EXPLORING REALISM AND TRUTH

by

Michael Lee Gifford

Defended on: July 20th, 2017

A dissertation submitted to the
Faculty of the Graduate School of
the University at Buffalo, State University of New York
in partial fulfillment of the requirements for the
degree of
Doctor of Philosophy

Department of Philosophy

For Nana

Special thanks to my committee:

David Braun, Chair

James Beebe

Neil Williams

And to the many other individuals who tirelessly provided support, guidance, and discussion.

Table of Contents

Abstract

This dissertation explores some important interconnections between the notions of Realism and Truth. Recently in analytic philosophy, there has been growing resistance to the idea that we might productively approach the realism issue via a consideration of the nature of truth. Michael Devitt, in particular, has argued forcefully and at length that realism is a metaphysical issue which ought to be settled before the semantical issue concerning the sorts of truth conditions our sentences may have. One goal of this dissertation is to establish that the realism issue may be fruitfully addressed through a study of language rather than metaphysics proper. The bulk of the dissertation is devoted to a careful study of one particular semantic approach to the realism question: That of Michael Dummett's.

The first chapter focuses on the dogmatic views of Michael Devitt. Devitt claims that we ought to "put metaphysics first," that is, answer metaphysical questions before semantical and epistemological questions. He also believes that we ought to put metaphysics first in the sense that our metaphysical conclusions should be given a certain kind of immunity from revision in light of what we may conclude in semantics and epistemology. Hence, Devitt argues, it would be a mistake to settle on an anti-realist conclusion based on the determination that our sentences have anti-realist (epistemic) truth conditions. I argue in the first chapter that Devitt's approach lacks sufficient motivation and furthermore requires us to take on a number of unsavory philosophical commitments.

Those who prefer a metaphysical, as opposed to semantic, approach to the realism issue are under special pressure to explain what exactly the nature of mind-independence is. This is because the metaphysical approach to the realism issue says that the question of realism comes down to the question whether what exists does so mind-independently. The second

chapter of my dissertation is devoted to an analysis of the notion of mind-independence. I criticize some common statements of realism in terms of mind-independence and attempt to offer a better version of the view. One major issue of common statements of realism in terms of mind-independence is that they leave it a mystery how we might be realists about minds. I overcome this problem by framing mind-independence in terms of what I call our *directed cognitive practices.*

The third chapter consists of a detailed study of the anti-realist views of Michael Dummett. Dummett happens to be an anti-realist, but his approach to the realism issue is consistent with realism. For Dummett, the realism question comes down to the question as to what sorts of truth conditions our sentences have. His realist claims that our sentences have truth conditions in such a way that they could be true or false despite our inability to ever come to know whether they were true or false. Dummett's anti-realist claims that our sentences could never have such truth conditions, and instead have truth conditions such that we must be capable, at least in principle, of coming to know the sentences' truth values. I criticize various interpretations of Dummett and respond to some common objections. Understanding Dummett's philosophy requires understanding what he means by the term 'undecidable', since he takes the realism dispute to be a dispute as to whether "undecidable" sentences have determinate truth values. I argue that the best way to understand the notion of undecidability in the context of Dummett's philosophy is as the lack of an effective decision procedure, that is, the lack of a procedure which is guaranteed to deliver an answer as to whether the sentence is true or false in a finite amount of time.

The fourth chapter focuses on the question as to whether a Dummettian anti-realist is under any special pressure to reject classical logic in favor of an intuitionistic logic. I conclude

that the Dummettian, who endorses a semantics in terms of proof-conditions rather than truth-conditions, must reject classical logic and instead endorse the logic of the intuitionists. The chapter is largely devoted to criticizing other attempts to discern in Dummett any particular argument for the conclusion that his anti-realist must reject classical logic. A major take-away of the chapter is that Dummett never settled on a particular notion of undecidability that could do all of the work he needed it to do in his arguments while at the same time remaining acceptable to both the realist and anti-realist. I also spend a good deal of time assessing attempts, most notably by Crispin Wright, to generalize the intuitionistic version of anti-realism in Dummett's vein. I conclude that the semantic approach to the realism issue is indeed fruitful and viable, but that the sorts of considerations which might lead one to endorse such a generalized intuitionism will likely be rooted in metaphysical concerns. My ultimate conclusion is that regardless of whether one construes the realism debate metaphysically or semantically, metaphysical considerations can and should influence our semantic conclusions while semantic considerations can and should influence our metaphysical conclusions.

Introduction

The purpose of this dissertation is to explore the structure of an idea that lies at the intersection of other prominent ideas. The idea to be explored is realism. The intersecting ideas are those of language, reality, and the concept of truth. Within the realism debate, there is another debate concerning how the first debate is best understood. That other debate is over whether realism ought to be construed as an issue in the philosophy of language or an issue in metaphysics proper. There is no real doubt that the realism issue is pertinent to metaphysics—it is, after all, supposed to be pertinent to practically every branch of philosophical inquiry. At the same time, there is a certain temptation to think that, at least when one is dealing with metaphysical realism, if not when one is dealing with any brand of realism whatsoever, priority ought to be given to metaphysical considerations. Such a temptation has led many to the conclusion that the realism debate is one that ought to be settled on a metaphysical basis alone. One of the main conclusions of this dissertation is that metaphysical considerations cannot carry all of the burden, even if we must conclude this while admitting that metaphysics can make the difference as to whether we ought to take a realist or an anti-realist attitude toward a given subject matter. This is no contradiction. The nature of the subject matter at hand may be such that certain metaphysical interests tip the scales toward either realism or anti-realism.

Rather recently in analytic philosophy, there has been harsh criticism of the thought that the question of metaphysical realism might be treated as an issue in the philosophy of language (Alston 1996, Devitt 1997, Horwich 1990). These philosophers explicitly resist what has come to be known, perhaps because Rorty once edited a volume with the name, as the "linguistic turn" in analytic philosophy. The resistance would have us turn back around. The thought that metaphysical realism should be left to the metaphysicians, while increasingly more popular

among analytic metaphysicians, threatens to carry out a full-fledged revolt against much of what a few of the undisputed founders of analytic philosophy took to be its most important insight: That many if not all of the most difficult philosophical questions stand to be solved, if not dissolved, by a careful analysis of language. I hope that, even if I am unable to convince anyone that a linguistic approach to the realism issue is the correct one, I am at least able to undo the current dogma among metaphysicians that linguistic considerations have no place in determining whether we ought to be metaphysical realists. And while these philosophers may trace their thought to Aristotle, who gave us the first systematic treatment of the very concept of being qua being, they tend to put too little emphasis on the linguistic considerations in the *Categories,* which set the stage for the depth of the *Metaphysics* (cf. Yu 2003).

Let us return to the realism debate. There are two main strands of thought in analytic philosophy about how to approach the realism question. One is linguistic while the other is metaphysical. According to the linguistic approach, the realism debate is best understood as a debate over the nature of truth. An anti-realist in this context claims that truth is somehow epistemic in nature—what is true depends in some way on what we are capable of coming to know. The realist in this context says that truth depends in no way on what we can know. Instead truth is seen as a property that a sentence can have independently of whether or not we might ever, under any circumstances, be in a position to find out that the sentence is true. Thus we get the claim that, for a realist, truth is evidence-transcendent or recognition-transcendent, while for an anti-realist we get the claim that truth is evidentially constrained in the sense that whether something is true depends on whether we are capable of coming to know that it is. This way of construing the realism issue is considered to be semantic or linguistic in nature because of the connection between truth and meaning. For instance, it is often thought that the semantics of

both formal and natural languages can be given in terms of truth conditions. The semantic construal of the realism issue is common among self-described anti-realists. Paradigmatic treatments of the realism issue centered on the nature of truth can be found in the works of Hilary Putnam (1981, 1983) and Michael Dummett (1978, 1991), both of whom advance their own versions of anti-realism understood as a commitment to an epistemically constrained notion of truth. The third and fourth chapters of this dissertation amount to a sustained study of Michael Dummett's anti-realism.

According to the metaphysical rendering of the realism debate, the issue comes down to whether what exists does so mind-dependently or mind-independently. Those who prefer such a construal of the debate tend, for the most part, to be metaphysical realists. These philosophers eschew the idea that the question of realism can be settled by attending to questions in the philosophy of language. The central proponent of this way of viewing things is Michael Devitt (1997, 2010). Devitt argues that we ought to, as he puts it, "put metaphysics first," and he believes that once we put metaphysics first, realism will follow inevitably. The first chapter of this dissertation is devoted to dismantling Devitt's doctrine. I conclude that Devitt does not provide sufficient support for his claim that the issue of realism must be understood as a metaphysical question rather than a question in the philosophy of language. Furthermore, I argue that a typical metaphysician looking for a sustainable version of realism in Devitt's philosophy is going to have to take on certain unsavory commitments. The first of these is a Quinean naturalism. Devitt follows Quine in his conviction that the only way we can come to know anything is through the empirical methods of science. The second commitment is a kind of Moorean commitment to common sense. Devitt believes that his naturalism and Mooreanism should lead us to settle metaphysical questions before settling semantic and epistemological

questions. He further believes that his approach makes metaphysical realism the clear winner in the realism debate. For he thinks that realism is supported by all of our normal, experiential evidence, whereas anti-realism is not. I argue that Devitt's conclusion is not warranted by the reasons he offers.

Devitt is reacting to the anti-realist arguments of Putnam and Dummett, among others, which proceed from the assumption that the realism debate is best understood as a debate in the philosophy of language over the nature of truth. My goal in the first chapter is not to prove Devitt absolutely wrong about realism so much as it is to make safe a discussion of realism in terms of language rather than pure metaphysics. I do not conclude that those who prefer to understand realism as a purely metaphysical issue are incorrect. On the contrary, I think that both the metaphysical and linguistic approaches to realism can give us interesting and useful results.

The second chapter is devoted to understanding the nature of mind-independence. Usually, not much is said about mind-independence by those who choose to discuss realism in terms of semantics. This is likely because their approach only indirectly seems to involve the notion of mind-independence inasmuch as they are asking whether what is true depends in any substantive way on what we are able to know, and what we are capable of knowing indisputably depends on what our minds are like. Unfortunately, philosophers who like to construe realism as a metaphysical issue don't say much about the nature of mind-independence either, despite the fact that their whole view is that the realism question is about whether what exists does so mind-independently.

One problem with common statements of what mind-independence amounts to is that such statements seem to rule out realism about the mental. In chapter 2, I argue that it is a

mistake to think that mind-independence precludes the existence of mental entities, and I further argue that the best understanding of mind-independence doesn't even rule out the possibility that everything we normally consider to be part of the so-called external world is actually mental in nature. For instance, I argue that it is consistent with a realist outlook framed in terms of mind-independence that the entire universe might be made up of ideas in God's mind. The reason we might be realist about the world in the relevant sense is that the world still wouldn't be constituted by our own mental activities. I claim that when we speak of mind-independence in the realism debate, the idea we are after is that the world is the way it is independently of what we may happen to think about it. Thus I conclude that the best way to understand mind independence is in terms of independence from our *directed cognitive practices.* I say that something is dependent on our directed cognitive practices if and only if the act of theorizing about, conceptualizing, or otherwise uncovering the nature of that thing is what lends the thing its existence or its nature. This way of understanding things allows someone who prefers to frame realism as the view that what exists does so mind-independently to be a realist about minds and mental things while saving the essence of the realist position. For instance, we might study the nature of someone's mind, which would clearly be mind-dependent in some metaphysical sense, without lending that person's mind its nature via the practice of studying it. One feature of my view is that we may also be realists in the appropriate sense without committing ourselves to a naturalistic view according to which everything in the world is physical rather than mental. Not everyone will agree that that feature is a benefit of my definition, but I think that it is.

Having cleared the way for a semantic discussion of realism and anti-realism in the first chapter, and having clarified the notion of mind-independence in the second chapter, in the third

chapter I move on to a detailed study of one particular semantic approach to the realism debate: The approach of Michael Dummett. Dummett's approach to the realism issue can be seen as the product of his study of three philosophies: Those of Frege, Wittgenstein, and the mathematical intuitionists.

From Frege, Dummett takes the idea that the understanding of a sentence is the understanding of the conditions for its truth. From Wittgenstein, Dummett takes the idea that meaning is use. From the intuitionists, Dummett takes the idea that the notion of truth is essentially constrained by the notion of proof. Putting these together, we get the view that what we understand when we understand what a sentence means is the method for verifying whether the sentence is true. Dummett thought not only that he was extending Frege's insights, but that he was offering a generalizable argument for intuitionistic logic that applied to mathematics as well as to other areas. The argument is supposed to be generalizable because it has to do not with the nature of mathematics but with the nature of meaning generally speaking. There is some irony in the fact that the intuitionists themselves did not think that their logic proceeded from considerations of language. Rather, they believed that it was a consideration of the nature of mathematics itself that should lead us to endorse their philosophy.

Dummett's concern, at least at first pass, is that there are sentences the meanings of which we definitely grasp but for which we cannot specify truth conditions. Hence, he says, we cannot specify any method for verifying whether a given one of these sentences is true. The sentences Dummett has in mind are the undecidables. Dummett thinks that we should understand the realism debate as a debate over whether undecidables have determinate, and in particular bivalent, truth values. Dummett's realist believes that bivalence, the principle that sentences are determinately either true or false, applies to undecidables while the anti-realist

believes that it does not. Because Dummett's version of the realism debate concerns the nature of truth, the two sides are referred to as semantic realism and semantic anti-realism, respectively.

Dummett is typically seen as advancing two distinct but related arguments for his version of anti-realism. One is known as the Manifestation Argument and the other is known as the Acquisition Argument (Cf. Wright 1993, p. 13) I perceive a third argument, which I call the Argument from Endowment, but I also claim that all three of the arguments can be understood in essentially similar ways. Dummett's reasoning is best seen as a reductio of the assumption that our sentences generally (and hence the undecidable sentences) have truth conditions such that they may be either true or false independently of our ability to learn what their truth values are. The idea that truth is such that it may apply to a sentence regardless of whether we can come to know that it does is commonly known as the view that truth is evidence-transcendent, or recognition-transcendent, or verification-transcendent. Given the assumption that undecidable sentences have such truth conditions, Dummett claims that, assuming also the theory that to understand a sentence, and hence to grasp its truth conditions, is to know how to use the sentence, we would not then be capable of manifesting our understanding of undecidables, learning what undecidables mean, or even giving undecidables the meanings that they are supposed to have.

A major thesis of this dissertation is that Dummett never succeeded in specifying a concept of undecidability acceptable to both the realist and anti-realist that would support his argument for anti-realism. In chapter 3, I painstakingly lay out my reading of Dummett. I also criticize some other prominent interpretations of Dummett and defend his arguments, as I understand them, against some prominent objections. I will not reproduce all of the details here, but I will make note of some major conclusions drawn. One conclusion is that objections to

Dummett from semantic externalism, and in particular from considerations of Twin Earth cases, miss their mark. Another conclusion is that, despite the temptation to understand undecidability as absolute unknowability, we ought instead to understand it the way that it is understood in logic and mathematics: As the lack of an effective decision procedure, i.e., a procedure which will deliver an answer as to the truth value of the sentence in a finite amount of time. Dummett seems to understand undecidability this way at some times and as utter unknowability at other times. Dummett himself, however, along with many others who have written on his philosophy, recognizes that an intuitionist cannot claim that there is any sentence P such that we cannot know P and we cannot know the negation of P. For, given the intuitionistic semantics in terms of provability, and the intuitionistic concept of negation according to which one proves ~P by proving that P is impossible to prove, the claim that we cannot prove P and cannot prove ~P amounts to the claim that ~P and ~ ~P, which is a contradiction. Acknowledgment of this fact has become common, and can even be found in Dummett, though, I submit, it is typically not taken to heart.

Chapter 4 deals with the following question: Is the semantic anti-realist, according to whom truth is evidentially constrained, under special pressure to endorse the intuitionistic logic which rejects the law of excluded middle and the rule of double negation elimination? This question might seem odd, and I believe that it is odd. The literature surrounding it seems largely to be the result of Crispin Wright's work following up on Dummett's major idea that intuitionism might be extended from the philosophy of mathematics to other areas. Again, I will not reproduce all of the details of chapter 4 here, but I will highlight some major conclusions. One conclusion is that some attempts, in particular those of Neil Tennant, Joe Salerno, and Jon Cogburn, to discern in Dummett an argument for why the semantic anti-realist must reject

classical logic in favor of intuitionistic logic do not work. I further conclude that an argument the essence of which can be found in both Wright and Salerno does work, but that the motivating force behind the argument's assumptions is ultimately metaphysical in nature.

The major point of the chapter is to unravel the various notions of undecidability which get thrown around in attempts to describe and defend a Dummettian generalized intuitionism. By the chapter's end, I offer three candidates for the concept of undecidability which is supposed to be relevant to the generalization of intuitionism. They are:

A. We cannot know P and we cannot know ~P.

B. We do not know that we can know P and we do not know that we can know ~P.

C. We do not have an effective procedure for determining whether P is true or ~P is true.

A significant portion of chapter 3 was devoted to rejecting A since it is unacceptable to an intuitionist, while the concept of undecidability at work in Dummett's argument must be one that an intuitionist can accept since the debate between the realist and the anti-realist is supposed to be over how to treat undecidables. B has been advanced by numerous parties in an effort to avoid the difficulties associated with A. However, I argue, B cannot do the work Dummett needs it to do in his argument. C, while it is in fact the standard notion of undecidability borrowed from logic and mathematics, also fails to motivate Dummett's overall conclusions. Hence, I claim, Dummett never specified a concept of undecidability that was acceptable to both the realist and the anti-realist while supporting his arguments for anti-realism. A generalized version of intuitionism might still be found acceptable, but consideration of the best versions of such a view on the table shows that what matters is not facts about meaning but rather prior facts about the metaphysical nature of the subject matter at hand. Such a statement might appear to be in contradiction with the stated aims of this dissertation. On the contrary, nothing has been said

to warrant the complete dismissal of Dummett's approach to the realism issue. What was found wanting was his attempt to establish semantic anti-realism, not his basic construal of the positions. In the first chapter, on Devitt, I advance the view that we shouldn't put any area of philosophy before any other when it comes to deciding the realism debate. In the final chapter, I conclude that considerations in both the philosophy of language and metaphysics should be taken into account, neither coming before the other.

Chapter 1: Reconsidering Devitt on Realism and Truth

1. A Short Introduction to Devitt's Views on Realism and Truth

On the face of it, there seems to be a strong connection between the issues of realism and truth. Not surprisingly, then, the nature of truth is often taken to be of paramount importance in settling the question of metaphysical realism. During the last half-century, a number of prominent philosophers have made theories of truth the centerpieces of their arguments against metaphysical realism (Dummett 1978, 1991; Putnam 1978, 1981; Wright 1992, 1993). Yet Michael Devitt, perhaps the most stalwart contemporary advocate of metaphysical realism in the face of anti-realist arguments of all sorts, tells us that realism has a kind of immunity from considerations concerning the nature of truth. My aim in this paper is to show that, as ingenious as Devitt's strategy may be, his approach to the realism debate is not without its problems, and in fact requires the acceptance of a handful of contentious philosophical theses which most metaphysical realists will want to avoid. Indeed, neither the realist nor the anti-realist is under any particular pressure to accept Devitt's contentious theses. Thus freed from the burdens of Devitt's philosophy, we should not be afraid to allow a conclusion in the debate over the nature of truth to lead us to any particular conclusion in the debate over metaphysical realism.

Devitt has erected two distinct lines of defense in an effort to protect realism from truth. The first line involves insisting that realism is a metaphysical thesis, rather than a semantic one (1997, pp. 3, 40). If Devitt is right about this, then settling on a theory of truth would not immediately show realism to be false by definition—some sort of further inference would be required. His second line of defense involves conservatively circumscribing the relevance of the truth debate to the realism debate even when some theory of truth is merely being used as

evidence for some realism thesis. In this way Devitt intends to limit the inferences we might reasonably make from a particular theory of truth to metaphysical anti-realism.

Devitt offers his own characterization of realism as a metaphysical thesis. Devitt's realism, which he calls "Realism," (with a capital 'R'), is stated as follows:

> **Realism:** Tokens of most current common-sense and scientific physical types objectively exist independently of the mental. (1997, p. 303)

As he explains, he includes the word 'most' to indicate that he is not committed to all of our current theories being one-hundred percent true, though he thinks "we are *more or less right* in the physical entities we posit" (p. 303). Devitt, of course, has intentionally worded his version of realism so that it is metaphysical in nature (because it is chiefly about existence) and does not, at least not on the surface of it, involve any particular notion of truth. This fact should not distract us from the fact that he takes the actual issue of realism, by any name, to be a metaphysical issue. In what follows, I will often quote Devitt using the term 'Realism' because he is talking about his particular view. The points made about the relationship between realism and truth, however, should be understood to apply generally to metaphysical characterizations of realism, and not just to Devitt's statement of his own view.

While Devitt does not deny that there are interesting connections between the debate over the nature of truth and the debate over realism, he only admits the existence of two possible routes from a theory of truth to metaphysical anti-realism. One of these routes goes from "relativism about truth" to "relativism about reality;" but Devitt sees no good reason to endorse truth-relativism, and so he maintains that realism is not threatened on this score, saying that the connection "is not very significant." (1997, p. 46) The other route from truth to anti-realism that Devitt considers involves an inference to the best explanation, and it goes as follows: Assume

that an epistemic theory of truth is correct and that all instances of the equivalence schema ('S' is true if and only if S) obtain. The best explanation for why these instances hold, given the epistemic theory of truth, would be that metaphysical anti-realism holds. Hence, anti-realism holds. (1997, pp. 44-45) Devitt says that "this abduction is the only significant respect in which a doctrine of truth is relevant even to the assessment of *Realism.*" (p. 95). Yet Devitt does not think that this is an abduction we should actually make, because he does not think that we should decide on any theory of truth until we have first settled the question of realism (see below).

The opposition to Devitt has mainly involved attempts to refute his claim that the correct characterization of realism is metaphysical, not semantical (Putnam, 2012. Appiah, 1991. Taylor, 2006). Little effort has been focused on undoing the heavy restrictions Devitt places on the ways in which the nature of truth may be important to the assessment of realism, even if it is not important to its definition. His major reason for thinking that our theory of truth should not influence our opinion on realism is that he thinks that we should "put metaphysics first," that is, we should settle the metaphysical question of realism before we start answering questions about semantics, which includes, for Devitt, questions about the nature of truth. This approach to the realism issue gives Devitt an easy response to anyone who claims that her theory of truth shows his realism to be false: That can't be, because we need to settle the realism question first and any plausible semantic theory will be built on the answer to that question as a partial basis.

