is severed. The *significance* of my own cognitive activity is removed altogether from the realm of the epistemic. Putnam rejects this view of language and insists that it is not consistent with the spirit of realism because it alienates speakers from their own language.

In the earlier book [*Reason, Truth and History*] I described current views of truth as 'alienated' views, views which cause one to lose one or another part of one's self and the world; in these lectures I have tried to elaborate on this remark, and on the connection between a non-alienated view of truth and a non-alienated view of human flourishing.³⁸

It may reasonably be argued that Putnam himself is responsible for this view of language. He is the one who helped create this revolution in the philosophy of language by arguing that "meanings just ain't in the head," and by insisting upon the causal requirement. Putnam, of course, stands by these claims. The extension of the term 'gold' is fixed by (possibly unknown) facts about objects to which I bear a certain causal relation. Putnam's whole point, however, is that this very reasonable theory of reference, when joined to a verificationist semantics, has none of the consequences just described. It is only when the causal requirement is joined with metaphysical realism that the speaker is alienated from her own thoughts. Internal realism is, for Putnam, a more "realistic" philosophical perspective than is causal realism.

On a similar note, Crispin Wright spoke of the "dignity that [realism] offers to our quest for knowledge." How does causal realism fare on this measure? The familiar rap against verificationism is that knowledge is too easily achieved. Since there is no room for radical skepticism, there is no risk of radical failure. The epistemological enterprise is thought, then, to be diminished in value, to lack the nobility that comes from facing a challenge where one risks complete failure. On the causal realist picture, however, the risk of radical failure seems to be missing as well. The reason that the brains in a vat are not engaged in the traditional epistemological enterprise which realists have found so ennobling is that the brains are caught in a *causal web*. Their words cannot escape the limits of that web and thus they cannot even think the kind of thoughts necessary for philosophical speculation of the traditional realist sort. They also cannot *fail* to be right in their beliefs about the world, assuming that those beliefs meet sufficient theoretical and operational constraints. Regardless of what goes on in their heads, their beliefs about "tables" and "chairs" *just will* be true beliefs about whatever reality is ultimately causing them to have table- and chair-experiences—whether that reality is a computer program or an evil demon.

We of course assume that our power to philosophize is not threatened in this way. After all, we bear the proper causal relation to real tables and real brains so that we are capable of the ennobling philosophical thoughts of a traditional realist sort. But in what relevant way does our situation differ from that of the BIVs? We are also caught in a causal web. We cannot speak of or otherwise think about anything that lies outside that causal web. We may *think* that our epistemological search is genuinely threatened with the possibility of a radical failure. We may think that our search for knowledge has a kind of nobility because we are willing to acknowledge the limitations of our cognitive capacities and to consider the implications of the fact that justified beliefs, even in the long run, may not necessarily be true. But all of this is unjustified sentimentality. There is a very real sense in which we cannot *fail* to hold mainly true beliefs in the way just described, any more than the BIV can. The meaning of our words will be causally tied to *whatever* reality is ultimately causing us to have the experiences that we do, and so we *just will* end up holding true beliefs about *that* reality. Thus, brains in vats will hold true beliefs about electrical impulses. And we, too, are likely to hold true beliefs. Certainly something of the “dignity” of the quest for knowledge is lost inasmuch as the wondrous powers of “causal relations of the appropriate type” will literally *carry* us up the mountain, regardless of how bad our epistemic situation might be. If there is more room for the “dignity” of a “real climb up the mountain” with causal realism than with verificationism, it is difficult to see how.

VI. WHAT IS “REALISTIC” ABOUT PUTNAM’S INTERNAL REALISM?

The primary objectives of this paper have now been met. I have argued that Putnam’s arguments do put pressure on any traditional realist who is committed to the causal requirement. Metaphysical realism can obviously be maintained come what may. Nonetheless, intuitions that many have considered to be all but constitutive of the “realistic spirit” must be forsaken. If this is accurate, then metaphysical realism together with a plausible theory of language results in a philosophical picture that is not fully “realistic” in the traditional sense. Admittedly, Putnam’s internal realism is not fully “realistic” in the traditional sense, either. The central question then is this: All things considered, is causal realism or internal realism more in keeping with the realistic spirit? I do not presume to answer this question. Certainly *any* answer would be controversial. In the brief discussion to follow, I shall do no more than to argue that Putnam’s internal realism keeps much of the spirit of *each* of the seven tenets of traditional realism. This will

not establish that internal realism is more “realistic” than causal realism; it will show, however, that internal realism has a reasonable claim to the attribution, “realistic.”