Before critiquing Devitt's approach to the realism issue, we will need to unpack his notion of "putting metaphysics first" and get clear on his reasons for endorsing it. For Devitt, to "put metaphysics first" is to "give a certain temporal and explanatory priority to metaphysical concerns." (1998, p. 499). Devitt's metaphysics is a naturalistic one; that is, he takes the evidence for it to come from our empirical observations. For him, naturalism is "the view that

there is only one way of knowing, the empirical way that is the basis of science" (2010, p. 110. Cf. 2010, p. 254). He is very clear about this, emphasizing that "knowledge can be *justified* only by experience…the *evidence* for it must be experiential" (2010, p. 254). Devitt is also clear that he is a naturalist in two senses: He is a metaphysical naturalist, one who endorses materialism/physicalism, and he is an epistemological naturalist, one who, like Quine, claims that the only way to know anything is empirically (2010, pp. 254-255).

Drawing on the work of Quine, Devitt says that "everything in the web [of knowledge/belief] can make a difference to everything else," (Devitt 1991, p. 75). This Quinean holism commits him to the view that, in principle, "evidence on the metaphysical issue of realism could come from anywhere, including from semantics," though he includes the important qualification that, "the empirical case for realism is much stronger than that for any epistemological or semantic thesis." (2010, p. 111n). Recall that Devitt takes realism to be a metaphysical issue. He thinks that we should settle the realism question first since, "we know far more about the world than we do about meanings," (2010, p. 163n) and, "…it is much easier to argue for realism than for any semantic doctrine." (1997, p. viii). Accordingly, he says, "the support for [realist] metaphysics outside semantics is so strong that it is scarcely conceivable that a semantic theory should overturn it" (1998, p. 500).

Devitt's attitude toward the realism issue also has a Moorean component. His naturalism leads him to metaphysical realism via experiential evidence, and his Mooreanism secures realism as part of common sense. As he says,

> *Realism* about ordinary objects is confirmed day by day in our experience. It is central to our whole way of viewing the world, the very core of common sense…*Realism* is much

> more firmly based than the epistemological theses…that are thought to undermine it. (2010, p. 62)

Devitt's focus in the passage just quoted is about threats to metaphysical realism from epistemology, but his remarks elsewhere make it clear that he takes the same attitude toward semantics (cf. the previous paragraph). And of course, given his naturalism, "the only way of knowing *anything* [including semantics] is the empirical way of science." (2010 p. 64) The general strategy here is to establish a realist metaphysics first, through empirical methods, and then, in the spirit of G.E. Moore, deny that philosophical conclusions reached in other areas could pose any conceivable threat. Devitt repeatedly harps on the point, as when he says, "Given this strong [empirical] case for *Realism*, we should give it up only in the face of powerful arguments against it and for an alternative" (2010, p. 104). This strategy precludes arguing from a theory of truth to any sort of conclusion in the realism debate.

I have likened Devitt's strategy to Moore's, as Devitt himself does (2010, p. 62). Regardless of whether the approach is truly Moorean, we can say the following about it: Devitt attempts to push the burden of proof onto the anti-realist by claiming 1) That we can be more certain of metaphysics than we can of either semantics or epistemology and 2) Realism has already been, and continues to be, established by our experiential evidence and considerations of common sense.

Devitt evokes Quine's use of an image of Neurath's, of a ship which the sailors must rebuild while staying afloat on it. The metaphor is intended to convey the idea that our theories are all revisable though they are not all revisable at once (2010, p. 110). Devitt claims that we should not rebuild our metaphysics while standing on the semantic part of the boat because, "epistemology and semantics are among the weakest places to stand." (p. 111). Devitt is

completely insistent that we put metaphysics before semantics and epistemology; he repeatedly insists that the latter two areas of inquiry are much more questionable and less secure from an empirical, naturalistic point of view. As he puts it,

> The argument for *Realism,* independent of semantics, is very strong. The argument for verificationism, independent of metaphysics, is very weak… I take the theory of language to be an empirical, conjectural, theory like all others. So there is no question of giving semantics an unearned privileged position in deciding what there is and what it is like. (2012, p. 109)

So Devitt insists that we begin with metaphysics. He also insists that when,

> we go on to seek empirical answers [to the questions of epistemology and semantics]… the theories that result have no special status. Indeed, given our lack of confidence in these areas, the theories should have rather a lowly status. To suppose that we can derive the right metaphysics from epistemology or semantics is to put the cart before the horse.

It is important to see that if Devitt's naturalism and Mooreanism are mistaken, the door is left wide open for a semantic attack on realism. Establishing that realism is a metaphysical thesis is not sufficient to deflect arguments against it from other areas of philosophy, including semantics and epistemology. Devitt needs his naturalism and Mooreanism in order to protect realism. He needs to be correct about the priority of metaphysics and the way that we learn about metaphysical facts.

We have focused on Devitt's doctrine of putting metaphysics first in the context of his argument for his version of metaphysical realism. However, the idea that we should put metaphysics first on its own threatens any semantic construal of the realism debate as a debate over the nature of truth. Such approaches are properly considered semantic despite the fact that

the question as to what truth is is properly viewed as a metaphysical question. Devitt himself argues that the nature of truth is a metaphysical issue (2010, ch. 8). When the realism debate is pursued as a debate about whether truth is somehow epistemically constrained, we are concerned with what sorts of truth conditions our sentences have. In this sense, the approach is semantic in character. Hence Devitt's resistance to the idea that the realism dispute is about whether truth is epistemic. To think of the dispute that way is to give priority to semantics.

I have attempted to cull from Devitt's writings a substantial reconstruction of his philosophy of putting metaphysics first in relation to his Realism. It should be kept in mind that his Realism is itself a piece of metaphysics. Devitt's epistemological naturalism and his Moorean commitment to common sense are supposed to support not just Realism, but the putting of metaphysics first generally. On my reading of Devitt, the Mooreanism and epistemological naturalism lead to the priority of metaphysics. It just so happens that Devitt additionally thinks that they support a realist metaphysic in particular.

Recall that Devitt puts metaphysics first both temporally and explanatorily. It often seems that he simply assumes the priority of metaphysics and works from there. For instance, things seem that way when Devitt lays out his Maxims in the first part of his *Realism and Truth* (1997). Closer analysis of his writings shows, however, that he does attempt to provide reasons for why we should put metaphysics first. Devitt appears to believe that his epistemological naturalism is sufficient to establish that we, as a matter of fact, do our metaphysics first. The idea is that empirical methods tell us first what is in the world and some of what it is like. Only later do we make progress at the more difficult tasks of analyzing language and epistemology. The Mooreanism supports the explanatory priority of metaphysics—Devitt thinks that, once we

have established what there is, we should lean on our metaphysics as we do semantics and epistemology rather than changing our metaphysics in light of semantics and epistemology.

One might try to understand Devitt in a different way. One might think that Devitt doesn't use his Mooreanism and epistemological naturalism as reasons for why we should give "temporal and explanatory priority" to metaphysics. Rather, this line of reasoning might go, Devitt simply puts metaphysics first and only then invokes Mooreanism and epistemological naturalism as he argues for his Realism, which just so happens to be a metaphysical thesis. The difficulty with reading Devitt in that way is that it leaves him practically if not entirely without an argument for a position which is supposed to reveal as wrong-headed so much of the work done in the analytic tradition since its inception.

Here is perhaps Devitt's most concise statement of his view about the priority of metaphysics:

> We should approach epistemology and semantics from a metaphysical perspective rather than vice versa. We should do this because we know much more about the way the world is than we do about how we know about, or refer to, that world…my view here reflects, of course, my epistemological naturalism. The metaphysics I want to put first is a naturalized one. (2010, p. 2)

In the passage just quoted, Devitt is clearly attempting to give us a reason for why we should put metaphysics first. The reader should recognize the theme concerning how much more we know about metaphysics than we do about epistemology and semantics from the various bits of Devitt I quoted above as I attempted to reconstruct his views. Part of the challenge has been to explain just why Devitt thinks that we know more about metaphysics than we do about other areas. I

have attempted to give an account based on Devitt's Mooreanism and epistemological naturalism, the latter of which he explicitly mentions in the most recent quotation.

In the section 2 I take issue with Devitt's epistemological naturalism, his Mooreanism, and the idea that those doctrines should lead us to put metaphysics, as opposed to something else or nothing at all, first.

2. Problems with Devitt's Approach

There are reasons to be worried about Devitt's approach to the realism issue. To begin with, his naturalism and Mooreanism are themselves highly contentious doctrines. This is worth noting because it makes clear the fact that, for the realist who might seek solace in Devitt's philosophy, Devitt's way to realism requires a lot of philosophical baggage which many realists will not be willing to carry. To be sure, this does not mean that Devitt is wrong about the priority of metaphysics. It does, however, mean that the run of the mill metaphysical realist should be wary of adopting Devitt's attitude toward the realism debate.

One issue is that many analytic metaphysicians take themselves to be engaging in an *a priori* metaphysics which is inconsistent with an epistemological naturalism like Devitt's. Thus for instance Trenton Merricks, in the preface to his *Objects and Persons*, tells us that "ontological discovery is not empirical." (2003, p. vii) Furthermore, even a metaphysician who embraced a naturalistic methodology would likely take issue with Devitt's claim that metaphysics is an easy, common-sense affair. Surely nothing so easy would be debated so intensely and at such length as metaphysics is. The idea that we learn about metaphysics through naturalistic methods and the idea that metaphysics is easier than semantics are both important aspects of Devitt's rationale for putting metaphysics first. And they would likely both be rejected by a majority of practicing metaphysical realists.

A realist about mathematics who wanted to go along with Devitt's idea that realism is a metaphysical thesis might defend the following similar view:

Tokens of mathematical types objectively exist independently of the mental.

While there are certainly naturalistic accounts of mathematics on offer, the field remains a place where naturalism is largely seen as unattractive because it seems to so many that we must come to know about mathematical objects via *a priori* methods. A mathematical realist of this sort would be unable to accept Devitt's naturalism. Hence, such a realist would not be able to accept Devitt's rationale for prioritizing metaphysics over semantics. Granted, this is not a problem for Devitt, but it does speak to the limitations of adopting his attitude toward the debates.

Another issue is, of course, that many metaphysical realists will not want to give so much priority to common sense; they will reject the Moorean component of Devitt's philosophy. Indeed, much of the work done by realist metaphysicians seeks to go beyond common sense and often ends up overthrowing our most basic common-sense convictions about the way the world is. These metaphysicians are realists in the sense that they believe they can, in principle, uncover deep truths about the nature of reality. And yet they have no particular reason to value a common sense ontology the way that Devitt does. One might think that these realists could just adopt Devitt's arguments for prioritizing metaphysics minus the commitment to common sense. But the commitment to common sense is actually crucial to Devitt's argument for why we should put metaphysics first. For imagine that he simply said that we should do metaphysics first, even if the resulting metaphysic ended up conflicting with common sense. He would have no good reason for then saying that we should hold on to the resulting metaphysical theory no matter what develops in semantics. For he could not then appeal to the authority of common sense.

Devitt's naturalism alone cannot push us to prioritize metaphysics, since his naturalism applies equally to metaphysics, semantics, and epistemology. As attractive as Devitt's approach to the realism debate may sound to many realists, most of them will be unable to follow his path to realism. Devitt's strategy is not easily transferrable to philosophical views other than his own.

To be fair, Devitt does not claim that every aspect of our common sense view of the world is immune from revision.[1] What Devitt does claim is that the basic tenets of his Realism as applied to ordinary objects are "central to our whole way of viewing the world;" they are part of the "core of common sense" (1997, p. 60). What Devitt has in mind are the views that everyday, ordinary, observable objects not only exist, but do so mind-independently. The idea that they might, for instance, exist mind-dependently is incredible to Devitt. So, of course, he thinks that the onus is on the anti-realist to explain how this could be the case. In the Quinean terms that Devitt prefers, Realism is closer to the center of the web of belief than semantics and epistemology, and therefore less susceptible to revision. So, one might say that Devitt's strategy of putting metaphysics first is really the strategy of putting metaphysics, and in particular *realist* metaphysics, closer to the center of the web.

Devitt has, in essence, told us that a realist metaphysics, according to which our everyday and scientific ontologies are populated with mind-independent entities, is much more adequately supported than any view in semantics which might be thought to overturn it. He repeatedly tells us that this support is based in our everyday experience of the objects which we admit into our ontology. That these objects exist and do so mind-independently is simply part of the view of the world which we have inherited from our earliest experiences of it. So Devitt puts his realist metaphysics near the center of the web. And so his realist metaphysics is more resistant to

[1] Thanks are due to an anonymous referee for pushing me on this point.

revision than, say, the semantical views which would fall closer to the outside. One wonders how Devitt can claim that epistemology and semantics are weak starting places while telling us that his naturalism, from which he derives his metaphysics, is itself a piece of epistemology (2010, p. 64, p. 77n22). And in fact his naturalism is in part an epistemological doctrine, as it has to do with the kinds of justification we can have for our beliefs. So it seems that in arguing for the view that we should put metaphysics first, Devitt has put epistemology first. Furthermore, Devitt claims that the burden of proof is on the anti-realist, but Devitt's approach to the realism issue seems to rule out the possibility that the anti-realist might ever be able to satisfy the burden. For, given how Devitt frames things, any semantical thesis that might overturn the realist metaphysic will be at odds with the common-sense view that the objects in our everyday ontology exist mind-independently. Devitt wants us to build our semantics on our metaphysics. If this metaphysics must be realist, as Devitt insists it must, given how we establish the metaphysics to begin with through experience, then no semantics which is inconsistent with it could ever obtain the requisite amount of plausibility. The worry here is that Devitt's argument establishes too much. If metaphysical realism becomes essentially impossible to overturn, this should make us question Devitt's approach.

Consider Devitt's Maxim 3: "Maxim 3: Settle the realism issue before any epistemic or semantic issue" (1997, p. 4). Devitt thinks we should accept Maxim 3 even though he offers a caveat which, I contend, leaves room for anti-realist rebuke. Here is the caveat:

> This maxim is oversimplified because realism, though largely metaphysical, *is* a little bit epistemic and semantic: the world must be independent of our knowledge of it and of our capacity to refer to it. So at least that much epistemology and semantics must be settled to settle realism. (1997, p. 4)

Recall how Devitt thinks that a realist metaphysic gets established: Via experience. But do our experiences tell us that what exists does so mind-independently? What would that experience be like as opposed to an experience that those objects have an existence which is at least partially mind-dependent? I contend that our experience is more or less neutral on this point, and that the question as to whether the denizens of our ontology exist mind-independently is a further question requiring further theorizing. So, I claim, metaphysical realism doesn't come along with common sense. The view that tables and chairs exist might be common-sensical, but any particular view about whether they do or do not depend for their existence on anything else does not. At best, perhaps experience does tell us that objects have objective existence, where that is understood to mean that their existence is not subjective—these objects do not exist for only one person, or anything like that. But to determine whether the objects of our experience have their existence tainted with some degree of dependence on our minds as investigators requires more than mere empirical experience.

It may be the case that putting semantics first is more in line with Devitt's framework than he is willing to admit. After all, Devitt wants to naturalize everything, including semantics and epistemology. If we see the study of language as an empirical endeavor, then putting semantics before metaphysics would seem to be consistent with Devitt's naturalistic approach. But he will resist prioritizing language, taking the Moorean line and saying that the realist metaphysic is simply part of good common sense. He prioritizes the everyday ontology of realism about ordinary objects on the basis that we first learn about such things early on in life through observation. It is worth noting, then, that we also learn about the meanings of words empirically and early on. At the very least, Devitt's naturalism commits him to as much, since it requires that the only way we can learn about anything is empirically.

We should wonder whether there is any good reason to prioritize common-sense metaphysics over common-sense semantics, or common-sense anything, for that matter. In short, Devitt has given us no good reason to think that metaphysics is prior to other kinds of inquiry with regard to common sense. In fact, one wonders why he would focus so much on the importance of a common-sense metaphysic while saying that semantics and epistemology are more difficult and that we cannot be as confident about what we learn in those areas. It should come as no surprise that Devitt thinks metaphysics is easier—he has made metaphysics commonsensical and semantics and epistemology highly theoretical.

Devitt may resist my criticism by claiming that there is no such thing as common-sense semantics and epistemology, or that common sense has no role to play in those areas. As for the former claim, it seems patently false. Surely there are commonsensical, perhaps even naïve, ideas pertaining to semantics and epistemology. Even the simplest thought about what words mean or what we are justified in believing could count. As for the latter claim, Devitt says nothing to justify the idea that common sense matters in metaphysics but neither in semantics nor epistemology.

So why should we put anything first at all? I've given reason to doubt that metaphysics should be prioritized the way that Devitt prioritizes it. In reality, we needn't prioritize any particular area over any other. And arguably, we shouldn't do so, since it would be very strange if the various areas of human inquiry settled so neatly together that we could, for instance, answer all of the metaphysical questions before we even began asking the semantical questions. Semantical considerations can be taken as evidence for metaphysical conclusions, unless we simply add Devitt's restriction that we should put metaphysics first. If we follow Devitt's line of reasoning, we are giving up hope of answering metaphysical questions through semantics

entirely, and this seems rather rash. Devitt hasn't given us any reason to think that one can't do metaphysics via semantics. Ultimately, he has only really said that we shouldn't, largely because he thinks that we can be much more confident about metaphysics than semantics.

A more nuanced and versatile approach would be more powerful. At the very least, it would be useful to take an approach which allowed us to fill in some of the gaps in our metaphysical knowledge via a consideration of how metaphysical reality and linguistic reality are related. Especially considering the fact that Devitt's approach of putting metaphysics first treats of metaphysical issues too lightly, taking questions about observable entities to be answered by recourse to common sense, we would do well to leave open the possibility of whether a commonsense metaphysic is sustainable and accurate.

In particular I would suggest that it is irresponsible to put off answering questions in semantics and epistemology until we have our metaphysics squared away. We should approach all areas with the same vigor, seeking the best answers that we can in each area which cohere as much as possible with answers in other areas. Such an approach is consistent with even a Devittian epistemological naturalism, since that view only requires that the questions in each domain be answered in a naturalistic way. At the same time, my approach promises to do justice to our everyday experiences that Devitt prizes so highly without miring us in dogmatism.

3. Conclusion

Devitt's take on the realism issue is complex and subtle. Besides offering positive arguments in favor of his Realism, Devitt has gone to great lengths to carefully scrutinize and attack the major anti-realist arguments of his opponents. Presently, I have not aimed to assess the overall merit of Devitt's defense of realism against those specific attacks—I have only been concerned to overthrow the most pervasive element of Devitt's general strategy: The doctrine that we must

put metaphysics first. At the same time, I do not pretend to have proven conclusively that Devitt is *wrong* to put metaphysics first, though a conclusive argument against either his naturalism or his Mooreanism would suffice to show that. In addition to criticizing Devitt's strategy of putting metaphysics first, I have attempted to offer an attractive alternative approach that does not favor any particular area of inquiry over another.

Chapter 2: Realism and Mind-Independence

1. Realism, Objectivity, and Mind-Independence

Realism is one of the most pervasive issues in philosophy. The question of realism comes up in essentially every area of philosophical inquiry, from metaphysics and ontology to semantics, ethics, and the nature of truth. And while very many talented philosophers have spilled much ink attempting to say what realism is, there is yet no general consensus as to what exactly realism is supposed to be. Furthermore, even those who have settled on some general characterization are faced with difficulties specific to their preferred definition of the notion. For instance, the notion of mind-independence is loosely defined at best though it is supposed to do much of the work explaining what realism is from a metaphysical point of view.

Despite the difficulties inherent in explicating realism and its opposite, anti-realism, there is significant agreement regarding the general idea. A realist in some domain will assert something along the following lines: The facts in that domain obtain, or the propositions in that domain are true, or the entities in that domain exist and have the natures they have, regardless of whatever we may happen to think, believe, or assert, and regardless even of what we might be capable of thinking, believing, or asserting. Few would argue that the foregoing description fails to capture the basic idea behind realism. The difficulties arise when we attempt to speak less loosely and in a way more befitting philosophical analysis of the relevant issues at stake.

One attempt to bring rigor to the realism debate is decidedly semantic in character. According to those who prefer a semantic construal of realism, the question of realism comes down to the question of whether our sentences, statements, propositions, and the like have the truth values they have in either an evidentially constrained or evidence-transcendent way. If the statements we make in some area have truth values according to whether we may, in principle,

come to know that they are true, then according to the semantic construal of realism we should take an anti-realist attitude toward the facts in that area. The notion of truth involved in this case would be epistemic, somehow constrained by the sorts of justification we can have for our beliefs and what we can come to know. Accordingly, if the statements we make in an area have their truth values completely independently of our ability to come to know whether they are in fact true or false, then we are working with an evidence-transcendent notion of truth and should be realists about that domain.

Another attempt to analyze the issue of realism focuses on issues in metaphysical reality as opposed to semantics and epistemology. According to this sort of definition of realism, the question comes down to what exists and whether it exists and has its nature independently of both our actual and ideal cognitive activities, with some statements of realism in terms of mind-independence going so far (perhaps inadvertently as a result of careless formulation) as to seemingly banish the mental entirely. The question of what exists is, of course, a difficult one on its own. But it is easier to answer than the most pressing question, which concerns whether what exists does so in a mind-independent way. In order to say whether it does, we would need to know what exactly mind-independence is supposed to be. Yet it is not entirely clear what mind-independence amounts to.

Closely related to the concept of mind-independence are numerous notions of objectivity. In short, something can be said to be the case objectively if it is the case regardless of what anyone happens to think or believe, what sorts of concepts they have at their disposal, etc. Hence it would be objectively true that Canada is north of the United States and Mexico is south of it. When it comes to matters of personal taste, however, there may be no objective facts. Hence many would resist the claim that blackberry pie is objectively delicious since there are

people who do not find blackberry pie delicious. There is clear overlap here with mind-independence—if nothing else, something's being the case in virtue of what someone happens to think would seem to make that thing the case only mind-dependently. A further question is whether Canada's and Mexico's cardinal orientation relative to the U.S. is a mind-independent matter—and this example helps bring out the fact that mind-independence and objectivity do overlap but also come apart from each other. For while we can agree on the objectivity of the facts, notions such as nationhood and even cardinal direction may be crucially dependent on our mental lives.