(M1) *Correspondence Truth*. Putnam has no problems with correspondence truth on a certain, commonsensical reading: ‘The tree *is* in the quad’ will come out true only if the tree really is in the quad; our statements will not be true unless the world is the way we say it is. His objection to correspondence truth traditionally conceived is that it naively assumes that it makes sense to talk about our words being firmly fixed to some uninterpreted, thing-in-itself reality. While Putnam does dismiss this account of truth, he does not dismiss our pre-philosophical convictions about the meaning of the word ‘true.’ Instead he argues that the actual role that the word ‘true’ plays in our discourse can be accounted for without appeal to the traditional realist picture. Let’s return to Wright’s metaphor of the mountain. If the realistic spirit requires that there be a mountain that is a genuine obstacle to be overcome, Putnam retains much of that. It *is* a real mountain! It is not *merely* a reification of our faculties because our faculties may, in the end, be incapable of carrying us all the way up the mountain. Achieving cultural consensus or consensus among any finite group of human beings is not sufficient for truth. Truth is a regulative ideal for Putnam. More than that, it is an ideal that is not necessarily achievable. The very real mountain that we are attempting to climb may be insurmountable, given the laws of physics and the limitations of human beings’ sensible and cognitive faculties. Of course, a great deal of what we believe is *true* on Putnam’s account. Our beliefs about tables and chairs, for example, are paradigmatic examples of beliefs that *do* meet our highest epistemic standards. However, while we have made considerable progress up the mountain, there is no guarantee that we will reach the summit. In the area of particle physics, for example, it is at least logically possible that at some deep level of investigation, there may simply be no theory which we are capable of grasping which will meet all the theoretical and operational constraints that we ourselves demand that a successful theory meet. To say that an epistemically ideal theory could not be false (which Putnam does say) is *not* to say that *in every conceivable area of inquiry* it is possible for human beings to achieve an ideal theory. Not only is ‘truth’ not a name for “the best that a culture can accomplish,” it is not even a name for “the best that human beings can accomplish.” Putnam’s view of truth, while not the doctrine that traditional realists have typically embraced, is nonetheless a view that makes sense of a good part of our discourse about truth, leaving the spirit of much of that discourse essentially intact.³⁹

(M2) *Semantic Realism*. It is often assumed that there can be nothing “realistic” about verificationist semantics because the heart of semantic realism is its countenance of “mind-independent objects”—and that is precisely what is left out of the verificationist account. One of the primary reasons that

we are committed to “mind-independent” objects is that we believe that physical objects are not causally dependent upon any mind, and thus that the following is true:

Dinosaurs existed before human beings did and would have existed even if human beings had never existed.

If it follows from verificationism that (a) ‘Dinosaurs existed 50 million years ago’ is nothing more nor less than a claim about the existence of certain bones to be found in museums, or that (b) it is a contradiction to say that dinosaurs would have existed even if humans had not (since there would be no humans around to do any verifying), then verificationism is not very realistic. Putnam’s verificationism does not have these consequences. His is an idealized verificationism where the truth of a statement will be determined by its verifiability under ideal epistemic conditions and in the long run. For Putnam the dinosaur statement is true because if a human being had been present 50 million years ago then the dinosaur would have been verifiable (i.e., if we had been there, we would have seen it). Dinosaurs are ontologically “independent” of human beings in this regard, and this is independent enough for much of what we want to say.

(M3) *Ontological Realism*. It is widely recognized that Putnam’s internal realism includes a commitment (at the empirical level) to the objects of both common sense and of science. Phenomenalism, for example, which is a monistic ontology (i.e., all objects are merely permanent possibilities of experience) cannot do justice to statements like “Micro-physical entities ultimately explain why tables and chairs possess the macro-properties that they do.” Internal realism, however, allows for a commitment to electrons as well as tables and chairs, together with all of the levels of explanation that are presupposed by an “objective” (as opposed to a merely operational) interpretation of scientific theories. Putnam’s empirical realism includes a full rich ontology, the envy of any metaphysical realist.

(E1) *Epistemological Realism*. Putnam’s complex ontology (as previously described) allows him to hold an epistemological position that shares much in common with the traditional doctrine of epistemological realism. We believe in the existence of the objects of common sense and of science and those beliefs are justified. Assuming that these beliefs are true (and on Putnam’s view the electron beliefs are certainly not guaranteed to be true) it is meaningful (and not trivial) to say that they are known to be true. Much of the epistemological orientation of traditional realism is preserved.