With these important concepts under our belt, let's take a look at some statements of the realism issue which make explicit use of them. In the late 1970s, William Alston offered this succinct definition of realism:

> Realism is here being understood as the view that whatever there is is what it is regardless of how we think of it. (1979, p. 779)

In this definition, there is no special distinction drawn between objectivity and mind-independence, though both notions seem to be in play. That the world does not depend on whatever anyone happens to think, but is some way regardless, suggests both mind-independence and objectivity, though Alston is not specific about the details. Notice that the focus on the existence and nature of things gives Alston's statement a distinctive metaphysical character, rather than a semantic character. Gideon Rosen provides this much fuller but also more abstract statement of the issue:

> What the realist mainly claims is the right to say things like this: 'Our discourse about X concerns a domain of fact that is *out there.* These facts obtain *anyway*, regardless of what we may think. When all goes well, inquiry into the disputed area *discovers* what is

> *already* there, rather than *constituting* or *constructing* its object...*the target discourse describes a domain of genuine, objective fact.*" (1994, p. 278)

Here too we have both notions. Objectivity is explicitly mentioned while mind-independence is at play when he brings up the idea that we discover what is out there rather than creating it, or perhaps lending it part of its nature, through the practice of inquiring about it. Penelope Maddy gives us the following statement of realism:

> To be a realist about medium-sized physical objects, the theoretical posits of science, or universals, is to hold that these entities exist, that they do so objectively—they are not mental entities, and they have the properties they do independently of our language, concepts, theories, and of our cognitive apparatus in general. (1992, p. 14)

A couple of aspects of Maddy's definition are important to bring out here. For one thing, she discusses both objectivity and mind-independence, rather than just one of those notions. Another point is that the second half of her statement appears to be explicating the notion of objectivity itself. This is interesting since in the latter portion of the definition she mentions both "mental entities" and "our cognitive apparatus in general," thus seemingly defining objectivity in part via the notion of mind-independence. Finally consider the following two statements of Devitt's:

> The general doctrine of realism about the external world is committed not only to the existence of this world but also to its 'mind-independence': it is not made up of 'ideas' or 'sense data' and does not depend for its existence and nature on the cognitive activities and capacities of our minds. (2006, p. 101)
>
> *Realism:* Tokens of most current common-sense and scientific physical types objectively exist independently of the mental. (1997, p. 303)

Note that Maddy and Devitt both associate the notion of mind-independence with a sort of complete independence from anything mental. Thus Devitt tells us that realism precludes the world being "made up of 'ideas'" and Maddy says that realism requires that the entities in question not be "mental entities." This is a major point to which I will return below.

So far my goal in discussing these various metaphysical statements of realism has been to highlight the role played by objectivity and mind-independence. These notions do not belong solely to metaphysical characterizations of realism, however. Part of my goal in this chapter is to establish that the concepts play an implicit if not an explicit role in semantic construals of realism as well. Before moving on to a discussion of how these notions figure in characterizations of realism which are not explicitly metaphysical, I will first spend some time critically analyzing the notion of mind-independence.

2. What is Mind-Independence?

The abstract nature of the notion of mind-independence has not been lost on philosophers. A number of them have attempted to dissect the concept and distinguish various versions of it. Many of the distinctions people have attempted to draw will not be discussed here as they are not very useful. Still, something can be said about mind-independence besides just that realism requires facts in a domain to obtain independently of minds.

Carrie Jenkins has noted a very informative distinction between two possible kinds of mind-independence. The first kind of independence is what she calls *modal independence*, according to which "something is independently the case just in case there is a possible world where that thing is the case although our mental lives are not [a certain way]," where the relevant property of our minds can vary from case to case. (2005, p. 200) Thus, if there is a possible world where dinosaurs exist on some alien planet, even though in that world we lack the ability

to conceive of dinosaurs, then the existence of dinosaurs would be modally independent of our minds.

The second kind of mind-independence that Jenkins discerns is what she calls *essential independence*, according to which "*p*'s being the case is independent of our mental lives iff it is no part of *what it is* for *p* to be the case that our mental lives be a certain way." (2005, p. 200) Thus we have the idea that, for anti-realism in some domain, our mental lives actually constitute the objects of inquiry in some essential way, or at least lend them some of their essential properties. Jenkins speaks of something's "being the case," which comes quite close to talk of a proposition being *true* mind-independently. This ought not to throw us off however, as I have already mentioned that mind-independence and objectivity both play an important role in semantical characterizations of realism. Still, the challenge to say something more about the nature of mind-independence is particularly pressing for those who wish to use the notion to characterize realism in a metaphysical way. This is because such theorists almost always make explicit reference to mind-independence without saying much about it and build their realisms (or anti-realisms) around it.

Accordingly, Devitt and Alston, both of whom endorse metaphysical characterizations of realism, have made use of a notion very close to Jenkins' essential independence. Thus Alston tells us,

> The kind of dependence that is incompatible with a realist status is what we may call *constitutive* dependence. If physical substances, space and time, universals, or whatever, depend on a relation to mind for being what they are, for their essential character, for their constitution, then they lack the kind of independence of mind that is required for

> realist status…what I have just been calling a 'constitutive' dependence of the nonmental on mind…is a reduction of the nonmental to the mental." (1996, p. 73-4)

Devitt essentially concurs:

> I mean *constitutive* independence…the known world is not *constituted* by our knowledge, by our epistemic values, by our capacity to refer to it, by the synthesizing power of the mind, nor by our imposition of concepts, theories, or languages; it is not limited by what we can believe or discover. (2012, pp. 113-114)

It was noted above that many of those who wish to characterize realism metaphysically and make use of the notion of mind-independence also make the further claim that the world itself, or at least the objects in it which we might want to be realists about, is not made up of mental entities in any way. There are a couple of difficulties with this idea. One of them has been brought up by a number of philosophers though not much has been done to solve the problem. This is the problem that insisting that realism requires absolute independence from the mental doesn't seem to leave room for realism about our own minds and our own mental states. (Devitt 1997, p. 14; Sayre-McCord 1988, p. 6). Yet surely if we want to be realists about anything it is our own minds, and as for mental states no acceptable generic formulation of realism should rule them out automatically since realism about mental states is not obviously incoherent.

I am attempting to give a formulation of the notion of mind-independence that allows us to be realists about minds and mental things. Some, such as Devitt and Sayre-McCord (cited above), legitimately raise the worry that casting realism in terms of mind-independence doesn't leave room for realism about minds. Others, I would suggest, are not necessarily concerned about the problem but their own formulations do fail to account for realism about minds and

mental things (perhaps this is the case with Maddy). My point is that the problem is ultimately illusory and avoidable as long as we properly state what mind-independence is. The following passage of Devitt's captures nicely the way of thinking about mind-independence I seek to overturn:

> For the realist the material or physical world he believes in has to exist not only objectively but non-mentally. We can roughly distinguish two aspects of this: (1) the world does not consist in mental objects of experience, neither in 'ideas' as idealists like Berkeley thought, nor in 'sense data' as many phenomenalists thought; (2) the world is not made up of minds, as Leibniz thought, nor of something ultimately spiritual as the absolute idealists thought. Both (1) and (2) are well enough covered by saying that, for the realist, the world exists *independently of the mental.*
>
> There is a minor problem about this characterization of independence. Part of the independent world that the realist believes in is made up of higher animals like people. It may be that the realist will think that an object would not be a person if he lacked a mind. That is the likely view of two common sorts of realist: non-eliminative physicalists and dualists. For such a realist, a person does not exist independently of the mental. The problem is only minor because the defence of a realist that did not cover the higher animals would be sufficient for our purposes. The belief in those animals would then be something additional to realism. I shall not complicate the characterization to take account of the problem. Rather I shall ignore the problem, assuming that realism does cover these animals. (1997, pp. 13-14)

I have quoted Devitt at length in order to forestall the thought that no one seriously thinks that the casting of realism in terms of mind-independence threatens to deprive us of the resources of

being realists about minds and mental things. Devitt says that he is going to ignore the problem. I aim to solve it for him.

Another problem with requiring realism in some domain to exclude the possibility of entities in that domain being somehow mental in nature is that this seems to just equate realism with a kind of naturalism, or at least require a realist to be a kind of naturalist. Yet surely there are perfectly sensible realists who reject any kind of naturalism about the area in question. Thus a realist about universals might claim that universals are ideal ideas in some sense, ideas which we can latch onto with our minds. This person would be a realist in thinking that universals exist mind-independently in the appropriate sense—they do not depend for their existence on our minds or our mental activities[2]. Still, these universals are typically taken to be ideas, perhaps ideas in the mind of God, which are not suited to naturalistic investigation and are not supposed to be part of the natural world. To say that any philosopher who rejects naturalism is automatically an anti-realist about universals is misleading at best.

Furthermore, not enough has been done to support the claim that realism precludes what we call the "external world" being constituted by mental entities of any kind. For consider that we were to one day discover that all of the entities we took to be physicalistic in nature were actually the mental products of an all-powerful divine being. Imagine further that this makes no massive difference to the practice of science as we know it, since we can just carry on studying phenomena in the universe as we have been doing. It would be quite odd if we were forced to say that we must be *anti-realists* about the external world in a situation like this. Quite the contrary, we should be realists about it since it is not dependent on *our own mental lives.* Here "our" should be understood to include any beings with minds sufficiently like the minds of

[2] Special thanks are due to David Braun for helping me to clarify this point not only on paper but also in my own mind.

human beings on Earth. Different sorts of beings may be realists about things we have no business being realists about.

The idea that realism requires independence from our own mental lives as inquirers, as opposed to any mental lives whatsoever, promises to dissolve not only the issue stated directly above concerning naturalistic physicalism, but also the issue regarding how to be realists about minds and mental states. For it allows us to admit minds and mental states into our ontology without making such things unreal in virtue of depending on our directed cognitive practices for their existence. Let us say that something depends on our *directed cognitive practices* if and only if the act of theorizing about, conceptualizing, or otherwise uncovering the nature of that thing is what lends the thing its existence or its nature. Consider for instance now a single person's mind. That person's mind is a mental thing. And yet if we go to uncover some facts about his mind, to determine its nature, we would not thereby be bringing his mind into existence. Nor would we necessarily be constituting with our own minds, through any cognitive act directed at his mind, the properties we took his mind to have. His mind would, on a realist view, have those properties anyway, regardless of how we thought of them or even if we ever did.

Contrast the sort of case just described with the following kind of case. On some anti-realist views, the world only is some way or another relative to a description or conceptual scheme. (Putnam 1983). Such anti-realisms say that there are numerous mutually inconsistent ways of describing or conceiving of the world, none of which necessarily gets at the way the world really is. The anti-realistic aspect of these theories derives from the fact that our directed cognitive practices, through which we seek to conceptualize the world around us, determine which facts about the world obtain. Thus, on this sort of a view, we can conceive of the world as

including the unobservables of science, but we will have to admit that the world doesn't *really* include them since the facts about unobservables will be dependent on the fact that we actually do divide the world up so as to include them.

I therefore maintain that we ought not think of the sort of mind-independence relevant to realism as one which banishes the mental altogether. A more useful version of mind-independence is one which makes something dependent on our minds when it depends on our directed cognitive practices for either its existence or its nature. This conception of mind-independence keeps the useful aspects of Jenkins', Devitt's, and Alston's constitutive independence while avoiding the pitfalls associated with an insistence on what essentially amounts to a form of physicalism.

To say that something depends for its existence or nature on our *actual* cognitive practices is certainly anti-realist. But this way of looking at things is rather restricted, as a perfectly sensible anti-realist may not be comfortable with the thought that what (anti-realistically) exists changes as our investigative practices change. And the very same anti-realist might wish to admit that there could be aspects of this universe which we will never investigate, and yet which, if we were to investigate them, would depend upon our minds in the way just explained. Thus we should say that it is not our actual cognitive practices but the cognitive practices we could conceivably engage in, given our nature, which matters for realism.[3]

[3] Thanks to David Braun for pointing out the importance of this distinction to me.

3. Semantic Characterizations of Realism and the Relevance of Mind-Independence to the Realism Debate Generally

Consider the following passage from Dummett, proclaiming realism to be semantical in nature:

> Realism is a *semantic* thesis, a thesis about what, in general, renders a statement in the given class true when it is true. The very minimum that realism can be held to involve is that statements in the given class relate to some reality that exists independently of our knowledge of it, in such a way that that reality renders each statement in the class determinately true or false, again independently of whether we know, or are even able to discover, its truth-value. (1982, p. 55)

This passage highlights Dummett's use of a notion of objectivity in his statement of realism. Notice that he does not say anything explicitly about mind-independence. In fact, Dummett and Putnam have both claimed that realist notions of mind-independence merely trade in metaphor. Thus Dummett says that, in the debate over realism (platonism) in mathematics, "The platonist metaphor assimilates mathematical enquiry to the investigations of the astronomer…[while] the constructivist metaphor assimilates mathematical activity to that of the artificer fashioning objects in accordance with the creative power of his imagination." (1983, p. 111) Putnam and Dummett see their semantic construals of realism as giving workable content to the notion of mind-independence, posing answerable questions where before we worked only with realist imagery. (Putnam 2012, p. 124).

Putnam, Dummett, and those who work in their vein do believe that realism issues are not only important but tractable. This sort of optimism is not to be found in the writings of Gideon Rosen, who, frustrated with what he sees as the impossibility of making sense of mind-independence in the context of a 20th century worldview, gives up on realism issues altogether:

> We *sense* that there is a heady metaphysical thesis at stake in these debates over realism… but after a point, when every attempt to say just what the issue is has come up empty, we have no real choice but to conclude that despite all the wonderful, suggestive imagery, there is ultimately nothing in the neighborhood to discuss." (1994, p. 279)

I do not concur with Rosen. Rather, I agree partially with Dummett and Putnam that one way of cashing out the abstract notion of mind-independence is to speak directly about properties of truth. Indeed, I wish to make the further claim that the way of Putnam and Dummett is only one of many ways we might give more precise meaning to the term 'mind-independent'. Much of the difficulty surrounding the use of the term, I contend, is rooted in the fact that it often goes underspecified.

There are many ways, as evidenced by the passages quoted in section 1, that the facts in some area of inquiry might fail to be mind-independent. The concept of mind itself is not entirely clear, nor has it been precisely specified. So it is of little use presently to look for a general, all-encompassing notion of mind-independence. We would do better to settle for particular identifiable instances. For even if we have to admit that there are hard cases, there is relative consensus that some things are rather clearly mind-independent and others are not. (Cf. Asay 2012, p. 380). Things are much the same in the realism debate generally. Again as evidenced by the passages quoted in the first section, attempts to state precisely what realism is almost always make use of some disjunction of possible views which would amount to realism in different ways. Thus one might be an anti-realist in an area because one believes that the facts in that area are dependent on our linguistic practices. We can talk of linguistic practices without speaking of generic mind-independence and yet we can recognize that facts which rely on such practices are somehow mind-dependent. We can also talk of our concepts and recognize that

anything which relies on them for its existence or nature is somehow mind-dependent. Similarly with semantic (or alethic) construals of realism and anti-realism. We can say, with Dummett, that facts in an area are mind-dependent because statements of those facts have their truth values in an evidentially-constrained way.

In all of these cases we have not had to lean on a generic notion of realism, nor have we had to lean on a generic notion of mind-independence. Rather we have isolated particular ways that facts in an area might be mind-independent or fail to be. Accordingly we should call ourselves realist or anti-realist with regard to those areas as we see fit. But we need not labor to define precisely every possible way that one can be a realist or an anti-realist. What matters is that in each case we have isolated that feature which, though things might not be the same in other areas, would suffice for realism, where the central issue of realism is mind-independence, and where mind-independence is understood in the way I have argued it ought to be. If I am correct about mind-independence, then Putnam and Dummett go too far when they attempt to *replace* that notion with a semantic notion about truth. Their semantic approaches centered on the relationship between truth and evidence simply spell out specific instances of mind-dependence and mind-independence based on what we could in principle come to know.

My account, while importantly different from Geoffrey Sayre-McCord's treatment of the realism issue, shares at least some of its motivation. According to Sayre-McCord, all that is required for realism is that

> 1) the claims in question, when literally construed, are literally true or false…and 2) some are literally true. (1988, p. 5).

Sayre-McCord claims his account has a number of virtues. Among these are the facts that he makes "no mention of objectivity," and that he avoids issues concerning being a realist about the

mind which stem from requiring complete "independence from the mental." (1988, p. 6). I have offered my own account of how to understand mind-independence in a way that avoids what is often seen as the insurmountable difficulty of allowing realism about minds. I have not tried to avoid the notion of objectivity, but as long as the way in which something is objective or fails to be is specified, the notion seems at least as workable as the notion of mind-independence. Sayre-McCord also takes it to be a virtue of his account that it allows different factors to make the difference in different areas depending on what in particular is at stake in each area. (1988, p. 6). I have endorsed a similar attitude toward the realism issue, in particular with regard to mind-independence. It is by specifying the ways, within a specific debate, that mind-independence is at issue that the notion becomes useful. Before such specificity, it is no more useful than an image. Furthermore, while Sayre-McCord's criteria may be necessary for realism, they do not seem to be sufficient. For instance, anyone working with an epistemic notion of truth should be classified as an anti-realist, but would still claim that things can be literally true or false, and that some things are literally (epistemically) true.[4] So Sayre-McCord's definition fails to capture some clear cases of anti-realism. My construal in terms of mind-independence properly understood avoids this shortcoming while helping itself to the benefits of his construal.

Something should be said about the notion in meta-ethics of "stance-independence" (Cf. Shafer-Landau 2003, p. 15).[5] Shafer-Landau has argued that the notion of mind-independence is of no use in characterizing realism in the ethical domain. Rather, he says, we should prefer the notion of *stance-independence.* Cashed out in these terms, realism in ethics amounts to the view that,

---

[4] Thanks to David Braun for pointing this out to me.

[5] Thanks are due to Brendan Cline and Alex King for pointing out the importance of this issue as well as for their helpful discussion.

> There are moral truths that obtain independently of any preferred perspective, in the sense that *the moral standards that fix the moral facts are not made true by virtue of their ratification from within any given actual or hypothetical perspective.* That a person takes a particular attitude toward a putative moral standard is not what makes that standard correct. (Shafer-Landau 2003, p. 15)

I do not wish to take issue with Shafer-Landau's characterization of moral realism. Rather, I only wish to argue that it is perfectly consistent with my account of how we ought to understand mind-independence and its relation to the realism debate generally, despite the fact that Shafer-Landau intends his stance-independence to replace mind-independence.

I have in one way offered a softening of the notion of mind-independence, since I have argued that mind-independence need not require a complete independence from the mental. I have also argued for the importance of saying precisely relative to what something is supposed to be mind-independent. Shafer-Landau's reasons for rejecting mind-independence as a useful notion in meta-ethics stem from his observation that,

> any plausible moral theory will claim that some moral truths depend crucially on an agent's mental states… and the moral status of an action may depend very importantly on how pleased or miserable it makes others, whether it prompts feelings of anger or empathy…etc. (2003, p. 15)

He also claims that, without minds, "there would be no moral facts (as opposed to moral standards)," though he qualifies this statement to mean that there would be "no particular instances of right and wrong," though there could still be "moral principles." (p. 15, esp. p. 15n2) So Shafer-Landau is not claiming that the existence of moral reality requires the actual existence of minds. Furthermore, I take it that most realists will be more interested in the real existence of

moral standards than they will be with any particular instances of wrongdoing, so I will not enter into a discussion of the relevance of mind-independence to the view that there are in fact instances of people doing right or doing wrong. Instead I will focus on Shafer-Landau's worry that moral truths may depend on facts about the mental, which leads him to abandon the notion of mind-independence in favor of stance-independence.

Shafer-Landau speaks of "perspectives," telling us that on a realist view, the correct moral standard is the one which is correct in a perspective-independent, or as he calls it, stance-independent, way. The perspectives that Shafer-Landau has in mind are moral perspectives, that is, those perspectives from which we take particular attitudes toward moral standards. And, on Shafer-Landau's version of realism, the correct moral standard is correct regardless of what attitude anyone may take toward it. As he says, "the truth of any first-order normative standard is not a function of what anyone happens to think of it." (p. 15). This notion is quite familiar to us by now, as it is simply the notion of objectivity adduced by numerous authors above in their descriptions of realism. The fact that "what anyone happens to think" is at issue suggests mind-independence, and in fact I will say that it is merely one of the numerous precisifications of mind-independence available to us. So, on my understanding of mind-independence, Shafer-Landau's stance-independence is just one way of cashing out that notion more specifically. In particular, it is a way of spelling out the sort of mind-independence that matters for realism about ethics, and it allows the realist to maintain that ethical truths do depend on minds and mental states in certain benign ways.

4. **The Extent of Mind-Independence**

I submit that my definition of the concept of mind-independence captures what is at stake in the realism debate while avoiding some (at least potential) pitfalls. Not enough has been said up to

this point about whether and to what extent mind-independence as I define it is compatible with the large stock of realisms and anti-realisms out there which make what exists and what it is like mind-dependent to varying degrees. My hope is that my version of the concept is compatible with a wide range of views. Here's a *prima facie* reason to think that it is: My view says nothing about how mind-independent or mind-dependent things are, it only attempts to capture what the proper notion of mind-independence is. So my definition of mind-independence is supposed to leave it open whether realism or anti-realism is the correct view. But ultimately my definition does offer some guidance in the debate because the background assumption is that we should be anti-realist about some domain to the extent that the objects in that domain are mind-dependent, and realist about them to the extent that they are mind-independent. All the same, my definition doesn't say anything about the degree to which the entities in any given domain actually do depend for their existence or their nature on our directed cognitive practices. My goal has just been to explain that we should think that those entities are mind-dependent to the extent that they do, and hence that we should be anti-realist about them to the extent that they do.

It would be helpful now to think of my preferred notion of mind-independence in the context of some important views in the realism debate. As indicated at the beginning of this chapter, there is some divergence as to what exactly realism and anti-realism are supposed to be. But when it comes to particular views, there tends to be significant convergence as to whether those views ought to be considered realist or anti-realist. Thus, for instance, while there is debate over whether we should understand the realism dispute as a dispute over the sorts of truth conditions our sentences have, nearly everyone agrees that within the Dummettian realism debate, the theorist who thinks that our sentences have evidence-transcendent truth conditions is properly considered realist, while the theorist who thinks that our sentences have truth conditions

such that we must be capable, at least in principle, of discovering whether those conditions are satisfied is properly considered anti-realist. The conviction goes beyond the fact that Dummett defines his semantic realism and anti-realism in that way. The thought is that, at least as long as we are framing realism and anti-realism in terms of truth conditions, Dummett is correct to divide the opposing sides up in the way that he does.