(E2) *The Coherence of Skepticism*. As has been argued previously, Putnam, like the causal realist, must reject Cartesian-style skepticism. This hardly means, however, that there is no coherent form that skepticism can take. We can be skeptical about ever arriving at an epistemically ideal theory (see the discussion of (M1) above). More than that, we can worry

about the coherent possibility that none of our present experiences of external objects are veridical—not in the transcendental way presupposed either by Cartesian skepticism or by Putnam's own BIV scenario, but in an empirical way. It is perfectly meaningful for Putnam to say: "I may have been kidnapped last night by evil scientists and plugged into a computer which is presently giving me experiences of trees that I am incapable of discriminating from 'real'-tree experiences." This is a hypothesis that is not only coherent on Putnam's view but which raises genuine epistemological worries. Given the intelligibility of this hypothesis, the epistemological enterprise is not an altogether trivial one.

(E3) *Anti-skepticism.* We have already seen how internal realism is anti-skeptical. We are justified in believing in the objects of common sense and of science. It is logically possible that I have been captured by evil scientists, but I am warranted in believing that that hypothesis is false.

(E4) *Anti-revisionism.* Putnam's strongest case against causal realism may well rest upon his claim that internal realism is less revisionist than causal realism. Of all the epistemological tenets of realism, Putnam has been most reluctant to give up "Anti-revisionism." He has been loath to advocate a philosophical revision of discourse that humans take to be not only *meaningful* but *valuable*. He believes that Rorty, Foucault, and Derrida are unacceptably revisionist, that they are gripped by the following idea:

The failure of our philosophical "foundations" is a failure of the whole culture, and accepting that we were wrong in wanting or thinking we could have a foundation requires us to be *philosophical revisionists*. By this I mean that, for Rorty or Foucault or Derrida, the failure of foundationalism makes a difference to how we are allowed to talk in ordinary life—a difference as to whether and when we are allowed to use words like "know," "objective," "fact," and "reason."

... I am not, in that sense, a philosophical revisionist.⁴⁰

While Putnam's position does not trivialize terms like 'know' and 'objective' he does admit that his view has revisionist implications. Discussing his rejection of the realism-idealism dichotomy, he confesses that it was a difficult thing to do. Nonetheless, the attempt to save all of the old intuitions is hopeless.

My rejection of these dichotomies will trouble many, and it should. Without the constraint of trying to 'save the appearances', philosophy becomes a game in which anyone can—and, as a rule does—say just about *anything*. Unless we take our intuitions seriously, we cannot do *hard* philosophy at all. So I respect philosophers who insist that the traditional dichotomies are deeply intuitive, and who 'need a lot of convincing' before they will give them up.

On the other hand, the whole thrust of Putnam's work over the last fifteen years is to convince us that the traditional realist paradigm is unstable. Revision is a last resort, but revision is unavoidable. The only question is: Which revision will you choose, and why?

But if philosophy which simply scorns our intuitions is not worth the candle, philosophy which tries to preserve *all* of them becomes a vain attempt to have the past over again.

... The task of the philosophers, as I see it, is to see which of our intuitions we can responsibly retain and *which* we must jettison in a period of enormous and unprecedented intellectual, as well as material, change.

If I reject the dichotomies I depicted, it is not, then, because I fail to recognize their intuitive appeal, or because that intuitive appeal counts for nothing in my eyes. It is rather because these dichotomies have become distorting lenses which prevent us from seeing real phenomena—the phenomena I have been describing—in their full extent and significance.⁴¹

Putnam is correct. The “alienated” view of language with which causal realism is saddled violates the spirit of realism with *extreme prejudice*. While verificationism also puts certain constraints on our cognitive activities, they are constraints that we can live with. At least the *meaning* of our discourse remains intimately connected to the epistemic features which make our use of that language intelligible.

Internal realism does not preserve every conceivable realistic intuition. It does, however, preserve a great deal of what realists have traditionally cared about. Internal realism may not be *more* realistic than causal realism, that is difficult to say. It certainly is, however, realistic in ways that causal realism is not.