So let us see, then, with regard to some specific realisms and anti-realisms, whether we should say that my definition of mind-independence would lead us to correctly view those philosophies as either realist or anti-realist. Here we will also be noting similarities and differences when it comes to how mind-independence is conceived according to these philosophies. We would do well to start with the classic Kantian picture. What follows is a rough and ready, but I think largely uncontentious, reading of Kant's views in the first Critique. I hope it will do for our current purposes.

In his *Critique of Pure Reason,* Kant tells us that the world as we know it is not entirely mind-independent, but it is not entirely mind-dependent either. For Kant, we must distinguish between the noumenal realm, which consists of things as they are in themselves, and the phenomenal realm, which consists of the objects of experience. The basic idea is that the noumena interact with our sensory apparatus which gives us raw sensible intuition unshaped by any concepts. In order for us to understand this raw sensible intuition, and to perceive objects as we perceive them, our minds shape the raw intuition according to our concepts. The result is the familiar, three-dimensional world of objects that we know. The phenomenal realm lies within the scope of our knowledge while the noumenal realm does not. We can't know whether the things as they are in themselves (the noumena) really correspond to the things as we perceive them (the phenomena) because the noumena aren't even possible objects of knowledge. All the

same, it is the noumena which are to some degree responsible for the phenomena as they have brushed up against our senses and provided us with the raw sensible intuition out of which, along with our concepts, we were able to construct the objects of experience. So the world that we know is partly mind-independent—our minds did not create the sensible intuition out of nothing—and partly mind-dependent—we added something to the raw intuition in order to construct the world that we know.

While few outside of the Kantian tradition subscribe to the details of the account just given, the general picture Kant left us with has been enormously influential and gives a moderate anti-realist some wiggle room between a strong realist view according to which our theories, when they succeed, can tell us how the world is in itself, and a radical constructivist view according to which the so-called "external" world is more or less entirely a human invention. Putnam sees himself as falling on this middle ground, and in his work in the late '70s through the '80s, the reign of his "internal realism," he often likens himself to Kant. Thus in the preface to his *Reason, Truth, and History* he tells us that he is after the view that "the mind and the world jointly make up the mind and the world" (1981, p. xi). Putnam's view is that we cannot describe the world as it is in itself, we can only describe it as it is within a particular conceptual scheme. There is, he would say, no "God's-eye view" of the world, no privileged vantage point from which we can assess the world independently of the concepts that we use to make sense of it (Putnam 1978, 1981, 1983). In that respect, Putnam's view is Kantian, halfway between the thesis that the world is extremely mind-dependent and the view that it is extremely mind-independent.

The general idea that the world as we know it is shaped by our own concepts, whether we have in mind Kant, the 80s Putnam, or just the generic Kantian view itself, accords well with my

definition of mind-independence. For my view says that the world we seek to discover is mind-dependent to the extent that it is shaped by our concepts (among other things). In short, the world we inquire about is mind-dependent to the extent that we as inquirers add something to that world as we seek to discover it. At the same time, my take on mind-independence does not imply any kind of Kantianism. If it turns out that the existence and nature of the entities in some domain of inquiry, say, microbiology, do not depend on our directed cognitive practices in any way, then we should be realists about microbiology because the subject matter of microbiology would be mind-independent in the way relevant to realism.

My definition of mind-independence is intended to keep us from falling into the trap of describing mind-independence in such a way that we can't be realists about minds while thinking that mind-independence or the lack of it is what makes the difference between realism and anti-realism. I submit that it adequately captures the important elements of what we want from a concept of mind-independence without building in, however inadvertently other theorists have built it in, the unacceptable conclusion that we can't be realists about minds. I furthermore claim that my definition of mind-independence supports some common realist and anti-realist intuitions. For instance, perhaps one will want to be an anti-realist about artifacts such as hammers. We made hammers, and they must be mind-dependent to the extent that they wouldn't exist if beings with minds had never made them to fulfill a particular purpose. Such mind-dependency is captured by the "conceptualizing" clause of what I said above, which I reprint here: "Something depends on our *directed cognitive practices* if and only if the act of theorizing about, conceptualizing, or otherwise uncovering the nature of that thing is what lends the thing its existence or its nature." The invention of the hammer, and subsequent creations of particular individual hammers, involve the conceptualizing of hammers. So there is mind-dependence.

But one can indeed study hammers, use hammers, take account of their properties, count them, etc., without thereby lending existence or a particular nature to hammers. So there is also mind-independence. In light of the foregoing, I submit that my definition of mind-independence accounts for the extent to which one should want to be a realist about artifacts such as hammers. Realism and anti-realism about fictional characters is more difficult, but I believe that my definition offers satisfying results in that area as well. The author of a fictional work conceptualizes (perhaps for the first time) a fictional character and thereby lends the character existence (if the character can be said to truly exist—I am committed to no view on this) and a particular nature (again, I am committed to no view as to whether fictional characters truly have determinate natures). So fictional characters have a certain sort of mind-dependence. All the same, we can study them, read about them, talk about them, etc., without thereby lending them their existences or natures. One might argue that we do change the natures of fictional characters by interpreting stories in certain ways, but this would do nothing to undermine my definition. In fact my definition would account, given such a view of the nature of literary interpretation, for why on that view fictional characters are even more mind-dependent that I originally suggested them to be. If literary critics, or we as readers, influence fictional characters in the way just described through the practice of interpreting a story, then my definition of mind-independence in terms of directed cognitive practices correctly pinpoints where the mind-dependence comes from: In the first place it comes from the author, and in the second it comes from the reader.

5. Conclusion

The notion of mind-independence is generally recognized as either explicitly or implicitly involved in statements of realism. The notion is typically not characterized specifically enough to be of any real use. What's more, some have worried that the notion of mind-independence is

so fraught with difficulty that we should abandon it. In particular, there have been concerns about how to understand mind-independence in such a way that it allows a form of realism about minds. I have offered an understanding of mind-independence that allows for such realism. A commitment to mind-independence is also often associated with a rejection of mental entities in general. I have argued that mind-independence should not commit us to such physicalism. Finally, mind-independence is a broad category which admits of many precisifications. I've argued that what matters is not necessarily that we are able to define mind-independence broadly, but that we are able to specify the particular versions of mind-independence that are relevant in each particular case.

Chapter 3: Dummett on Realism and Truth

1. Introduction

The point of this chapter is to get clear on Dummett's views on realism and truth. In addition to laying out what Dummett says about realism and truth in general, and, importantly, his argument for semantic anti-realism, this chapter will also assess the secondary, largely interpretative, literature on Dummett. The end goal is to acquire not only an understanding of what Dummett actually said, but to settle on what I see as the most plausible version of his overall argument, one which may be richly extended. For Dummett's work does not merely advance semantic anti-realism as a philosophical artifact to be gaped at while it sits idly on a shelf. Dummett attempts to overhaul our view of language at its foundations and to set it on what he sees as a more stable course. If he is wrong in the details, we should still try to salvage what is right and move forward.

For Dummett, realism and anti-realism are best understood as views about the nature of truth. The realist about a given subject matter believes that truth, as it applies to sentences about that subject matter, is *evidence-transcendent*, while the anti-realist believes that truth, as it applies to sentences about that subject matter, is *evidentially constrained*. The notion of truth which applies to a class of sentences is evidence-transcendent just in case those sentences are capable of having some truth value or other independently of our ability to ascertain whether they have any particular truth value. If the sentences of a class cannot have a truth value independently of our ability to ascertain whether they have any particular truth value, then the notion of truth which applies to that class is evidentially constrained. While some take issue with this so-called "semantic" characterization of the realism dispute, it seems clear that the characterization in terms of truth is at least one way of codifying the realist intuition that reality

is the way it is independently of our ability to discover it, and of codifying the anti-realist intuition that we cannot make sense of the world being some way if we are unable, at least in principle, of ever coming to know how it is.

Thus Dummett tells us,

> Realism I characterize as the belief that sentences of the disputed class possess an objective truth-value, independently of our means of knowing it: they are true or false in virtue of a reality existing independently of us. The anti-realist opposes to this the view that sentences of the disputed class are to be understood only by reference to the sort of thing which we count as evidence for a sentence of that class. (1978, p. 146)

And, similarly,

> The realist and the anti-realist may agree that it is an objective matter whether, in the case of any given sentence of the class, the criteria we use for judging such a sentence to be true are satisfied: the difference between them lies in the fact that, for the anti-realist, the truth of the sentence can only consist in the satisfaction of these criteria, whereas, for the realist, the sentence can be true even though we have no means of recognizing it as true. (1978, p. 147)

Dummett's characterization of the realism dispute is often framed in terms of truth conditions. We will say that a sentence has a realist truth condition if and only if the condition which must obtain in order for it to be true might obtain even though we were incapable of determining whether it did, and that a sentence has an anti-realist truth condition if and only if it cannot be true unless we are capable, at least in principle, of coming to know that it is true. Thus characterized, the realism dispute in some area becomes the debate as to whether sentences in that area have realist or anti-realist truth conditions.

Dummett sometimes speaks as though what matters to the realism dispute is the nature of truth and the sorts of truth conditions sentences have, and other times as though the central issue is rather about bivalence and the law of excluded middle. It would be fair to say that Dummett sees these issues as so closely intertwined that he does not think there is a significant difference between them, at least not as far as the realism debate is concerned. It remains to be seen whether Dummett's attitude toward these issues is appropriate.

Dummett's program has a negative component and a positive component. He is more certain about the negative component than he is about the positive one—he is convinced that a realist theory of meaning based on truth-conditions cannot work, and he does much to advance an alternative, which he admits might ultimately fail as well. That alternative is his verificationist theory, which he models directly and deliberately on intuitionism in mathematics. Dummett's overall anti-realist point is captured nicely in the following sentence: "A theory of meaning in terms of [realist] truth-conditions cannot give an intelligible account of a speaker's mastery of his language." (Dummett 1993, p. 74).

A few remarks on terminology are in order. Dummett's views on the theory of meaning are typically characterized as being about a theory of understanding (Cf. Alston 1996, p. 106. Devitt 1997, p. 268). This terminology is natural, but since for Dummett, to understand a sentence is just to know what it means, and vice versa, we should not feel any special pressure to frame things in terms of understanding rather than knowledge. As we will see below, this issue is somewhat dicey, but I maintain that my usage is harmless.

Some of Dummett's interpreters have claimed that Dummett's goal was not to establish that the truth-conditions of our sentences are epistemically constrained, but rather to establish that the notion of truth-conditions should be replaced in our theory of meaning with a notion of

evidential support. Though I generally agree with this interpretation, I see no harm in continuing to speak of truth and truth conditions as long as we are precise about whether these conditions themselves are evidence-transcendent or evidentially constrained (that is, realist or anti-realist). In other words: We can speak of truth and truth-conditions in both cases, but it will turn out that the anti-realist works with an evidentially-constrained notion of truth based on warranted assertibility. I take this way of speaking to be justified by Dummett's own remarks in the preface to his *Truth and Other Enigmas,* among them*:* "The problem is not whether meaning is to be explained in terms of truth-conditions, but of what notion of truth is admissible" (1978, p. xxii).

Semantic realism is associated with the following four theses concerning sentences about the subject matter for which realism is claimed to hold:

1R. The appropriate notion of truth for the sentences is evidence-transcendent.

2R. The sentences have realist, and not anti-realist, truth conditions.

3R. The principle of bivalence holds for the sentences.

4R. The law of excluded middle holds for the sentences.

Semantic anti-realism, in turn, is associated with the following four theses concerning sentences about the subject matter for which anti-realism is claimed to hold:

1AR. The appropriate notion of truth for the sentences is evidentially constrained.

2AR. The sentences have anti-realist, and not realist, truth conditions.

3AR. The principle of bivalence does not hold for the sentences.

4AR. The law of excluded middle does not hold for the sentences.

We will begin with rather simple versions of Dummett's arguments and allow for complications as we go along. While he is typically understood as presenting two different

arguments against semantic realism, one focusing on how we acquire language and the other focusing on how we manifest our knowledge of language, the manifestation argument is by far the most discussed. Importantly, the centerpiece of both arguments is Dummett's insistence on the publicity of meaning. Some would even say that the acquisition argument and the manifestation argument are essentially the same (see Gardiner 2000, p. 30). Dummett's arguments, in their simplest forms, are best understood as instances of *modus tollens.* We will focus our attention for now on the argument from manifestation, bringing the acquisition argument, and something else I will call the endowment argument, into the picture later on.

The Manifestation Argument

M1. If undecidable sentences have realist truth conditions, then we are able to manifest a grasp of those conditions.

M2. We are not able to manifest a grasp of the realist truth conditions of undecidables.

M3. Therefore, undecidable sentences do not have realist truth conditions.

In what follows, I will attempt to work out the contentious points. Perhaps the most difficult and tangled issue involves the proper understanding of undecidability. In the process of coming to understand this concept and the role it plays in Dummett's arguments, we will be elucidating other important concepts as well.

Undecidability is centrally important precisely because of its epistemic nature. Undecidability is a brand of unknowability, and as we will see, the debate between the semantic realist and the semantic anti-realist in a particular area of discourse comes down to the debate over whether the sentences in that area have truth values which are essentially knowable. Anti-realist truth is essentially knowable, while realist truth is not. Oversimplifying somewhat, we can say that if a sentence is a candidate for truth or falsity despite being undecidable, then

realism holds for that sentence. Apply this criterion to other sentences in turn, and one will arrive at realism (or anti-realism) for a whole discourse. Dummett takes the question of which sort of truth applies in an area to be answered by a determination of what sort of meanings the sentences in that area could possibly have. The meanings that our sentences can have are, for Dummett, determined by the way that we use them. What we learn when we learn a language is how to use it, and what we manifest to others when we use language is our particular understanding of it, that is, our knowledge of how to use it.

Dummett's over-arching semantic anti-realism is generated via analogy with intuitionism, and more broadly, constructivism, in the philosophy of mathematics. Dummett is interested in importing, more or less wholesale, the motivation for mathematical intuitionism into non-mathematical areas. The following passage illuminates this motivation particularly well:

> Constructivist philosophies of mathematics insist that the meanings of all terms… must be given in relation to constructions which we are capable of effecting, and of our capacity to recognise such constructions as providing proofs of those statements… The most powerful form of argument in favour of such a constructivist view is that which insists that there is no other means by which we can give meaning to mathematical expressions. We learn, and can only learn, their meanings by a training in their use… there is no means by which we could derive from such a training a grasp of anything transcending it, such as a notion of truth and falsity for mathematical statements independent of our means of recognising their truth-values. (Dummett 1975, p. 301).

Restrictions on our ability to recognize truth values, then, are directly connected to the meanings that our sentences have, and therefore to the uses we make of them. To understand Dummett's

reasoning, we must understand what the relevant epistemic restrictions are. Hence we must get clear on undecidability.

2. Undecidability

2.1 What is Undecidability and How Does it Figure in the Dialectic?

It is not news that Dummett's use of the notion of undecidability is idiosyncratic. The question is how we are to understand the notion in the context of his arguments. Dummett has borrowed the notions of decidability and undecidability from logic and mathematics, where those notions receive a precise treatment. Alfred Tarski, for one, writes that a formalized theory T is decidable if there is a decision procedure, that is, "a method which permits us to decide in each particular case whether a given sentence formulated in the symbolism of T can be proved by means of the devices available in T (or, more generally, can be recognized as valid in T)" (Tarski 1953, p. 3). Later in the same monograph, this method is said to be "mechanical" (p. 13). The mechanical nature of the procedure is crucial to a proper understanding of decidability. Similarly, Rogers writes that a theory T is decidable if, "there exists an effective procedure—i.e., an algorithm—for determining whether an arbitrary formula A of T is a theorem of T" (Rogers 1971, p. 215), telling us further that, "these procedures are effective or algorithmic in the sense that they provide us with instructions for ascertaining something or other in a systematic, step-by-step manner" (Rogers 1971, p. 18). Thus we have a clear conception, in mathematical logic, of how a *formalized theory* can be said to be decidable.

Sanford Shieh complains that, "this notion of undecidability is defined with respect to a formal system, so a formula cannot be undecidable in this sense simpliciter, without reference to a formal system, and, moreover, every formula is trivially decidable with respect to a set of axioms that includes it" (1997, p. 57). Thus Dummett is supposed to be somewhat out of line

when he applies the undecidability concept to individual sentences independent of any particular formalized theory. We can, however, easily apply the notion to individual sentences within a formal system, as Rogers explains: "Any sentence A of a theory T which is such that neither A nor ~A is a theorem of T is said to be *undecidable* in T" (Rogers 1971, p. 201). Hunter speaks more generally in a way that easily permits application of the decidability concept to individual sentences all on their own: "In logic and mathematics, an *effective method* for solving a problem is a method for computing the answer that, if followed correctly and as far as may be necessary, is logically bound to give the right answer (and no wrong answers) in a finite number of steps" (Hunter 1973, p. 14). All that we must do, then, is take the problem at hand to be determining whether a given sentence or its negation is true. If we possess an effective procedure for determining such a thing, then the sentence is decidable. If we do not possess such a procedure, then the sentence is undecidable. Shieh himself offers a similar definition of undecidability for sentences, according to which, "a sentence is undecidable just in case it is not presently decided, and, we don't have an effective procedure for [deciding] it" (1997, p. 61). This understanding of undecidable sentences is also endorsed by other commentators on Dummett (Cf. Kapsner 2015, p. 15; Tennant 1997, p. 161), and meshes well with what Dummett himself says about the notion, though as we will see Dummett sometimes appears to equivocate on two distinct notions (Cf. Tennant 1997, p. 183). I will be proceeding under the presumption that Shieh's formulation is a good one, though we will soon see that the formulation is not without issue. We will work with the following definitions of decidability and undecidability in the context of the semantic realism debate:

undecidable sentence: A sentence is undecidable if and only if neither it nor its negation is currently proven and we possess no effective procedure for proving either it or its negation.

decidable sentence: A sentence is decidable iff either it or its negation is currently proven or we possess an effective procedure for proving either it or its negation.

Note that what matters for decidability in the relevant sense is whether we possess an effective procedure, not whether an effective procedure merely exists. There will be more to say about this later, but for now let it suffice to say that, as far as what we are able to know is concerned, the existence of an effective procedure doesn't do us any good if we aren't able (in the relevant sense of 'able') to carry it out, which we can't do if we are not in possession of it.

We might have said that we either do or do not possess an effective procedure for determining whether the sentence is true or false, but this formulation would seem to build in an assumption of bivalence, an assumption we must avoid when giving definitions which are supposed to be acceptable to realists and anti-realists alike. As we have defined undecidability, that there are undecidable sentences is a fact which both the realist and the anti-realist must accept—the existence of undecidable sentences is common ground between them. Alternatively, in order to avoid an assumption of bivalence, we may have loosened things up a little and framed the concept in terms of determining which truth value, if any, the sentence had. But this semantic way of dealing with the matter in terms of truth values is, overall, less faithful to the original conception of decidability borrowed from mathematical logic, where the emphasis is on what is provable, not on the semantic values of expressions. Nevertheless, in what follows we will often speak of what we can establish as true, what we can know as true, etc. This way of speaking does not conflict with our definitions, since to prove something, generally speaking, is

to prove that it is true (though it is worth noting that Dummett, and others after him, have suggested an alternative to realism in terms of falsification rather than verification). As we seek to extend the intuitionistic/constructivist model of anti-realism from the philosophy of mathematics to the empirical domain, our interest will be centered on what we can establish as true, i.e., what we can prove.

Early on, we mentioned that, for Dummett anyway, the question as to whether bivalence applies to a given class of sentences and the question as to whether those sentences have realist truth conditions appear to be one and the same. Those who have come after him have questioned such a tight connection between the two notions and have asked whether the anti-realist need reject bivalence (Wright 1993, Tennant 1997, Rasmussen and Ravnkilde 1982.). Such thinkers may take issue with the following definition of bivalence, since, as we will see, its modal elements make the link to evidence-transcendent truth rather straightforward, but the definition is independently defensible.

> **bivalence:** The principle of bivalence holds for a class of sentences iff each sentence in the class has just one of the values true or false in each possible world.

As Dummett emphasizes, bivalence is a semantic principle. The law of excluded middle, however, is a logical law (Dummett 1978, pp. xviii-xix). If we accept bivalence for a class of sentences, then excluded middle follows. Bivalence, however, does not necessarily follow from excluded middle. From the fact that at least one of either P or its negation is true, we cannot conclude that there are only two possible truth values, much less that just one of them is had by each P.

> **excluded middle:** The law of excluded middle holds for a class of sentences iff for each sentence φ in the class, (φ v ~ φ) is satisfied in each possible world.

In the literature on our topic, the distinction between bivalence and excluded middle is typically glossed over, most often without consequence. It is not uncommon to see bivalence formalized in the following way, though such a formalization clearly is better understood as capturing the law of excluded middle:

$$\Box(\varphi \vee \sim\varphi)$$

For our purposes, there is no harm in expressing bivalence this way, especially since the logical law follows from the semantic principle. In fact, it will be useful to have this formalization in our toolkit. In logic, a *class of expressions* is said to be decidable if there is an effective method for determining, for any given expression, whether it is a member of the class. Dummett's terminology, perhaps confusingly, permits talk of *classes of undecidable sentences.* It is important to keep in mind that, in this context, Dummett takes the concept of undecidability to apply to individual sentences according to our definition, and he then wishes to group these sentences together into classes based on subject matter. Whether these classes are themselves decidable, that is, whether there is an effective method for determining, for any sentence, whether it belongs in the class, is not at issue. Each of the resulting classes is what he terms a "disputed class" (Dummett 1978, p. 146). They are disputed because the realist and anti-realist disagree as to how we should treat them. The realist applies bivalence and evidence-transcendent truth to them, despite their undecidability, whereas the anti-realist does not. Those who have questioned the anti-realist's need to reject classical logic essentially argue that the anti-realist might apply bivalence to a class of sentences while still endorsing an evidentially constrained notion of truth for the class.

It is not appreciated often enough that, whatever notion of undecidability one takes Dummett to be employing, it must be one that both the realist and the anti-realist can accept.

This is because the dispute between the realist and the anti-realist only arises for undecidables—the realist wishes to apply bivalence to them while the anti-realist does not. Dummett's arguments against realism and for anti-realism take it as a premise not only that undecidable sentences exist, but that we know what they mean. It is the questions of how we could possibly know what they mean, how we could endow them with the meanings they have, how we could convey those meanings to each other, and the different answers offered by the competing views which mark the differences between the realist and the anti-realist. The dual facts that both the realist and anti-realist admit the existence of undecidables, and that the dispute between them only emerges for undecidables, is not always sufficiently taken into account, or even recognized, in the vast literature on our topic. These points deserve further attention, which they will receive later on. For now they should simply be borne in mind.

The semantic realist, for most intents and purposes, subscribes to all of 1R-4R above. Thus, for some class of sentences toward which the realist takes a realist attitude, each of 1R-4R should be claimed to hold. The realist endorses evidence-transcendent truth for the class:

> **evidence-transcendent truth:** Evidence-transcendent truth applies to a class of sentences iff those sentences can have a truth value independently of our ability to ascertain which one they have.

As with the definitions of decidability and undecidability, we need to avoid building in an assumption of bivalence. The matter is less pressing here, since clearly evidence-transcendent truth is intended to be a realist notion. All the same, for precision, we will avoid building presumptions about how many truth values there are into any of our definitions that are not specifically concerned with how many truth values there are. Earlier, it was pointed out that the

realism debate is often framed in terms of truth conditions. Our notion of evidence-transcendent truth can be easily transformed into a realist notion of truth conditions:

> **realist truth conditions:** A sentence has a realist truth condition iff the condition which must obtain in order for it to be true might obtain even though we were incapable of determining whether it did.

Our anti-realist endorses 1A-4A above (though see the next chapter for some debate over this point). Correspondingly, we can define the anti-realist notions of evidentially constrained truth and anti-realist truth conditions as follows:

> **evidentially constrained truth:** Evidentially constrained truth applies to a class of sentences iff those sentences cannot have a truth value independently of our ability to ascertain which one they have.
>
> **anti-realist truth conditions:** A sentence has an anti-realist truth condition iff the condition which must obtain in order for the sentence to be true is one whose obtaining we can, at least in principle, recognize when it does obtain.

These anti-realist ideas are commonly associated with a modal conditional, often referred to as the *knowability principle:*

> **knowability principle:** $\Box\forall\varphi\,(\varphi \rightarrow \Diamond K\varphi)$

Where the universal quantifier is taken to be restricted to the disputed class in question, the meta-variable φ is taken to range over sentences (including their negations), and the predicate K is taken to mean 'is known'. The principle just says that, necessarily, if something is the case then it is possible for us to know that it is. As we will see, some work is required to determine the relevant sense of "is possible." The anti-realist accepts, while the realist denies, the knowability principle. That is, the realist endorses the negation of the knowability principle:

$$\sim\Box\forall\varphi\,(\varphi \rightarrow \Diamond K\varphi)$$

Importantly, the modal operators in this case are not epistemic modal operators.

We will continue to speak of a sentence being 'knowable' in such a way that does not presuppose its actually being true. If a sentence is actually false, then in some sense, of course, we cannot know it because it is not true. What we mean when we say that a sentence is knowable is just that, whatever truth value it may have, it is possible for us to determine that it has that value. Likewise, to say that a sentence is unknowable is just to say that, whatever truth value it may have, it is not possible for us to determine that it has that value.

Dummett's commentators often write as though an undecidable sentence is one whose truth value is unknowable, or potentially unknowable, or one whose truth is at least potentially recognition-transcendent (the term "verification-transcendent" is also common). Crispin Wright, for instance, speaks of the challenge posed by sentences which may be "undetectably true" (1993, p. 248). Dummett himself writes that we can be sure that "every statement is either true or false…only if…we can form a conception of what it is for a statement to *be* true, or to *be* false, even when we are incapable of recognising it as such" (1978, p. xl). Other examples abound, some of which will be adduced below.

Consider now a sentence which satisfies our definition of undecidability. Is this sentence one that we cannot know? Not necessarily. Indeed, if a sentence is unknowable, then it is, *a fortiori*, undecidable, since there clearly could be no effective procedure for coming to know it, given that it is impossible to know. And if the sentence is decidable, then it is, *a fortiori*, knowable, since not only could we come to know it, but we in fact have an effective procedure for coming to know it. None of this forces the conclusion that we cannot come to know a sentence which is undecidable (on these matters compare Tennant 1997, p. 185 and Shieh 1997,

p. 62). All that undecidability requires, as we are understanding it, is that we don't currently know the sentence or its negation, and we have no effective procedure that would guarantee that we could, in a finite amount of time, come to know either it or its negation. It is completely consistent with this idea that it be possible, at any given time, for us to come upon reason sufficient to establish the sentence in question or its negation. So if our possession of an effective procedure for deciding a sentence is what matters, construing undecidability as a strong form of unknowability appears quite inappropriate.

Neil Tennant explicitly defines undecidability in terms of effective decision procedures, as we have here (Tennant 1997, p. 161). He then goes on to interpret Dummett as equivocating between two notions of undecidability, one involving effective decision procedures and the other involving general unknowability (Tennant 1997, p. 183). The charge of equivocation is a subject for our discussion of whether the anti-realist should reject classical logic. According to the latter understanding of undecidability, a sentence "is not recognizably true, if it is true, and is not recognizably false, if it is false" (Tennant 1997, p. 183). Joe Salerno tells us that an undecidable sentence is one that "cannot in principle be proved or disproved" (2000, p. 213). Similarly, Jon Cogburn works with the notion that, possibly, there is a sentence such that neither it nor its negation is knowable (2005, p. 24). This general conception of undecidability can be represented symbolically in the following way:

$$\exists\varphi(\sim\Diamond K\varphi\wedge\sim\Diamond K\sim\varphi)$$

Similar notation is in fact employed by Salerno and Cogburn (the latter using subscripts to denote different concepts of knowability, a point to which we will return). Later I will take issue with their use of the existential quantifier, but for our current purposes it makes no difference. There are two major problems with this conception of undecidability as utter unknowability.

The first is that, as all three of the authors mentioned in the foregoing paragraph themselves point out, such a formulation is intuitionistically inconsistent with the anti-realist's coveted knowability principle. Hence, a semantic anti-realist of the sort that matters to us cannot accept that there are any undecidable sentences as long as the undecidables are considered to be completely unknowable One reason that this is a problem is that, on a very reasonable understanding of the debate as Dummett frames it, the realist and the anti-realist only disagree about how to treat undecidables. At the very least, it is commonly understood that the debate between them is only discernable with regard to undecidables. Salerno and Cogburn have tried to get around this problem in different ways. Tennant and Salerno have taken the result to show that Dummett has failed to effectively argue for the conclusion that a semantic anti-realist ought to reject classical logic. What impetus the anti-realist may have for rejecting classical logic deserves a large discussion of its own. For the time being I will restrict my attention to the question of which notion of undecidability we ought to employ in framing the semantic realism debate.

The second major problem with conceiving of undecidability as unknowability is that for an intuitionist, to be justified in asserting $\sim\varphi$ it is sufficient to show that it is impossible to prove (or know) φ. If we knew that there were undecidable sentences in the sense under consideration, such that it were impossible to prove φ and impossible to prove $\sim\varphi$, then we would be justified in asserting both $\sim\varphi$ and $\sim\sim\varphi$. So this conception of undecidability leads intuitionistically to absurdity all on its own. Hence it should come as no surprise that one can derive absurdity from this formulation of undecidability plus the knowability principle. One could intuitionistically derive absurdity from undecidability qua unknowability plus anything at all.

One finds the sort of reasoning just given in Dummett. "It is impossible, therefore, that we should ever be in a position to assert, of any statement A, that A is (absolutely) neither provable nor refutable" (1977, p. 17). Similar results are discussed in Williamson (1994) and Kapsner (2015). Williamson shows that an intuitionist cannot, while remaining consistent, say that any given proposition will never be decided, that is, that it will be neither proven nor refuted at any time (1994, pp. 136-137). His reasoning is different from, though similar to, the reasoning I have just led us through. It should be noted that if, as the understanding of undecidability qua unknowability entails, an undecidable sentence is such that it is impossible to either prove it or refute it, then it is also such that it will never be either proven or refuted (compare Williamson p. 136). Thus, even if the logic I offered above fails, the intuitionist who wishes to adopt the current notion of undecidability is in trouble.

It was noted above that nothing in our definition of undecidability, based on the possession of an effective decision procedure, entails the impossibility of coming to know an undecidable sentence. Dummett himself would appear to endorse an understanding like ours, as when he says, "of course, for any given undecidable sentence, the possibility may well be open that we may find ourselves in a position to recognize that the condition for its truth is satisfied, or that it is not" (1993, p. 46). This sort of statement should lead us to be skeptical of an identification of undecidability with unknowability. Here is another example illustrating a conception of undecidability like ours, based on the possession of an effective decision procedure:

> An undecidable sentence is simply one whose sense is such that, though in certain effectively recognizable situations we acknowledge it as true, in others we acknowledge

> it as false, and in yet others no decision is possible, we possess no effective means for bringing about a situation of one or other of the first two kinds. (Dummett 1973)

Still, Dummett can be caught speaking as though undecidable sentences are actually unknowable. Consider for instance, the following quotation:

> It is when the principle of bivalence is applied to undecidable statements that we find ourselves in the position of being unable to equate an ability to recognize when a statement has been established as true or as false with a knowledge of its truth-condition, since it may be true in cases when we lack the means to recognize it as true or when we lack the means to recognize it as false. (1993, p. 62)

Saying that we "lack the means" to determine the truth value of a sentence might not mean that we are incapable of determining it—perhaps all Dummett wants to get across is that the application of bivalence to undecidables entails that they might be true, and might be false, despite our lacking an effective procedure for determining their truth value. This sort of interpretation makes sense in light of Dummett's (admittedly not so clear) expositions of what is problematic about undecidables. The problem with them is supposed to be that "we have no guarantee either that a situation [where we can recognize them as true or as false] will occur, or that we can bring about such a situation at will" (Dummett 1973, p. 468). In other words, we cannot be sure that we would land upon an answer in a finite amount of time. All the same, we cannot be sure that we would not come upon the answer in finite time either. That is what follows from our lacking an effective decision procedure (compare Shieh 1997, p. 61 and Kapsner 2015, p. 16). Again, it is not being claimed that we *could not* determine the answer—this idea is incomprehensible to the intuitionist, as we have seen. Our notion of undecidability

should be intuitionistically acceptable, as one based on the possession of effective decision procedures is.

Other passages from Dummett, on the other hand, make it less mysterious how his interpreters would come to impute to him the view that undecidables are unknowable, as when he says that an undecidable sentence has a truth condition "which we are not, in general, capable of recognizing as obtaining when it obtains" (1993, p. 64). Similarly, regarding a sentence which is "effectively undecidable," he writes, "the condition which must, in general, obtain for it to be true is not one which we are capable of recognising whenever it obtains, or of getting ourselves in a position to do so" (1983, pp. 106-107). While these passages present some difficulty for an effective procedure reading of Dummett's undecidability, we can take his use of "in general" and "whenever" to imply that, in some cases, we might not be guaranteed to recognize the obtaining of the relevant condition in a finite amount of time, though this does not mean that there are not other circumstances under which we could. This understanding of Dummett meshes with our simple lack of an effective decision procedure, that is, lack of a guarantee that we could, in finite time, determine the sentence's truth value, and is consistent with undecidable sentences being knowable. Textual evidence that Dummettian undecidability just is unknowability is therefore lacking, even aside from the difficulties laid out earlier.

Ultimately, it may be the case that Dummett indeed ambiguously employs the term 'undecidable', and perhaps even equivocates on two different senses of the term in order to get his reasoning to work, as Tennant claims. Regardless, we are interested in the best version of Dummettian reasoning. With that end in mind, given the difficulties plaguing an unknowability conception of undecidability, we will continue to employ 'undecidable' and 'decidable' as we have defined them.

2.2 The Relativity of Undecidability

We have said that an undecidable sentence is one for which we possess no effective decision procedure. Recall that the semantic realism debate is supposed to be over how to treat undecidables. The realist says that undecidables have determinate (in particular, bivalent) truth values, where the anti-realist says that we are unwarranted in thinking that undecidables have determinate truth values of any kind. This is because the realist thinks that all of our sentences have determinate truth values completely independently of our ability to come to know what those values are, while the anti-realist is only willing to admit that decidable sentences have determinate truth values. Thus the anti-realist is unwilling to endorse the principle of bivalence and law of excluded middle for undecidable sentences. The semantic anti-realist attitude toward undecidables can be understood as a certain brand of agnosticism according to which we do not know that we can't know (as this would be contradictory on the anti-realist semantics), but we also do not know that we can. The agnostic, intuitionist anti-realist outlook is explored in detail in the next chapter.

As Kapsner points out, decidability and undecidability are for Dummett "epistemic concept[s]" (Kapsner 2015, p. 16). What matters is not whether we can construct a proof of a sentence, but rather whether we know that we could construct such a proof. Possession of an effective decision procedure is reason enough, for the anti-realist, to believe that the sentence to which the procedure applies has a determinate truth value. It is worth noting that decidability is something that a sentence might lack at one time and then acquire at a later time. Say that we have no effective decision procedure for a sentence S currently, and we also have no reason to believe that we could come up with a procedure—perhaps we could, but we do not currently know that. The anti-realist says in such a case that S has no determinate truth value (or, perhaps

better, that we are not warranted in asserting that it does). If we did have a procedure, or otherwise knew that we had, at the very least, an effective means for producing one, the anti-realist would be comfortable thinking that S had one truth value or another determinately. The realist would also be comfortable thinking so—but that is only because the realist thinks that every sentence, undecidable or not, has a determinate truth value. Again, the debate is over the status of undecidables, not the status of decidables. The key is understanding why the anti-realist wants to treat the undecidables differently, and that is our challenge.

A realist will think that an undecidable has a determinate truth value in every case. Hence the realist does not need to worry about what exactly counts as an undecidable, except perhaps to the extent that she wants to make her debate with the anti-realist clear. A realist can even admit that there are undecidables of the sort the anti-realist cannot countenance, to wit, sentences which can be neither proven nor refuted. The realist can believe in the existence of utter unknowables while thinking that those sentences have determinate truth values nonetheless. Things are much more complicated for the anti-realist. The anti-realist does not believe that undecidables have determinate truth values. In order for the anti-realist point of view to be at all tempting, we must figure out why the anti-realist is so opposed to the realist on this point. By the end of this dissertation, I hope to convince the reader that no one has to date suggested a concept of undecidability that can support Dummettian anti-realism.

Dummett typically speaks as though the relevant notion of undecidability is not only undecidability in principle, but also undecidability for humanity generally:

> Many features of natural language contribute to the formation of sentences not in principle decidable: the use of quantification over an infinite or unsurveyable domain (e.g. over all future times); the use of the subjunctive conditional, or of expressions

> explainable only by means of it; the possibility of referring to regions of space-time in principle inaccessible to us. (Dummett 1993, p. 46)

Consider the following purported example of an undecidable sentence:

> 'Caesar sneezed when he crossed the Rubicon'.

It is not entirely clear that the undecidability of that sentence should be understood as in-principle undecidability or merely circumstantial undecidability. For all we know, one day we will construct a time machine that would allow us to go back and see whether Caesar in fact sneezed when he crossed the Rubicon. At the same time, for all we know, travel to the past is logically impossible—I take it that no one has ever given an incontrovertible proof that travel to the past either is or is not logically consistent. So what exactly is undecidability "in principle" supposed to be? I would suggest that Dummett typically uses the phrase 'in principle' to indicate that he does not mean undecidability "in practice," where something is undecidable in practice if it is not within our actual means to land upon an answer as to the sentence's truth value in our lifetimes (see the discussion of strict finitism in the next section).

Things get more complex still. Just how difficult does it have to be for us to determine the truth value of a sentence in order for it to count as undecidable in practice? In 1961, John F. Kennedy gave a speech in which he said that the U.S. should put a man on the moon. At the time, it was not possible in practice to get to the moon and back safely—the whole goal was to make that possible in practice. And yet, in some sense we did have the means in practice of getting someone to the moon and back in one piece. The means was to continue to develop the requisite technology and methods. So, while in 1961 we did not have a procedure which we could implement for safe moon travel, we had a procedure for getting a moon-travel procedure. Likewise, sometimes we may not currently possess an effective decision procedure though we

know that we could get one, perhaps because we have an effective procedure for producing an effective decision procedure. Recall that the importance of the notions of decidability and undecidability here is epistemic twice over: The importance for the anti-realist surrounds not just certain restrictions on what we are capable of coming to know, but also what we currently know about what we are capable of coming to know. According to Dummett's anti-realist, we are not warranted in thinking that a sentence has a determinate truth value unless we have assurance that a decision as to what that truth value is would be reached in a finite amount of time. One may be concerned that the requirement that the relevant decision procedures be effective in the sense of halting in finite time is irrelevant. This is a real concern for the Dummettian anti-realist, and I eventually conclude that talk of effective decision procedures is problematic for the anti-realist who wishes to extend the intuitionistic model to the contingent, empirical realm (for discussion see section 3 and section 5, esp. beginning on p. 186).

Given that the issue for the anti-realist is whether we are warranted in thinking that a sentence has a determinate truth value, it is not enough that the procedure exist—we must know the procedure and be able, in some sense, to implement it. Exactly what restriction should be placed on our ability to implement it is unclear and underspecified in the literature. One suggestion is pursued in the next section which pertains to strict finitism. But regardless of the restriction one chooses, it must be admitted that most restrictions are going to be relative to individuals, current states of information, evidence available or potentially acquirable, etc. Kapsner, inspired by Brouwer, gives us the example of "there are seventy consecutive 7s in the decimal expansion of π" (Kapsner, 2015 p. 15). We currently possess no effective procedure for determining whether that sentence is true. All the same, if we were to sit down and calculate the expansion of π, for all we know we might come across seventy 7s. And at that point we would

have an effective decision procedure for determining the truth value of the sentence in question. The procedure is to sit down and calculate until you reach seventy 7s, and we know that it will terminate in a finite amount of time because we performed the calculation in a finite amount of time. Until we performed the calculation, of course, we did not know whether our efforts would ever terminate with success, and so an anti-realist would tell us that we were not warranted in thinking that 'there are seventy consecutive 7s in the decimal expansion of π' had a determinate truth value. After we performed the calculation, we became so warranted. The sentence became effectively decidable for us.

Now imagine that I, in my office alone, have performed the calculation in question and have thus concluded that the decimal expansion of π contains seventy consecutive 7s, while no one else has ever performed the calculation. Pretend in addition that I am an anti-realist of the Dummettian bent. I should then think that 'there are seventy consecutive 7s in the decimal expansion of π' has a determinate truth value—in particular I should think that it is true since I found seventy consecutive 7s. But would that sentence be decidable for anyone else? Things become odd here for the anti-realist. Say that the sentence is only decidable for me, while it remains undecidable for everyone else. After all, according to the Dummettian anti-realist, no one else has a good reason to think that the sentence has a determinate truth value. So they should not believe that it does, and since the sentence is undecidable for them, the Dummettian has to say that it in fact does not have a determinate truth value for them. At the same time, the sentence has a determinate truth value for me. Such relativity is odd, but unavoidable for the anti-realist as long as the sentence is only decidable for me.

Say that when I perform the calculation the sentence becomes decidable for at least some others, perhaps only people near me whom I am likely to talk to, or perhaps everyone in my

linguistic community, if not everyone in the universe. If that is the case, the anti-realist has to say that 'there are seventy consecutive 7s in the decimal expansion of π' has a determinate truth value even though only one person (me) is warranted in thinking that it does. The oddness here is that a sentence can be decidable for someone even though that person has no reason to think that it is, but decidability was supposed to be an epistemic notion about what one is warranted in believing. A similar issue arises with epistemic conceptions of truth according to which truth is cashed out in terms of justification. No one wants to say that what is true depends on what we are currently justified in believing. Rather, epistemic truth theorists usually say that what is true is what we could achieve enough evidence to believe, or perhaps what we would finally believe at the end of theorizing—that way, something might be true even though we were not currently justified in believing it. Similarly, a Dummettian can say that something can be decidable for you even if you are not currently aware that it is. Indeed, one might say that, upon my successful calculation of seventy consecutive 7s, everyone acquired an effective procedure for determining whether 'there are seventy consecutive 7s in the decimal expansion of π' is true—they need only ask me and I will tell them to perform however many steps it took me to get to the result. Thus people might be in possession of a procedure without realizing it. That is not to say that they ought to believe that 'there are seventy consecutive 7s in the decimal expansion of π' has a determinate truth value, though it is true for them that it does.

If the decidability of a sentence can vary from person to person, or community to community, there will also be difficulty in cases where, for instance, someone is told that I have an effective procedure for deciding a sentence though they are not themselves in possession of that procedure. For in that case, if the person is a Dummettian anti-realist, she should believe that the sentence has a determinate truth value for me. But how could she admit as much without

thinking that it also has a determinate truth value for her, which she should not think since she does not possess the procedure? The solution I suggest is that she *does* have an effective procedure for acquiring the procedure in question, which ought to be enough to warrant her in thinking that the sentence has a determinate truth value. The procedure she possesses simply involves learning the decision procedure for the sentence from me.

Ultimately, a generalized intuitionism of the type Dummett envisioned may not be entirely coherent, or it may be coherent while its adherents are forced to take on some costly intellectual commitments. Here I have only been concerned to illustrate what some of the difficulties with applying the notion of undecidability to sentences in the contingent, empirical realm might be and suggest some potential solutions. I take on the issue in more depth at the end of the next chapter, which deals again with the concept of undecidability in the debate over whether a Dummettian anti-realist ought to endorse non-classical logic.

3. Strict Finitism and Meaning-Relevant Epistemic Constraints

Undecidability is one form of unknowability—it is one form of constraint on what we are capable of coming to know. What we ought to ask at this point is why this particular form of unknowability, the lacking of an effective decision procedure, is so crucially relevant to what our sentences can mean than other forms of unknowability. If you, possessed of all normal human faculties, are walking in Central Park and want to know whether your roommate is currently in the apartment you share only a few blocks away, in only the most extraordinarily uninteresting sense are you not able to know whether your roommate is home. In any interesting sense, you are quite capable of finding out whether he is home by calling, or by going there to see. We can imagine a not-so-far-fetched case where you are actually being held prisoner in Central Park, forbidden from contacting anyone or communicating with anyone in any way; in such a

situation, there is an interesting sense in which you are unable to find out whether your roommate is home. However, we do not think that the limitations which have been placed on your epistemic capacities in this case are relevant to the kinds of meanings our sentences can have the way that Dummett thinks that undecidability is relevant to those meanings.