NOTES

1. Research for this paper was supported in part by an ISU University Research Grant, for which I am grateful. I would like to thank the following people for helpful comments: Ann Baker, Larry Bonjour, Pat Franken, Patrick Murphy, Hilary Putnam, Mark Siderits, Bob Steinman, Mark Timmons, Charles Travis, and especially Thom Carlson.
2. Here is one passage where Putnam suggests, explicitly, that he may be more realistic than some metaphysical realists: “That truth *is* a property . . . is the one insight of ‘realism’ that we should not jettison. But Hartry Field shows signs of being inclined to jettison this insight, although he calls himself a ‘metaphysical realist’ and says that I am a ‘nonrealist.’ Could it be that I am more of a realist—though not a ‘metaphysical’ one—than Field, after all?” *Realism with a Human Face*, James Conant, ed. (Cambridge, Mass.: Harvard University Press, 1990), 32.
3. *The Many Faces of Realism* (Lasalle, Ill.: Open Court, 1987), 1.

4. "Semantic Dualism." (submitted for publication).
5. For those tempted by Michael Devitt's claim that realism is exclusively an ontological doctrine not a theory about truth or meaning see note 36 below for an extended discussion of the matter.
6. In *Reason, Truth and History* (Cambridge: Cambridge University Press, 1981), 49, Putnam suggests that the metaphysical realist also believes "There is exactly one true and complete description of 'the way the world is'." I do not include this requirement because few philosophers think that it is a necessary condition for being a realist, and it plays no crucial role in the arguments to follow.
7. There is some justification for claiming that Putnam may not necessarily reject (M3), ontological realism. It all hinges on how the qualifier "independently of any mind" is cashed out. Since I interpret "independently of any mind" as invoking a realist semantics, Putnam is not an ontological realist.
8. *Meaning and the Moral Sciences* (Boston: Routledge & Kegan Paul, 1978), 129.
9. *Philosophy of Logic* (London: George Allen & Unwin Ltd, 1971), 65–66.
10. I have benefited greatly from discussions with Mark Siderits about the model-theoretic argument.
11. "Realism and Reason" in *Meaning and the Moral Sciences*, 123–38, "Models and Reality" in *Realism and Reason: Philosophical Papers*, vol. 3 (Cambridge: Cambridge University Press, 1983), 1–25, and "A Problem about Reference," chapter 2 of *Reason, Truth and History*, 22–48.
12. Carsten Hansen, "Putnam's Indeterminacy Argument: The Skolemization of Absolutely Everything," *Philosophical Studies* 51 (1987): 89.
13. David Lewis, "Putnam's Paradox," *Australasian Journal of Philosophy* 62 (1984): 225 (my emphasis).
14. *Realism and Reason*, xii.
15. *Reason, Truth and History*, 3, 51.
16. Lewis, op. cit., 226–29.
17. "How to be an Antirealist," *Proceedings and Addresses of the American Philosophical Association* (September 1982): 61.
18. *Realism and Truth* (Princeton: Princeton University Press, 1984), 190.
19. "Putnam's Model-Theoretic Argument Against Metaphysical Realism," *Analysis* 44 (1984): 137.
20. Lewis, op. cit., 231–32.
21. The second disjunct in this quote (*Reason, Truth and History*, 16–17) suggests that a *direct* causal connection is not a *necessary* condition for referring to concrete objects. This is borne out later when he says "The idea that causal connection is necessary is refuted by the fact that 'extraterrestrial' certainly refers to extraterrestrials whether we have ever causally interacted with any extraterrestrials or not!" (ibid., 52).
22. Among those who consider an interpretation of the BIV argument that is similar to Brueckner's are: John D. Collier, "Could I Conceive Being a Brain in a Vat?" *Australasian Journal of Philosophy* 68 (1990): 413–19; Paul Coppock, "Putnam's Transcendental Argument," *Pacific Philosophical Quarterly* 68 (1987): 14–28; J. Harrison, "Professor Putnam on Brains in Vats," *Canadian Journal of Philosophy* 17 (1987): 427–36; Gary Iseminger, "Putnam's Miraculous Argument," *Analysis* 48 (1988): 190–95; Michael Kinghan, "The External World Sceptic Escapes Again," *Philosophia* 16 (1986): 161–66; Jane McIntyre, "Putnam's Brains," *Analysis* 44 (1984): 59–61; and James Stephens and Lilly-Marlene Russow, "Brains in Vats and the Internalist Perspective," *Australasian Journal of Philosophy* 63 (1985): 205–12.
23. Brueckner, *Ibid.*, 167.
24. Brueckner actually does believe that the argument rests on a "trick." He says: "Given the presuppositions of our anti-skeptical argument, it is difficult to avoid an uneasy feeling that

there is some trick involved in the reasoning of the last paragraph. To see that there is a trick . . . note that . . .”

“Brains in a Vat,” *Journal of Philosophy* 83 (1986): 164.