Clearly there are lines to be drawn, but where are we to draw them? To the extent that this question has been pursued, it has been in the context of a view known as *strict finitism*. Strict finitism is the view that the epistemic constraints which matter to the meanings of our sentences are not just those constraints limiting what we could calculate in principle, but also those constraints limiting what we can actually calculate in practice (Dummett (1975) distinguishes "in principle" from "in practice." Tennant (1997) uses "in principle" and "feasibly" to capture the same idea). For if, as Dummett contends, what matters to meaning is the use we are capable of making of our expressions, then why, says the strict finitist, should we think that an evidence-gathering procedure which we might employ for an infinite amount of time without arriving at an answer is any more relevant to what our sentences could not mean than an evidence-gathering procedure which, though guaranteed to terminate in finite time, might take longer to produce a decision than there is time left before the human race becomes extinct, or even than there is time left before the universe collapses back in on itself? A procedure which is guaranteed to terminate in finite time is guaranteed to produce an answer *in principle.* It is not guaranteed to produce an answer *in practice.* The strict finitist simply pushes the reasoning of the intuitionist farther than the intuitionist cared to push it. Indeed, if actual use it what matters to our giving, learning, and conveying meaning, the strict finitist's question is a very good one.

Tennant has an interesting answer to these finitistic worries. He employs the notion of "scaling" to deal with the kinds of cases to which the strict finitist wishes to extend intuitionistic

reasoning. This sort of scaling simply permits whatever increases in storage capacity, computational speed, etc. may be necessary to allow a human being, or perhaps a mechanical computer which employs the sorts of computational procedures humans are capable of, to arrive at an answer as to whether a given sentence or its negation is true (1997, pp. 146-147). In essence, Tennant allows for the overcoming of limitations of space and time which otherwise set constraints on our ability to employ effective decision procedures. This sort of scaling is justified, he says, because "the actual limits to effective human thought…that we are thinking of ourselves as transcending here are not limits to the kind of thinking we may do, but only limits on how much of that kind of thinking one could do" (1997, p. 147). It is important to see that while Tennant's scaling technique works for decidable sentences whose decision procedures might have taken us an incredibly long time to carry out, it does not work for sentences for which we lack an effective decision procedure. And this result is desirable, as it allows us to draw the appropriate line of knowability precisely where the intuitionist draws it, between what is effectively decidable and what is not.

So we have some kind of an explanation for why we should halt intuitionistic reasoning where the intuitionist does. What we still lack is a real explanation for why the constraints imposed by the lack of an effective decision procedure are indeed relevant to meaning as opposed to, say, the constraints imposed by a notion of undecidability understood as unknowability. So far the best reasons we have are these: An anti-realist needs to accept that there are sentences which are not effectively decidable, an anti-realist cannot accept a strong form of unknowability, and it is crucial to the whole semantic realism debate that both sides accept the existence of undecidable sentences. All signs point to those undecidables being sentences without effective decision procedures. Our position is in line with Kapsner, who tells

us in his discussion of the intuitionistic unacceptability of what he calls "absolute undecidability," that is, the idea that we will never decide a given sentence, that "the most an intuitionist can mean when she says that a statement is undecidable is that *at the present moment* there is no known decision procedure" (2015, p. 16).

At this point one may begin to wonder how effective procedures could be so relevant to the realism debate, in fact how the debate between the realist and the anti-realist could be over how to treat undecidables as we have defined them, when it is plain that our definitions of the realist and anti-realist notions of truth involve not effective procedures, but rather the stronger notion of unknowability. The easy answer is that the semantic realist and anti-realist disagree over whether truth is essentially knowable. And they disagree over how to treat undecidables because while they both agree on what undecidability is, they part ways when it comes to whether to apply bivalence to undecidables. The anti-realist does not think we are warranted in assuming that a sentence is either true or false when we have no guarantee that we will *ever* come to know it (despite their refusal to say that we are *incapable* of coming to know it). The realist thinks that truth comes completely apart from what we are able to know, and so is perfectly comfortable applying bivalence even to an undecidable. The harder answer is that, aside from the fact that anti-realists cannot work with a strong notion of unknowability, it is not entirely clear that the lack of an effective decision procedure has all of the relevance to the realism debate that Dummett has thought it to have.

One of the attractions of the effective procedure reading of undecidability is that it does not conflict with the anti-realist's knowability principle. Hence the anti-realist can accept it as a notion of undecidability. It was pointed out earlier that decidability on this reading entails knowability. Is the realist now unable to accept our form of undecidability, if the opposite of

that is decidability? Note that if this was going to be a problem, it was already going to be a problem with the reading of undecidable as unknowable, for on that reading it was even more straightforwardly the case that a decidable sentence was knowable. The thing is that it was not going to be a problem, because P's being knowable doesn't entail that, necessarily, if P then P is knowable. Hence the knowability of a given sentence, even its knowability due to its being effectively decidable, does not entail the knowability principle and does not contradict semantic realism. The realist can admit that there is an effective procedure for deciding P, while maintaining that P would have been the case regardless of whether we had any ability to know it. The anti-realist, on the other hand, maintains that P could not have been the case unless we had some ability to come to know it. This is exactly what we should expect, given our formulation of the knowability principle.

As I've construed Dummett's views, the realism debate is essentially over whether undecidables have realist truth conditions. Dummett thinks we should not apply bivalence to undecidables if they have realist truth conditions, and he does not see how they could. Without bivalence, we are not forced to accept the law of excluded middle for undecidables, and Dummett thinks we ought to reject it. It should be kept in mind that Dummett does endorse bivalence and the law of excluded middle for decidable sentences, but decidable sentences are not what Dummett's realism debate is about. Interestingly, the dispute between the semantic realist and semantic anti-realist is not discernible for decidables. One drawback of Dummett's approach is that the disagreement over what kinds of truth conditions *decidables* have never manifests itself in a way that is relevant to the realism debate. Since they are decidable, the realist and anti-realist agree to apply bivalence and the law of excluded middle to them. Any further disagreement falls by the wayside.

Dummett summarizes the sorts of features which make sentences undecidable:

> Many features of natural language contribute to the formation of sentences not in principle decidable: the use of quantification over an infinite or unsurveyable domain (e.g. over all future times); the use of the subjunctive conditional, or of expressions explainable only by means of it; the possibility of referring to regions of space-time in principle inaccessible to us. (1993, p. 46)

Given that the Dummettian realism debate only arises for undecidables, we are left to conclude that the realism debate is itself the product of the same features of natural language which give rise to undecidable sentences. This result is somewhat odd.

There are legitimate worries as to whether Dummett's treatment of the realism debate gets things right, or is even fruitful in its own right. Crispin Wright has canvassed some ways in which Dummett's approach is inadequate when it comes to capturing the essence of the realism debate in certain areas (Wright 1992). While Wright has carried on in the spirit of Dummettian anti-realism, seeking to improve upon Dummett's work, others have more forcefully rejected Dummett's overall approach. There are genuinely interesting questions to be asked of anyone who chooses to conceive of the realism debate in Dummett's terms. Thus in the end, though we will have good reason to think that Dummett's arguments against semantic realism and for semantic anti-realism do not work, we will be left with a fruitful way of thinking about the issues.

4. Dummett's Arguments, Again

In the literature surrounding Dummett's work there has arisen a consensus that Dummett provides two arguments against semantic realism, the manifestation argument and the acquisition argument. The manifestation argument is by far the most discussed, though the two arguments

have a great deal in common. It comes as no surprise that some have questioned how distinct the arguments really are. What ties the arguments together is an emphasis on the public use of language. Dummett nearly dogmatically associates such use with our capacity for epistemic appraisal. The world presents us with a situation, and we have an ability to determine whether a given expression is apt to be deployed in that situation. Learning to appropriately deploy expressions is what learning a language is, for Dummett. Demonstrating to others that we think particular expressions are, or are not, assertible under given circumstances is how we convey meaning and manifest understanding. Tied as our use of language is to our epistemic abilities, and given that use determines meaning, Dummett furthermore asserts that our sentences cannot have meanings which are not essentially tied to our epistemic abilities. The upshot of this idea is that our sentences cannot have meanings such that they might be true despite our inability to tell whether they were. The relevant lack of ability is supposed to be captured in the concept of undecidability as Dummett is understanding it, though as we will see it is not entirely clear how Dummett understands the concept. As we have seen, what makes this lack of ability meaning-relevant is contentious. It is an issue we will continue to explore once we have a firmer grip on the structure of Dummett's reasoning.

We have here three interrelated ideas: How we learn the meanings of expressions, how we convey those meanings to others, and how we endow expressions with meaning in the first place. All three of these ideas are prominent in Dummett's presentations of his anti-realist views, though only the first two have arguments named after them—the arguments from acquisition and manifestation, respectfully. We can call the reasoning associated with the third idea the *endowment argument.* In reality, all of these arguments are essentially based on the same central considerations: Meaning is use, and use is rooted in epistemic capacities.

The following passage clearly illustrates what I am calling Dummett's argument from endowment:

> The realist holds that we give sense to those sentences of our language which are not effectively decidable by appealing tacitly to means of determining their truth-values which we do not ourselves possess, but which we can conceive of by analogy with those which we do. The anti-realist holds that such a conception is quite spurious, an illusion of meaning, and that the only meaning we can confer on our sentences must relate to those means of determining their truth-values which we actually possess. Hence, unless we have a means which would in principle decide the truth-value of a given statement, we do not have for it a notion of truth and falsity which would entitle us to say that it must be either true or false. (1978, p. 24)

Here is another example, which the reader will recall from above. This example illustrates both the endowment and acquisition ideas:

> Constructivist philosophies of mathematics insist that the meanings of all terms… must be given in relation to constructions which we are capable of effecting, and of our capacity to recognise such constructions as providing proofs of those statements…The most powerful form of argument in favour of such a constructivist view is that which insists that there is no other means by which we can give meaning to mathematical expressions. We learn, and can only learn, their meanings by a training in their use…there is no means by which we could derive from such a training a grasp of anything transcending it, such as a notion of truth and falsity for mathematical statements independent of our means of recognising their truth-values. (1975, p. 301).

In his essay "The Reality of the Past," Dummett offers us the following statement of his acquisition argument:

> [The anti-realist] maintains that the process by which we came to grasp the sense of statements of the disputed class, and the use which is subsequently made of these statements, are such that we could not derive from it any notion of what it would be for such a statement to be true independently of the sort of thing we have learned to recognise as establishing the truth of such statements…In the very nature of the case, we could not possibly have come to understand what it would be for the statement to be true independently of that which we have learned to treat as establishing its truth: there simply was no means by which we could be shown this. (1978, p. 362)

And later,

> The only notion of truth for past-tense statements which we could have acquired from our training in their use is that which coincides with the justifiability of assertions of such statements, i.e., with the existence of situations which we are capable of recognizing as obtaining and which justify such assertions. (1978, p. 363)

Dummett's emphasis on the primacy of use, and the essential connection between use and our capacity for epistemic appraisal, runs unabated through his writings on realism, both in mathematics and elsewhere. Here is another example of his acquisition argument. The reader should keep in mind that Dummett takes what he says about mathematics to be more or less directly transferrable to the non-mathematical realm:

> When we learn…the language of a mathematical theory, what we learn to do is to make use of the statements of that language: we learn when they may be established by computation, and how to carry out the relevant computations, we learn from what they

> may be inferred and what may be inferred from them, that is, what role they play in mathematical proofs and how they can be applied in extra-mathematical contexts, and perhaps we learn also what plausible arguments can render them probable. These things are all that we are shown when we are learning the meanings of the expressions of the language of the mathematical theory in question, because they are all that we can be shown. (1975, p. 99)

Here is a classic presentation of Dummett's manifestation requirement:

> The meaning of a mathematical statement determines and is exhaustively determined by its *use.* The meaning of such a statement cannot be, or contain as an ingredient, anything which is not manifest in the use made of it, lying solely in the mind of the individual who apprehends that meaning…an individual cannot communicate what he cannot be observed to communicate. (1983, p. 98).

The requirement plays an essential role in the manifestation argument, where Dummett contends, in essence, that we cannot be observed to communicate realist truth conditions, since realist truth conditions transcend all that is determined by use, where use is understood as essentially tied to our capacities for epistemic appraisal. So while it is interesting, and perhaps useful for understanding, to separate these three threads in Dummett's line of thinking (acquisition, manifestation, endowment), ultimately what matters is that, for Dummett, facts about our epistemic capacities as humans combined with truisms about language and use make it the case that our sentences can only have anti-realist, and not realist, truth conditions. Let us now turn to an analysis of Dummett's arguments.

4.1 Acquisition and Endowment

One of Dummett's favorite examples of an undecidable sentence involves a man, Jones, who never had a chance to act bravely in his life and who is now dead. If we wonder, now, whether 'Jones was brave' is determinately either true or false, Dummett thinks we should say that it is not, that we should reject a realist semantics for that sentence. For the truth or falsity of that sentence will come down to the truth of one of a pair of subjunctive conditionals: 'If Jones had encountered danger, he would have acted bravely' and 'If Jones had encountered danger, he would not have acted bravely.' And what evidence could possibly be available to us, Dummett says, for either one of these conditionals? None, he thinks. Or, at the very least, certainly not enough to make either one of them effectively decidable. So, then, to think one of the conditionals true, and hence to think that 'Jones was brave' has a determinate truth-value, we must violate, "the principle that a sentence can be true only if there is something in virtue of which it is true" (1978, p. 23).

Dummett's invocation of the principle is somewhat stunning, as he eventually realized, since he after all says that it is a central tenet of realism (1978, p. 14), not to mention that the principle has a distinctive realist ring to it. In a postscript to the essay in which he initially laid out the problem, Dummett attempts to allay our worries, clarifying first that he meant we would be violating the principle that a sentence cannot be true if there is, "nothing of the sort which we ordinarily use as evidence for the truth or falsity of such a sentence," and second that the weaker principle, "that a sentence cannot be true unless it is in principle capable of being known to be true," is what should really have been under consideration (1978, pp. 23-24). This latter claim is itself stunning, as Dummett seems to be saying that even a realist must admit that truth is evidentially constrained. But this is not what he is saying. Dummett claims that *both* the realist

and the anti-realist would accept the latter principle, though they would each draw different conclusions from it. He says, "the anti-realist uses it to infer that the sentence ['Jones was brave'] is not necessarily either true or false; the realist uses it to infer that that which makes it true or false cannot be identified with that by which we recognise it as true or false, when we are able to do so." (1978, p. 24) Dummett is also careful at this point to note that, with regard to the latter principle, "the anti-realist interprets 'capable of being known' to mean 'capable of being known *by us*', whereas the realist interprets it to mean 'capable of being known by some hypothetical being whose intellectual capacities and powers of observation may exceed our own'" (p. 24). The idea, it would appear, is that if a sentence is true, then there must be *some* logically possible being, an omniscient one, perhaps, who could acquire evidence for the sentence and come to know it. As long as we are not restricted in our imagining of what sorts of epistemic powers this being might have, the idea is innocuous enough.

According to Dummett, both parties should agree that the subjunctive conditionals under discussion above are not effectively decidable. The question is what conclusion should be drawn from this fact. Another one of Dummett's favorite examples involves sentences about the past. Sentences about the past would count as undecidable in virtue of referring to "regions of space-time in principle inaccessible to us" (Dummett 1993, p. 46). Ultimately, Dummett believes we should give up a realist semantics for a verificationist alternative. His proposal is that,

> *We no longer explain the sense of a sentence by stipulating its truth-value in terms of the truth-values of its constituents, but by stipulating when it may be asserted in terms of the conditions under which its constituents may be asserted.* The justification for this change is that this is how we in fact learn to use these sentences. (1978, 17-18)

Dummett unabashedly mimics the mathematical intuitionists both in giving his reasons for rejecting the realist view and in laying out his anti-realist alternative, as he admits without reservation: "What I have done here is to transfer to ordinary sentences what the intuitionists say about mathematical sentences" (1978, p. 17). Dummett thinks that, in general, our grasp of the meaning of a sentence cannot amount to a grasp of the condition under which the sentence is true, because for those sentences in particular which are not effectively decidable, we never could have come to acquire the realist notion of truth which is supposed to apply to them, and they never could have been endowed with the kind of meaning the realist thinks they have.

While numerous extant reconstructions of Dummett's argumentation contain many steps, his arguments are best understood simply, as instances of *modus tollens* which conclude with the rejection of realist truth conditions for undecidables. We have already seen the manifestation argument. Here it is again with the others.

The Manifestation Argument

M1. If undecidable sentences have realist truth conditions, then we are able to manifest a grasp of those conditions.

M2. We are not able to manifest a grasp of the realist truth conditions of undecidables.

M3. Therefore, undecidable sentences do not have realist truth conditions.

The Acquisition Argument

A1. If undecidable sentences have realist truth conditions, then we are able to learn the realist truth conditions of undecidable sentences.

A2. We are not able to learn the realist truth conditions of undecidable sentences.

A3. Therefore, undecidable sentences do not have realist truth conditions.

The Endowment Argument

E1. If undecidable sentences have realist truth conditions, then we can endow undecidable sentences with realist truth conditions through the way that we use them.

E2. We cannot endow undecidable sentences with realist truth conditions through the way that we use them.

E3. Therefore, undecidable sentences do not have realist truth conditions.

The first premise of each argument always has the same character—it presents something that Dummett takes to be a fact about language in general. He thinks that we must be able to manifest in use our knowledge of meaning (which Dummett sees as understanding), we must be able to learn the meanings that our words have by observing others' use of them, and that we give meanings to the words of our language through the way that we use them. Each of these alleged facts can be seen as placing a particular requirement on the theory of meaning, a requirement that Dummett claims is impossible to meet in the case of undecidables with realist truth conditions. The second premise of each argument simply tells us that the requirement cannot be met. Dummett's reasons for thinking that the requirements cannot be met are obscure—he often seems to think it obvious. The obscurity is exacerbated by Dummett's use of the concept of undecidability, which as we have seen is somewhat obscure itself. In what follows we will continue to work towards clarifying Dummett's reasoning.

4.2 Manifestation

Little in Dummett's writings has vexed philosophers so much as his manifestation argument for semantic anti-realism. The argument smacks of behaviorism and positivism. Even if one gets over that, it can still be difficult to see why exactly Dummett insists that knowledge of

meaning must be fully manifested (or at least manifestable) in behavior. Furthermore, even if one is able to come to terms with all of this, one is left with the feeling that Dummett has simply begged the question against his realist opponents, building an essentially anti-realist requirement into the theory of meaning.

Dummett's goal is to construct an adequate theory of meaning for our language, one which specifies what it is that we know when we know the meanings of our words. As a theoretical apparatus, a theory of meaning for a language should of course be informative to those studying the language, as we are, but a theory of meaning as Dummett intends it is much more. A theory of meaning in Dummett's sense is what we know when we know a language, and it is in virtue of two speakers sharing a theory of meaning for a language that they are in fact able to communicate. This much is clear from Dummett's writings:

> A speaker's mastery of his language consists, on this view, in his knowing a theory of meaning for it: it is this that confers on his utterances the senses that they bear, and it is because two speakers take the language as governed by the same, or nearly the same, theory of meaning that they can communicate with one another by means of that language. (1993, pp.100-1)

Dummett often says that knowledge of meaning is best understood as a "practical ability" to use the language, and that the sort of knowledge one has when one knows a language should be seen as, "*implicit* knowledge: knowledge which shows itself partly by manifestation of the practical ability, and partly by a readiness to acknowledge as correct a formulation of that which is known when it is presented," though the speaker would not necessarily be capable of explicitly formulating what it is that is known (1993, 96). What we as theorists are attempting to do is to

model the competent speaker's linguistic competence (see Khlentzos 2004, p. 51, for a discussion of modeling of this sort):

> What [the competent speaker] has when he knows the language is practical knowledge, knowledge how to speak the language: but this is no objection to its representation as propositional knowledge…what we seek is a theoretical representation of a practical ability. (Dummett 1993, p. 36)

Even though Dummett typically sees linguistic competence as a sort of practical ability, one for which we can produce a theoretical representation, he nevertheless has said that we should conceive of linguistic competence as a species of knowledge, knowledge which must be implicit: "A theory of meaning will, then, represent the practical ability possessed by a speaker as consisting in his grasp of a set of propositions…the knowledge of these propositions that is attributed to a speaker can only be an implicit knowledge" (1993, p. 36).

In the preface to *The Seas of Language,* Dummett recants his idea that an adequate theory of meaning should hold linguistic knowledge to be implicit, since a speaker's willingness to acknowledge the propositions of the meaning theory as correct, "tells us nothing about how the knowledge is applied when the occasion for its application arises: it therefore fails to explain what the philosopher seeks an explanation for" (1993, p. xi). Dummett furthermore is careful to explain that linguistic knowledge shouldn't be seen as a "pure practical ability," and neither as, "explicit knowledge of a theory of meaning," but rather as, "a species of knowledge intermediate between pure practical knowledge and pure theoretical knowledge" (1993, p. x).

It is disappointing to see Dummett remove the notion of implicit knowledge from its prized place in his theory of meaning without offering a replacement, or at least a stand-in. This is partly because he has gotten so much use out of the concept in advancing his overall project

with regard to the theory of meaning. It is also too bad in light of Michael Devitt's accusation that Dummett endorses the view that linguistic knowledge is propositional *knowledge-that* rather than practical *knowledge-how*, given that the conception of linguistic knowledge as implicit knowledge of a meaning theory offers the Dummettian a plausible way around the objection (Devitt 1997). It is too bad, again, given that Mark Gardiner's and Drew Khlentzos' otherwise rather convincing (and separate) defenses of Dummett's manifestation requirement make essential use of the claim that knowledge of meaning is implicit (Gardiner 2000, Khlentzos 2004). As both authors point out, and as we will see, Dummett himself plainly says that manifestation is required when implicit knowledge is at issue. We lose that justification for Dummett's manifestation requirement if we turn our backs on the concept of implicit knowledge. Thankfully, Dummett himself denies that "anyone is an authority on whether his present views are closer to the truth than his earlier views" (1993, p. xii), and we will be at liberty to improve upon his views as we see fit.

All the same, Dummett proclaims in the very same preface, "I continue to believe that any knowledge attributed to a speaker as constituting a component of his knowledge of a language must be manifested in his employment of that language" (1993, p. xii). There are two aspects of Dummett's views on this matter which have continuously baffled his critics and which we presently ought to address: First, Dummett's insistence that any adequate account of meaning must explain how knowledge of meaning is manifested in linguistic behavior, and, secondly, his insistence that such knowledge must be *fully* manifested in said behavior. We will approach the first one first, since the explanation of why Dummett thinks linguistic knowledge requires manifestation at all leads simply and directly into an explanation of why he thinks that the knowledge must be manifested fully.

> Here is perhaps Dummett's most concise sentence of his manifestation requirement: A meaning-theory, being a theoretical representation of a practical ability, must not only say *what* a speaker must know in order to know the language, but in what his having that knowledge consists, that is, what constitutes a manifestation of it. (1993, pp. 115-116. See also p. 37)

A notable aspect of the above quotation is that in it, Dummett appears to equate a statement of what linguistic knowledge consists in with a statement of what counts as a manifestation of that knowledge. Perhaps Dummett thinks that an explanation of what is known, given in terms of how that knowledge is manifested, is always a sufficient explanation, or perhaps he only thinks that such an explanation is sufficient for particular cases. Regardless, it is clear that he thinks the speaker's knowledge must be explained in some way. It is not enough to simply specify what is known—we must also say something about what it is to have such knowledge. With regard to the central question of whether realist truth conditions are what competent speakers grasp when they grasp the meaning of a sentence, Dummett says that, "a theory of meaning which takes [realist] truth as its central notion has to supply an explanation of what it is to ascribe to someone a knowledge of the condition which must obtain for a sentence to be true" (1993, p. 44). As we will see, Dummett doesn't think that the realist can provide such an explanation, but for the time being we can simply take note that merely ascribing the knowledge isn't enough for Dummett—we must also say in what such an ascription consists.

So far we have been told that any meaning theory worth its salt will tell us "what it is to ascribe" knowledge of meaning to someone, and also what it is to have such knowledge. Presumably these are just two different ways of getting at the same thing: If we can explain what it is to have knowledge of meaning, then we can explain what we are ascribing to someone when

we say of that person that he knows what his words mean, and vice versa. Dummett is also clearly concerned with what justification we can possibly have for ascribing knowledge of meaning to a speaker, as when he says, "A speaker, to be credited with an understanding of the expressions of a language, must manifest his grasp of connections between them" (1993, p. ix). Some commentators have suggested that the issue here is that ascriptions of knowledge when there is no manifestation will be "vacuous" (Kirkham 1989, p. 215. Khlentzos 2004, p. 66). Kirkham at another point says the question is whether we can, "meaningfully ascribe" knowledge of meaning to a speaker (1989, p. 216). Talk of vacuity and meaninglessness issues in thoughts of the behaviorism and positivism which we are trying to avoid. Dummett's talk sounds dangerous, but I do not believe that it is, or, if it is, then we do not need to carry on his project in a dangerous way. The key notion, I would like to suggest, is simply that language is a vehicle of communication. In order for a speaker to communicate something to us, we must be able to witness him communicating it. This requirement holds as much in normal conversation as it does when we, as semanticists, are laying down the strictures for an adequate theory of meaning. We must specify what it would take to justifiably ascribe knowledge of meaning to a speaker. So, I want to suggest, the idea that knowledge of meaning must be manifestable in use is not in itself as objectionable as many of Dummett's commentators have found it. What may be so objectionable, however, is the restricted way in which Dummett conceives of language use as tied to epistemic capacities.

Elsewhere, Dummett outright says that it is the fact that the knowledge is implicit which creates the demand for an account of what understanding consists in based on manifestation (1993, pp. 37, 45). Given what Dummett has said more recently about the theoretical inadequacy of ascribing implicit knowledge to a speaker, along with the fact that he still sees the

manifestation requirement as holding (see above), we ought not assume that the requirement is dependent upon the picture of linguistic competence as implicit knowledge. And in fact there is no need to, as Dummett gives us a stronger reason which applies regardless of whether we give up on implicit knowledge, and which makes perfect sense of his argument from manifestation. Dummett's requirement of manifestation grows entirely out of his conception, borrowed from Frege, of meaning as essentially communicable.

That Dummett adverts to the essential communicability of meaning, and even that he does so in support of the manifestation requirement, is by no means lost on his commentators. In my view, however, not enough has been said to promote the idea that Dummett bases his manifestation requirement entirely on the thesis that meaning is essentially communicable, and that this is, furthermore, all the support the requirement needs (essentially the same view I am promoting here can be found in Khlentzos 2004, p. 56, though see below for brief discussion). On this point, Dummett's response to a possible objection to his view concerning implicit knowledge is particularly illuminating. In the relevant passage, Dummett is considering the idea that what he has said about a speaker's linguistic knowledge being implicit falls prey to Frege's argument against psychologism, that is, "the explanation of sense in terms of some inner psychological mechanism" (Dummett 1993, p. 102). Dummett goes on to outline Frege's argument, focusing ultimately on

> the assumption, which is, indeed, required if we are to be able to communicate by means of our utterances, that we are talking the same language, a language that we both understand: but that in which our understanding of the language consisted would lie open to view, as Frege maintained that it does, in our use of the language, in our participation in a common practice. (Dummett 1993, p. 102)

The argument is basically that, if meaning were just some inner mechanism, we would not be able to communicate with our language because we could not be sure that we meant the same things by our words. Dummett then points out that the very same argument could be used against his view that speakers have implicit knowledge of the theory of meaning for their language. Luckily, Dummett has a way around this, for, as he explains,

> I said earlier that implicit knowledge ascribed to the speakers must be manifested in their *use* of the language, and that it is part of the business of a philosopher of language to explain in what specific feature of this use a speaker's knowledge of each particular part of the theory of meaning is so manifested. (1993, p. 102)

So it is because meaning is essentially communicable that language is to be seen, "not just as a means of expressing thought, but as a *vehicle* for thought" (Dummett 1993, p. 99), something which contains the thoughts we wish to convey, that Dummett requires that the implicit knowledge of the meaning theory must be manifested in the use speakers make of the language. But the essential communicability of meaning applies to more than just linguistic competence understood as implicit knowledge of meaning. I contend that Dummett intends it to apply to linguistic competence understood in any way whatsoever. Even if he does not, such a route is open to us as we assess the viability of Dummett's overall project.

Dummett adduces similar considerations in defense of his claim that speakers must not only manifest whatever linguistic knowledge they can be said to have, but must also manifest such knowledge fully. His remarks on the issue are helpful not only for understanding his justification for the requirement, but also the precise nature of the requirement. The basic problem with requiring full manifestation, as Dummett seems to understand, is that I may attach a meaning to an expression which, under the right circumstances, would be manifested in my use

of the expression but which, as things actually turn out, is never fully manifested. So Dummett considers the suggestion of

> weakening the requirement thus: any difference in the meanings attached to an expression by two speakers must be *capable* of being manifested by some differences in the uses they make of it. That is to say, there are things they *might* say that will manifest the difference; but there is not in general any guarantee that they *will* say them, in which case the difference in their understanding of the expression will never come to light. (1993, p. xiii)

Dummett then goes into a technical example from an earlier publication on Gödel's theorem, and concludes that,

> to make understanding rest upon an inner mental grasp of a mathematical structure that can never be made fully explicit is to render it incommunicable in just that way objected to by Frege in psychologistic accounts of understanding: that is why I made the stronger claim that understanding must be capable of being *fully* manifested in use. (p. xiv)

Dummett's remarks here are confusing. His example was supposed to explain why it is not enough that knowledge of meaning be merely *capable* of manifestation through verbal behavior, but he himself concludes his counter-argument with talk of understanding (that is, grasp of meaning) being capable of being manifested. He then goes on, appearing to play directly into the hands of his objectors:

> It is plain that it cannot be demanded that every aspect of a speaker's understanding of an expression should in fact be manifested by something he says or does; the most that can be required is that there is something he would say or do, should the occasion arise, that would manifest it. It is not enough, however, to require, of every aspect of his

> understanding, that it be capable of being manifested. Rather, we are entitled to require that it be possible that every aspect of his understanding should have been manifested; for, if not, there can never be conclusive evidence for the attribution to him of any specific understanding of the expression. (p. xiv)

So what is going on here? Does Dummett agree with his objectors or not? Let us distinguish two possible objections to which Dummett could be responding. Dummett, it seems, is responding to one very precise form of the objection to his full manifestation requirement. That is the form which says that, even under appropriate circumstances, there is no guarantee that I will manifest my linguistic knowledge. That knowledge is merely capable of being manifested, and it may or may not be, so the objection goes. Dummett's response is that we *must* think that, under the appropriate circumstances, (assuming, it must be said, that I am engaging in any linguistic behavior at all), my linguistic understanding *would* be manifested by the use I in fact make of the language. My actual linguistic behavior on those occasions must manifest the meaning I actually attach to the expressions I use, and any meaning I actually attach to the expressions I use must be capable of being manifested in the use I make of them. There are subtleties to be worked out here, but I believe what Dummett is resisting is the idea that I might actually engage in linguistic behavior, even a great deal of it, and yet systematically fail to manifest the meaning I attach to my expressions because, after all, that meaning is only *capable* of being manifested. To think this, Dummett would say, is to think that meaning is not essentially communicable, that language is not the vehicle of our thought. As we've seen, he finds this conclusion unacceptable.

The other way of understanding what Dummett's objectors are saying is simply this: There is no guarantee that all of my linguistic knowledge will ever be manifested, just because I

may never be in a position to manifest it, or I may choose not to, etc. And this much, it seems, Dummett would agree with, but he does not see the objection to be damaging. We should not see it as damaging either. In order for his argument from manifestation to go through, Dummett does not need it to be the case that every speaker actually finds occasion in his life to manifest all of the linguistic knowledge he possesses. The above attempt constitutes the only sense I can make of Dummett's remarks concerning objections to his full manifestation requirement. Beyond that, the only recourse I can see is concluding that what he has said in print on the point is in fact not coherent and even contradictory.

So, I conclude, Dummett's reasons for requiring full manifestation are just the same reasons for which he requires manifestation at all. For Dummett, a speaker cannot have a grasp of the meanings of his words which is not potentially manifestable in his linguistic behavior. We can understand the requirement this way: A speaker who was sincerely attempting to express himself with his words (barring cases where he is not of sound mind, drugged, etc.) would be able to demonstrate whatever grasp of those words he had through his use of them. Perhaps he would express himself poorly at first, choosing words which he then saw fit to take back, etc., but it could not be the case that, despite his use, he still attached meanings to the words which we could not, at least in principle, read off of the way that he used them. As Dummett continuously harps, the speaker's use must give us all the evidence we need to ascribe a particular grasp of meaning to him. In fairness, Dummett's writings provide some evidence that he does not see things exactly the way I have set them up here. His remarks concerning advances he takes himself to be making on Frege with regard to implicit knowledge being the source of the manifestation requirement are a case in point (1993, p. 85). But we have already seen that Dummett clearly no longer thinks the manifestation requirement depends upon his conception of

implicit knowledge. And at any rate, our ultimate goal here is the best version of Dummett's views, even if we must go beyond what Dummett would have, at any point in his long career, agreed to.

As the foregoing discussion makes clear, the manifestation argument is concerned with what a speaker's knowledge of meaning amounts to, that is, what an ascription of knowledge of meaning to the speaker consists in, and what could possibly justify us in attributing to him a grasp of a particular kind of meaning. The question is going to be: Could what a speaker knows (or understands) when he knows (or understands) the meaning of a sentence be that sentence's realist truth condition? And Dummett's answer is going to be: No.

Dummett is arguing that a realist theory of meaning based on truth conditions is not viable. He concludes from there that our sentences cannot have realist truth conditions, and hence that semantic realism is false. Since a central tenet of semantic realism is that our sentences have truth conditions which objectively obtain independently of whether we are able to determine whether they obtain or not, Dummett calls the difficulty facing the realist, "the difficulty about a theory of meaning based on the notion of [realist] truth which arises from the fact that the truth of many sentences of our language appears to transcend our powers of recognition" (1993, p. 52). The sentences Dummett has in mind, of course, are undecidables.

Dummett points out that there is "no difficulty in stating what constitutes a speaker's knowledge of the condition for the truth of a sentence" in certain cases, and the cases he has in mind are those such that "the condition in question is one which [the speaker] can be credited with recognizing whenever it obtains" (1993, p. 45). In such cases, we can just say that the speaker's knowledge of the truth condition "consist[s] in his capacity… to evince recognition of the truth of the sentence when and only when the relevant condition is fulfilled" (p. 45). Still,

Dummett recognizes that such cases will be few and far between, so he offers a generalization of the approach which can be applied to effectively decidable sentences generally. The effectively decidable sentences are just those sentences "for which a speaker has some effective procedure which will, in a finite time, put him into a position in which he can recognize whether or not the condition for the truth of the sentence is satisfied" (p. 45). The generalization of the approach is this: We can say for any such effectively decidable sentence, "that the speaker's knowledge of the condition for it to be true consists in his mastery of the procedure for deciding it, that is, his ability, under suitable prompting, to carry out the procedure and display, at the end of it, his recognition that the condition does, or does not, obtain" (pp. 45-46).

Notice the terminology Dummett is employing here. He has said that knowledge of the truth conditions will be something the speaker can "evince" and "display," something the speaker can be "credited with recognizing." This is no accident. For recall that, whatever we say the speaker's knowledge consists in, it must be knowledge that we could in principle be justified in ascribing to him based solely on his use of the language.

Given that we have a general approach for attributing knowledge of realist truth conditions to speakers, specified so far for all effectively decidable sentences of the language, what exactly is the problem for the realist? As Dummett tells us, the problem is that, "natural language is full of sentences which are not effectively decidable, ones for which there exists no effective procedure for determining whether or not their truth-conditions are fulfilled" (p. 46). The culprit sentences are the same sentences that caused problems when it came to explaining how we could ever come to acquire a realist notion of truth and endow our sentences with realist truth conditions. The issue is not that we cannot come to be justified in believing that an undecidable sentence is true. It is, rather, that, "for such a sentence, we cannot equate a capacity

to recognize the satisfaction or non-satisfaction of the condition for the sentence to be true with a knowledge of what that condition is…because, by hypothesis…it may obtain or fail to obtain," independently of our being able to ascertain whether it does (p. 46). This matters because it means that, for the undecidable sentences, we cannot extend the analysis given above of what a competent speaker's understanding of the truth conditions of the effectively decidable sentences consists in.

Dummett continues:

> Whenever the condition for the truth of a sentence is one that we have no way of bringing ourselves to recognize as obtaining whenever it obtains, it seems plain that there is no content to an ascription of an *implicit* knowledge of what that condition is, since there is no practical ability by means of which such knowledge may be manifested. An ascription of the knowledge of such a condition can only be construed as *explicit* knowledge, consisting in a capacity to *state* the condition in some non-circular manner; and that, as we have seen, is of no use to use here. (p. 46)

It is clear why Dummett's proposed approach fails when it comes to sentences which are not effectively decidable. The extension of his original proposal required the possession of an effective procedure which would allow the speaker to recognize when the truth condition of a sentence had been fulfilled and when it had not. For the undecidables, of course, the speaker possesses no such procedure, so the approach cannot be extended to cover them. This is the essence of his manifestation argument against semantic realism. The thrust of the argument is that a realist meaning theory in terms of evidence-transcendent truth conditions cannot satisfy the manifestation requirement, which Dummett thinks any plausible theory of meaning will adhere to, as there will be sentences, the undecidables, for which speakers could never manifest their

grasp of the appropriate truth conditions. Thus, Dummett concludes, competent speakers of our language cannot, in general, be said to know realist truth conditions, and hence semantic realism must be false.

One immediate concern is why Dummett thinks that, even if we give up on thinking of knowledge of the truth conditions of undecidables as implicit knowledge, we cannot see such knowledge as explicit, that is, as "knowledge which is manifested by [the speaker's] ability to state [the truth condition]" of the sentence (1993, p. 45). He says in the above passage that, "as we have seen," such knowledge, "is of no use to use here." In fact, we have not been shown that. Dummett's discussion up to the passage in question has indeed explained how knowledge of meaning cannot, in general, be explicit knowledge, as then the explanation of what our knowledge of meaning is would be circular (p. 45). Commentators such as Khentzos and Gardiner, both of whom, in their reconstructions of Dummett's argument, make important use of the idea that the knowledge in question cannot be explicit, cite nothing other than the passage above itself, which clearly only tells us that we have already seen how it cannot be explicit.

It is worth noting that explicit knowledge of meaning already meets Dummett's manifestation requirement as a matter of course, since that knowledge would be directly manifestable in the speaker's ability to state what that knowledge is. So while Gardiner (2000) and Khlentzos (2004) see the manifestation requirement as falling out of what Dummett says about, on the one hand, the inadequacy of construing linguistic knowledge in general as explicit knowledge, and on the other, what he says about requiring manifestation for implicit knowledge, it turns out that manifestation is required for both sorts of knowledge. It is just that the requirement is trivially met in the case of explicit knowledge, as becomes apparent if we revisit a passage of Dummett's we have already been over, albeit in pieces:

> A theory of meaning which takes truth as its central notion has to supply an explanation of what it is to ascribe to someone a knowledge of the condition which must obtain for a sentence to be true. If the sentence is of a form which a speaker can come to understand by means of a verbal explanation, then there is no problem: his knowledge of the truth-condition of the sentence is explicit knowledge, knowledge which is manifested by his ability to state that condition. (1993, pp. 44-45)

Dummett's reasoning is based on a conception of use which leads almost immediately to Dummett's anti-realist conclusions. Many have contended that his manifestation requirement suffers from either an unacceptable behaviorism, or question-beggingness, or both. It is important to see that, while both accusations may be fair, any purported behaviorism would not be the source of any alleged question-begging. If anything, a positivistic behaviorism would amount to a form of anti-realism about meaning, or perhaps about understanding or knowledge of meaning. This kind of view may be damaging to Dummett, but does not beg the question against the realist when it comes to the conclusion Dummett is actually attempting to establish, that is, anti-realism about truth conditions.

The real point at which Dummett may have rigged his reasoning against the realist is in his insistence that the meaning of our sentences must be determined, via our use of them, by our capacities for epistemic appraisal. There is nothing about the requirement that meaning be manifested in use, taken all by itself, which would lead to Dummett's conclusion. Rather, it is Dummett's commitment to the idea that our use of language is inherently epistemological that makes for the difficulty. Those aspects of use pertaining to our ability to determine whether certain assertions would be warranted are paramount for Dummett. Thus it is difficult to see

how Dummett's reasoning is supposed to work against a semantic realist who doesn't already accept, with the anti-realist, that truth is epistemically constrained.

> Neil Tennant characterizes Dummett's manifestation requirement as follows:
>
> **MR**$_T$: Understanding is fully manifestable in the public exercise of recognitional skills. (1997, p. 177).

Once one accepts that *recognitional skills* are what matter not only to the manifestation of understanding, but also to the acquiring of understanding, and to the giving of meaning to expressions, it is plain how one might become concerned as to how we could be said to understand undecidable sentences, learn what they mean, and in fact give them the meanings they are purported to have. Crispin Wright puts the problem this way:

> If truth in general is evidentially unconstrained, then—depending on its subject matter—knowing the truth-conditions of a sentence may require an understanding of how it could be undetectably true. And how could that knowledge consist…in any ability whose proper exercise is tied to *appreciable* situations? How can knowing what it is for an unappreciable situation to obtain be constituted by capacities of discrimination exercised in response to appreciable ones? (1993, p. 248).

What is not so plain is why a realist would accept such a strong connection between meaning and epistemic capacities. Tennant even contends that the manifestation requirement amounts to little more than the requirement that all truths be knowable (1997, p. 182). In other words, the manifestation requirement is essentially the knowability principle in different garb.

Edward Craig lambasts Dummett's overall approach, saying that, "one will get philosophically important conclusions out of a theory of meaning only in proportion as one feeds philosophically controversial assumptions into it" (1982, p. 564). Indeed, Dummett has built some highly

controversial assumptions into his reasoning. Still, to accuse him of begging the question against the realist would be a somewhat exaggerated response. Dummett has not outright assumed the falsity of semantic realism as a premise in his argument against semantic realism. What he has done, rather, is attempt to show that the onus is on the realist to defend the view that our sentences can have realist truth conditions while respecting what Dummett takes to be incontrovertible facts about language use. Those facts lead very quickly to the rejection of realism and the endorsement of anti-realism.

Alexander Miller, in his rich study of the manifestation argument, distinguishes what he calls "strong" and "weak" versions of the argument. According to the strong version, semantic realism must be rejected because it fails the manifestation requirement. According to the weak version, there is nothing manifest in the use speakers can make of sentences which requires that those sentences be understood as having realist truth conditions *rather than anti-realist truth conditions* (2002, p. 360). Miller argues that the weak version of the manifestation argument is not only unmotivated to begin with, but also answerable by the realist. He lets the strong version of the argument fall by the wayside since, as he sees it, anti-realists have as a matter of fact weakened the strength of their claims in the face of realist criticism—criticism which Miller implicitly accepts.

The three versions of Dummett's argument I laid out above should all be understood in Miller's strong sense. Their aim is to establish the unacceptability of semantic realism. Miller's attempts to undermine the weak manifestation argument threaten stronger versions of Dummett's arguments as well. Miller begins with some considerations due to John McDowell:

> There is a truistic connection between the notion of the content of an assertion and a familiar notion of truth…the connection guarantees, as the merest platitude, that a correct

> specification of what can be asserted, by the assertoric utterance of a sentence, cannot but be a specification of a condition under which the sentence is true. (McDowell 1981, p. 229).

McDowell goes on:

> Now if a sentence lacks an effective decision procedure, then the condition which any competent speaker knows that he would be asserting to obtain if he used the sentence in order to make an assertion…is *ex hypothesi* not a condition whose obtaining…a competent speaker can be sure of being able to put himself in a position to recognize. Thus, without lapsing into psychologism, we seem to have equipped ourselves with a kind of realism: a description of linguistic competence which makes central use of the idea that speakers have a knowledge of conditions which they are not, in general, capable of recognizing whenever they obtain. (1981, p. 231).

McDowell's idea, as Miller emphasizes, is that if we understand undecidables, which all parties agree that we do, then we understand the sorts of conditions which would have to obtain in order for us to be justified in asserting an undecidable sentence, and since, in the case of an undecidable, it is understood that the condition in question may obtain though we may never come to recognize that it does, in that sense the sentence has, and we grasp, an evidence-transcendent truth condition. But to admit that our sentences have evidence-transcendent, i.e. realist, truth conditions is just to admit realism for those sentences.

The main trouble here is with the notion of evidence-transcendence, which is currently being deployed in such a way that it coincides with undecidability. Since anti-realists are also supposed to be able to admit the existence of undecidables, such a rendering of evidence-transcendence is not felicitous. Recall that our notion of undecidability has to do with the lack of

effective procedures, and not the strict lack of an ability to know. Speaking of undecidability as unknowability is only acceptable if one is very careful to circumscribe the precise way in which lack of an effective procedure places constraints on our ability to know. Crispin Wright sums up what he thinks we ought to say in response to McDowell's line of thinking:

> It is therefore quite consistent to hold that S is not effectively decidable, that we understand its assertoric content to be that conditions obtain which are necessary and sufficient for its truth, and that we do *not* understand what it would be for those conditions to obtain undetectably…somehow McDowell has—perfectly question-beggingly—run together lack of effective decidability with the capacity to be undetectably true. (1993, p. 19n)

Miller thinks that he has caught Wright missing the point. He claims that Wright is making too much of the idea that "we do *not* understand what it would be for those conditions to obtain undetectably." Miller asserts that what matters is not whether we understand what such undetectable obtaining would be, but whether what we grasp is a truth condition which might obtain undetectably (p. 369). In my view, Miller makes too much of Wright's phrasing. Wright is essentially correct that realism is not to be had so easily.