25. Putnam, *Reason, Truth and History*, 22.
26. Peter van Inwagen, “On Always Being Wrong,” *Midwest Studies in Philosophy, Volume XII: Realism and Antirealism*, Peter A. French, Theodore E. Uehling, Jr., and Howard K. Wettstein, eds. (Minneapolis: University of Minnesota Press, 1988), 95–112.
27. *Reason, Truth and History*, 22.
28. *Ibid.*, 104 and 106.
29. *Realism with a Human Face*, 42.
30. Review of *Reason, Truth and History* in *The Philosophical Review* 93 (1984): 277.
31. “Realism, Antirealism, Irrealism, Quasi-Realism,” *Midwest Studies in Philosophy, Volume 12: Realism and Antirealism*, 25.
32. If one is interested in a taxonomy of epistemological realisms, Roger Cornman, in *Perception, Common Sense & Science* (New Haven: Yale University Press, 1975), serves up a cornucopia.
33. *Meaning and the Moral Sciences*, 125.
34. Realism has *not* traditionally opposed revision in our ontological commitments when motivated by scientific, rather than philosophical, considerations. The realistic spirit inspired a rejection of beliefs in animism, heavenly bodies, and countless other entities assumed as uncontroversial for hundreds of years.
35. This is something of an overstatement. In the philosophy of perception, for example, the terms “representational realism” and “direct realism” continue to refer to the same epistemological doctrines that they always have. Nonetheless, outside of an explicitly epistemological context, to call a contemporary philosopher a “realist” is usually to attribute to her a commitment to one of the *metaphysical* doctrines discussed below.
36. Michael Devitt rejects the view that (M1) and (M2) are essential to realism and has gained some support for his claim that realism is strictly an ontological doctrine, according to which one can be committed to (M3) without any commitment to either (M1) or (M2). On this view, one can be an ontological realist *without* holding a substantive view of truth and without interpreting any statement realistically. He insists that those, like Quine, who eschew semantics altogether and thus deny that the realism-antirealism dispute is even intelligible can themselves be ontological realists. I find this baffling. On my view, a necessary condition for being an ontological realist about kittens is to interpret (at least some) kitten statements realistically and to believe that so interpreted they are (correspondence) true. If you don’t believe that plants exist, then you can’t be a botanist; if you don’t believe in correspondence truth and you don’t think that statements express realist truth-conditions, then you can’t be a realist. Putnam is correct; metaphysical realism does presuppose a substantive view of truth and meaning. While it is not possible to justify this claim here, I am convinced that Devitt fails in his attempt to find anything in virtue of which Quine will qualify as an ontological realist. In his definition of “ontological realism,” Devitt interprets the phrase “independently of the mental” in such a way that it is *not* necessarily an invocation of correspondence truth or a realist semantics. Thus, Quine can be committed to (M3) without being committed to (M1) or (M2). The problem with this move is that if you do *not* give (M3) a strong, semantic reading then Putnam himself will be happy to embrace (M3). Putnam, after all, believes that dinosaurs existed long before human beings ever arrived on the scene—and *would have* existed even if human beings hadn’t. The kind of empirical realism that is left after substantive theories of meaning and truth are abandoned is too pale and uncontroversial a doctrine to deserve the title “ontological realism.”
37. “The Epistemic Route to Anti-realism,” *Australasian Journal of Philosophy* 66 (1988): 164. Heil offers his own solution to this problem, attempting to show that the causal realist can have introspective access to the content of his own thoughts. The content of this second-order thought, however, will also be determined *externally*: “Presumably, the

content of this second-order, introspective thought is itself determined externally and non-epistemically. Its being about my thought that *p*, requires that certain circumstances external to it obtain . . . it evidently follows that I can, after all, have an introspective grip on the contents of my thought that *p*" (171). While this is a response available to the causal realist, it is difficult to appreciate why it is thought to dispel the worry. A second-order thought, the content of which is itself externally determined and thus non-epistemic, is just one more thought from which I am epistemically alienated. For every thought that Heil produces, the same objection can be raised.

38. *The Many Faces of Realism*, 1.
39. Assuming that this is an accurate account of Putnam's view of truth (and, admittedly, it may not be), one is tempted to say the following: If, in some area of inquiry, there is no ideal theory accessible to human faculties, then there will simply be *no truth* in that domain. This is, of course, quite different from saying that truth exists but is beyond our reach. I leave it to the reader to decide how "realistic" this result is. [I am grateful for discussions with Mark Timmons who helped me to see this point.]
40. *Realism with a Human Face*, 20.
41. *The Many Faces of Realism*, 28–30.