Miller says that, given any particular sentence for which we lack an effective decision procedure, and given the fact that "we understand the assertoric content of [the undecidable sentence] to be that conditions obtain which are necessary and sufficient for its truth," the anti-realist cannot, while remaining consistent, assert that, "it is not the case that our understanding of [the undecidable sentence] consists in our grasp of truth-conditions which, if they obtain, we may be incapable, even in principle, of detecting" (pp. 368-369). The reason is supposed to be that the undecidable, whose truth conditions we do grasp, is just the sort of sentence whose truth

conditions "we may be incapable, even in principle, of detecting." The essence of the reasoning here is that if we grasp the truth conditions of undecidables, then what we grasp must be realist truth conditions, since by definition an undecidable is a sentence which might "be true without our being able to recognise that it is so" (Miller 2002, p. 369).

In response to Miller and McDowell, it should first be pointed out that the truth condition which an anti-realist would take an undecidable to have would be an anti-realist truth condition such that the sentence could not be true unless we were capable, at least in principle, of coming to know its truth. McDowell may be correct to point to a platitudinous connection between assertion and truth conditions, but he does not get to say that those truth conditions must themselves be realist in character. Furthermore, it is worth noting that, precisely because the anti-realist takes the truth conditions of our sentences to be anti-realist in character, the anti-realist is not willing to say that an undecidable is assertible, because its undecidability tells us that we do not know whether we will come across a proof of either it or its negation. At any rate, what an anti-realist would be asserting by asserting a sentence with an anti-realist truth condition would be that the sentence is knowably true. This is decidedly not the same as asserting the truth of a sentence whose truth condition, "we may be incapable, even in principle, of detecting."

Finally, we have witnessed yet again the conflation of undecidability, where that is taken to entail the lack of an effective decision procedure, with unknowability in principle. Return to the alleged inconsistency Miller points to. It is claimed that an anti-realist's denial of semantic realism, that is, her denial of the view that our sentences have realist truth conditions, contradicts her claim that we grasp the truth conditions of undecidables. The anti-realist does think that we can grasp the truth conditions of undecidables, but those truth conditions are understood anti-realistically. The anti-realist is only in trouble here if undecidability is equated with the sort of

evidence-transcendence we see in our definition of evidence-transcendent truth. But that is not how an anti-realist understands undecidability, and it is also not how the realist, or anyone, should understand undecidability.

5. The Anti-Realist Alternative

It is worth noting again that Dummett's program has a positive component and a negative component. Practically all of what has been covered so far relates mainly to the negative aspect of his work. So far we have seen why he thinks that a realist theory of meaning based on truth conditions cannot work. There is a tendency in the literature to run Dummett's destructive arguments against the possibility of a realist meaning theory together with his constructive arguments in favor of an anti-realist alternative. To be fair, he tends to give them both in the same breath. But the failure of his verificationist semantics, or the fact that it stands in need of improvement, does not necessarily spell the doom of his negative campaign against a realist semantics. I have already, in quoting Dummett above, given the reader more or less all one needs to envision how Dummett sees an anti-realist meaning theory proceeding. I will say only a little more to complete the picture. Returning again to his discussion of the manifestation requirement, Dummett tells us:

> This requirement calls in question the feasibility of any model of understanding, any theory of meaning, according to which the understanding of a sentence consists, in general, of a knowledge of its truth-conditions, when the notion of truth is construed as satisfying the principle of bivalence and as, in general, given independently of our means of recognizing truth (1993, pp. 116).

His solution is, of course, to abandon the realist notion of evidence-transcendent truth and the principle of bivalence. And so he proclaims:

> We are thus in the position of having to abandon, for certain classes of sentences, the principle of bivalence…we shall have to construct a semantics which does not take, as its basic notion, that of an objectively determined truth-value at all (1993, p. 66).

Nowhere is the influence of intuitionism on Dummett more apparent than with regard to his anti-realist semantics. He very deliberately models his semantics on that of the intuitionists:

> The intuitionistic explanations of the logical constants provide a prototype for a theory of meaning in which truth and falsity are not the central notions. The fundamental idea is that a grasp of the meaning of a mathematical sentence consists, not in a knowledge of what has to be the case, independently of our means of knowing whether it is so, for the sentence to be true, but in an ability to recognize, for any mathematical construction, whether or not it constitutes a proof of the sentence. (1993, p. 70)

This much about Dummett's anti-realism has been prominent since the beginning. He wishes to replace the realist notion of truth with an evidentially-constrained notion, and so he also wishes to replace the realist notion of what it is we grasp when we grasp the meaning of a sentence with an anti-realist one. Again drawing from intuitionism, under the new semantics, "our understanding of a sentence consists in a capacity, not necessarily to find a proof, but only to recognize one when found" (1993, p. 70). Thus our ability to grasp an objective truth-condition is replaced with our ability to recognize what counts as evidence in support of a sentence, evidence that would warrant that sentence's assertion.

The project of constructing a viable anti-realist semantics is a perilous one. Dummett is not convinced that his verificationist route is the correct one. As he admits, "whether a plausible theory of meaning in terms of verification can be constructed, I do not know," and he even begins to spell out an alternative in terms of falsification (1993, p. 76). Still, he holds fast to his

conclusion that a realist semantics cannot be made to fit with an acceptable theory of meaning for our language.

6. Devitt's Dummett

The ground I have covered is not new. A number of philosophers have tread a similar path before, seeking an adequate reconstruction of Dummett's arguments, regardless of whether they support his ultimate conclusions. My own interpretation of Dummett holds a great debt to the work of these philosophers, despite the criticisms I bring to bear on their interpretations and objections. The work of Dummett's most ardent critic, Michael Devitt, deserves special attention due to its uniqueness. Devitt's version of Dummett's argument consists of three premises, which he labels A, B, and C:

> A: The Realism dispute is the dispute about whether sentences have realist (evidence-transcendent) or only verificationist truth conditions.
>
> B: The dispute about truth conditions is the dispute about whether the competent speaker's understanding is realist (evidence-transcendent) or only verificationist.
>
> C: The competent speaker's understanding is only verificationist. (1997, p. 260)

Devitt argues extensively that A is false. I have already criticized his views on the nature of the realism debate, and will not rehash my criticisms here. Still, even if Devitt is right about the realism debate, it is clear that Dummett's argument is about *something,* and we can assess it based on how successful it is at establishing its conclusion concerning truth conditions, even if, as Devitt insists, this does not get us anywhere with regard to Devitt's own version of realism. Dummett's versions of realism and anti-realism are semantic. Devitt's are metaphysical. It is an interesting question whether they ultimately have the same subject matter: realism with a *lowercase* 'r'. But even before we answer that question, we can evaluate the arguments for and

against the different versions of realism and anti-realism on their own merits. In what follows we will take A for granted.

It is easy to see how Devitt's reconstruction of Dummett's argument is supposed to work. From A, B, and C together, it follows that sentences have only verificationist (i.e., evidentially constrained) truth conditions and hence that semantic realism is false. We are assuming A. Where, then, does Devitt think that Dummett goes wrong after A? Devitt takes immediate issue with B: "How could a semantic dispute about the truth conditions of sentences be a psychological dispute about the competent speaker's understanding? How could disputes about such different sorts of property be the same?" (p. 268). As mentioned earlier, for Dummett, understanding a sentence is just knowing what it means, and vice versa. He tends to use the expressions interchangeably. Still, something should be said in response to Devitt, because Dummett's view that understanding and knowledge of meaning should be seen as the same undergirds his entire approach and in fact explains much of what is contentious about his arguments. We will put off a discussion of the issue until we come to the problem of Twin Earth.

Devitt believes that he has isolated an ambiguity in Dummett's premises. So he offers us two distinct versions of both B and C. The first versions of B and C attribute to Dummett the *propositional assumption,* according to which the knowledge a speaker has when he understands a language, "is propositional, or theoretical: it is knowledge-that and not, for example, mere knowledge-how… an L-speaker's understanding of a sentence of L consists in his knowing that the sentence is true-in-L in such and such circumstances" (Devitt 1997, p. 268) Hence we get

> B_1: The dispute about truth conditions is the dispute about whether the competent speaker *knows* realist (evidence-transcendent) or only verificationist truth conditions. (p. 268)
>
> C_1: The competent speaker *knows* only verificationist truth conditions. (p. 269)

Regardless of whether Dummett subscribes to the propositional assumption, or would subscribe to B_1 and C_1, we should note that, at the very least, his subscription to B_1 and C_1 would make some sense of what Devitt found initially puzzling, the idea that a debate over the truth conditions of sentences amounts to a debate about the nature of the competent speaker's understanding of those sentences. For, as Devitt points out, in order for the competent speaker to know something, it must be the case (p. 268). So it would seem that determining what kind of truth conditions a speaker knows will settle what kind of truth conditions sentences actually have. I would contend that a speaker could know the truth conditions of a sentence without knowing whether they are realist or anti-realist in nature, but I will not pursue that line here.

Devitt rejects the propositional assumption which he attributes (somewhat hesitatingly) to Dummett. Hence he rejects B_1 and C_1. As explained in section 1, I do not see Dummett's talk of knowledge to be problematic. It is not clear that Dummett indeed endorses the propositional assumption. One obvious reason for not saddling Dummett with the assumption is that Dummett commonly speaks of knowledge on the part of speakers, but also frames their linguistic competence in terms of a practical ability. Devitt realizes this, but points again to Dummett's use of the words 'knows' and 'knowledge', accusing him of using "weasel words" to avoid full commitment to the assumption (Devitt, p. 269). I stand with Kirkham (1989) in holding that attribution of the propositional assumption to Dummett doesn't take into account the exact nature of Dummett's project. Recall that in framing a theory of meaning for our language in terms of

what speakers know, Dummett is only looking for a "theoretical representation of a practical ability" (Dummett 1993, p. 36). Another reason to think that Dummett does not need the propositional assumption is that when he speaks of knowledge, he tends to see such knowledge as knowledge of how to use the language. This is not propositional knowledge-that, but practical knowledge-how. Furthermore, Dummett, at least more recently, sees the "practical/theoretical dichotomy" as too crude for his purposes (p. xi). The upshot of all of this is that, interesting as his points may be, we do not need to be concerned with Devitt's reasons for rejecting the propositional assumption.

Having already rejected A, Devitt's rejection of every single step in Dummett's argument would be complete, except that Devitt entertains the idea that there are more plausible versions of B and C. These are supposed to be based on Devitt's own view that linguistic competence is a practical capacity rather than a knowledge-*that* anything. Here I would like to register the preliminary complaint that these new versions of B and C do not obviously have anything to do with practical capacities, though admittedly we can agree with Devitt that they "differ from B_1 and C_1 in making no mention of knowledge" (p. 270). Thus he gives us:

> B_2: The dispute about truth conditions is the dispute about whether the sentences understood by the competent speaker *have* realist (evidence-transcendent) or only verificationist truth conditions.
>
> C_2: The sentences understood by the competent speaker *have* only verificationist truth conditions. (p. 269)

Here the emphasis has shifted from *what* the competent speaker knows to direct talk of what sorts of truth conditions are had by the sentences that speaker does understand. As Devitt himself notes, this makes B_2 trivial—the dispute about truth conditions is obviously the dispute

about what kinds of truth conditions sentences have (p. 270). Such triviality should make us question whether B_2 could really have been what Dummett had in mind. Devitt says almost nothing about B_2 and neither shall we. At the same time, Devitt has much to say about C_2, which he rejects, as he must, since C_2 is tantamount to semantic anti-realism. Hence, even granting A (which Devitt of course does not), he thinks that both of his versions of Dummett's argument fail. Recall that the first version was taken to have failed simply on account of the falsity of the propositional assumption (if one had not already taken it to fail on account of the falsity of A).

On any of Devitt's versions of Dummett's argument, the C premise is the most important, and the most contentious, since it states Dummett's verificationism. In fact, neither Dummett nor his other interpreters make much, if anything, of Devitt's A and B—the bulk of what Dummett says of interest is contained in the C premise. It is worth noting that, arguably, Dummett takes something like C_2 to follow from something like C, though Devitt casts these as two different versions of the same thing. Dummett's argument is, I would say, best understood as a *transcendental argument* for semantic anti-realism. Facts about us as human beings and about what counts as an acceptable theory of meaning place certain constraints on what the meanings of our sentences could be, given the presumption that we do understand the sentences of our language. One way of summing up Dummett's argument is as follows: Since we could only come to understand sentences with verificationist truth conditions (in essence, Devitt's original C), and we do understand our language, our sentences must have verificationist truth conditions (Devitt's C_2). So Devitt's attempted precisification of C is not incredibly apt.

I have emphasized that one might view the problems Dummett presents for a realist semantics as insurmountable without at the same time endorsing his verificationism—at the very least, one need not accept every aspect of Dummett's verificationist semantics in order to

appreciate the reasons he gives for why he thinks a realist theory of meaning cannot work. Devitt's in-depth discussion of C_2 is ultimately an attack on Dummett's positive program. In fact, Devitt presents Dummett's (alleged) argument for C_2 almost entirely as a positive argument in favor of verificationism rather than as a negative argument against realist semantics. Hence, I will put off grappling with the many serious objections Devitt raises to C_2. For now let it suffice to say that Devitt's arguments against Dummett's verificationism do not show that Dummett's arguments against semantic realism do not work.

7. Dummett on Twin Earth

Semantic externalism presents a *prima facie* challenge for the Dummettian. If our relationship to our own language is as Dummett says it is, if the facts about how we learn meanings, convey meanings to others, and give meanings to our words in the first place are as Dummett takes them to be, then it is, at the very least, not obvious how we can account for the sorts of Twin Earth cases that semantic externalists hold up against their internalist opponents. The externalist challenge was raised against Dummett in the 1980s by Colin McGinn (1982) and Michael Devitt (1997), and has been more recently articulated at length by Panu Raatikainnen (2010). Responses to the challenge have been rather slim and weak. Crispin Wright (1993) writes off the challenge as misguided in no more than one page. Alexander George (1984) spends just over a page defending Dummett, concluding that no strong case was made against Dummett to begin with. Mark Gardiner (2000) also discusses the issue for approximately a single page. He concedes that the challenge is real but claims that Dummett should just reject semantic externalism.

The rejection of semantic externalism should not be taken so lightly. But the moves that might be made on Dummett's behalf have yet to be fully articulated. Indeed, there is no very full

discussion of all aspects of the externalist challenge to the Dummettian currently on the market. Here I will attempt to spell out the challenge in its details and provide a fuller account of how the Dummettian could respond than has heretofore been attempted.

At the outset, it is crucial that we first separate Dummett's negative and positive views. His negative view is that semantic realism must be wrong. His positive view is that some form of verificationism (or perhaps falsificationism) is correct. It is important to distinguish Dummett's attack on realist semantics from his promotion of a particular brand of anti-realist semantics for two reasons. The first reason is that we need to be clear about how Twin Earth cases threaten Dummett's positive and negative programs, respectively. The second is that while a given response might be sufficient to defend one of Dummett's programs, it might not be sufficient to defend the other. While Twin Earth cases might lead us to conclude that Dummett's verificationism cannot be correct, that would not immediately imply that Dummett's reasons for rejecting semantic realism aren't any good.

McGinn aims his critique at something like what I have called the endowment argument, saying that he seeks to, "rebut [Dummett's] criticism of realist semantics by showing that it depends upon a conception of what it is that confers content upon sentences which need not and should not be accepted" (McGinn 1982, p. 113). As we have seen, Dummett would say that use determines meaning, and that use is inherently tied to our capacities for epistemic appraisal. McGinn puts Dummett's idea this way:

> In sum, grasp of meaning can be nothing other than mastery of use, but use is a matter of recognising and responding to assertibility conditions: truth conditions, however, are inherently unconnected with use, at least when they are verification-transcendent, and so cannot contribute to the determination of content. (McGinn 1982, p. 114)

The basics of Twin Earth cases are quite familiar, but it will be useful to review them with an eye toward attacking Dummett's views. Imagine that it is 1750, before the dawn of modern chemistry. Water is H_2O, but no one on Earth knows it because science at the time is not sufficiently advanced for anyone to know that. Now imagine that, in 1750, there is another place almost exactly like Earth, populated by beings like us, with the same knowledge of chemistry that Earthlings would have had in 1750. Call that place 'Twin Earth', and call its inhabitants 'Twin Earthlings'. The crux of the case is that on Twin Earth, what they use the term 'water' to refer to is a substance which is almost exactly just like water, with one crucial difference: It isn't H_2O, it's something else, which we can call XYZ. And yet, if an Earthling were to go to Twin Earth, she would discover that Twin Earthlings also have a clear liquid they call 'water' which rains from the clouds, flows through the streams, quenches thirst, etc. In 1750, so the case goes, an Earthling would have had every reason to believe that the clear substance on Twin Earth was water. Externalists present this case in the hopes that we will agree that, even though no one in the case knows that 'water' as spoken by an Earthling refers to the substance H_2O while the same term spoken by a Twin Earthling refers to the substance XYZ, 'water' still refers that way. The extension of 'water', it would seem, is fixed by more than what is, so to speak, "in the head" (cf. Putnam). The extensions of terms like 'water' seem to be fixed, at least in part, by environmental factors, in particular such factors as being appropriately causally connected to certain substances in the speaker's environment.

A distinction from the philosophical literature on natural kinds, borrowed from Locke, is helpful here. Let the *nominal essence* of a natural kind like H_2O be whatever set of epistemic criteria normal speakers use to determine whether something is properly referred to as 'water'. And let the *real essence* of a natural kind be whatever microstructural properties it has, such as

being H2O, that make it what it is rather than something else. Then, arguably, we would want to say that Earthlings and Twin Earthlings use the term 'water' in such a way that they associate the same nominal essence with the term (the way the stuff looks, feels, tastes, etc.), while in fact they are referring to two distinct substances with different real essences. One of the liquids is H2O while the other is XYZ. Externalists would further contend that if an Earthling were to go to Twin Earth in 1750 and proclaim, "This is water," he would have said something false. For the Earthling's word 'water' refers to H2O, while the stuff all around him on Twin Earth would actually be XYZ.

Given Dummett's conception of use based on our capacities for epistemic appraisal, McGinn contends that the use the Twin Earthlings put 'water' to is the same as the use regular Earthlings put it to. His concern is that, "speakers on earth and twin earth thus acquire the same recognitional capacities and manifest them in the same conditions of evidence, but their sentences do not mean the same: they have the same assertibility conditions but different truth conditions" (p. 116).

Raatikainnen argues along similar lines, telling us that the arguments for semantic externalism "entail that there is a definite sense in which we do *not* actually know the meanings of many of our words" (p. 6), and that furthermore, "the idea that understanding just is knowledge of meaning…just cannot be correct, if the basic lesson of semantic externalism is true" (p. 6). He, like McGinn, claims that Earthlings and Twin Earthlings would use words in the same way despite those words differing in reference and hence, as far as they are concerned, meaning. The upshot of this claim about use is supposed to be that use alone cannot account for all the relevant facts about what our words mean. Thus Raatikainnen says that,

> One simply could not determine, on the basis of observable linguistic behavior of the language-users in 1750, whether our 'water' and their 'water' had the same meaning or not. The manifestable use of the two linguistic communities would be exactly the same. So would any explicit verbalizable knowledge of meaning. (p. 7)

We now turn to Devitt's take on the issue. According to Devitt, Dummett's view of language requires that, "someone who is competent with a term must have the ability, roughly, to manifest behaviour leading to recognition of its referent" (1997, p. 279). According to an externalist (causal-theoretic) view, Devitt says, "a speaker can use such a term to refer even if he is almost entirely ignorant about its referent. He may not be able to describe, recognize, or know how to track down the referent" (p. 279). The problem is supposed to be that, for instance, in Twin Earth cases, a descriptivist has no way of accounting for the fact that the Earthling said something false when on Twin Earth he uttered 'this is water'.

Part of Devitt's critique, as with Raatikainnen's, is that Twin Earth cases seem to show us that understanding and truth conditions come apart. This is supposed to pave the way for the possibility that a sentence might have evidence-transcendent truth conditions. Thus Devitt says,

> A person, even a whole speech community, may understand, say, 'Catiline is *F'* or *'x* is an echidna', without being able to establish whether those statements are true or false…so if these arguments are good, the whole anti-realist case collapses: people *can* understand statements that have truth conditions transcending their epistemic capacities. (p. 280)

So we now have on the table the following externalist problems for Dummett:

1. Earthlings and Twin Earthlings use the term ‘water’ in the same way, and hence manifest the same understanding, but they refer to different substances with the term and therefore mean different things by it.
2. Understanding a sentence cannot be the same as grasping its truth conditions, because in light of Twin Earth cases someone might understand ‘water’ without knowing its referent and hence without knowing how the term contributes to the truth conditions of whole sentences containing it.
3. Twin Earth cases show that, contrary to Dummett’s main contention against the semantic realist, we can understand sentences whose truth values we cannot ascertain.

We will begin with the third alleged problem for Dummett. For the moment, we will set aside concerns about whether what matters for Dummett is an outright inability to know a sentence’s truth value or simply the lack of an effective procedure for coming to know the sentence’s truth value.

In the Twin Earth example where the Earthling travels (somehow) to Twin Earth in 1750 and falsely utters, ‘This is water’, we believe that the Earthling understands the sentence he used though he is unable to determine whether it is true or false given that he is unable to know whether the liquids on the two planets actually have the same real essence. Thus, according to the critics we have been discussing, Dummett is wrong to think that truth is epistemically constrained. Devitt’s ‘echidna’ and ‘Cataline’ examples are supposed to supply further evidence in support of the same conclusion.

Crispin Wright boldly claims that semantic externalism provides no reason at all to reject Dummett’s arguments against semantic realism, telling us,

> The suggestion that, for instance, I may—by ordinary criteria—fully understand the sentence, 'This ring is fashioned from 16 carat gold' without thereby knowing what, essentially, has to be the case in order for it to be true, carries no implication that the essence of gold cannot be identified or that samples of gold cannot be recognized qua instances of that essence. (1993, p. 34)

Wright is here claiming that semantic externalism does not undermine evidentially constrained truth. We should note that Wright's successful defense of Dummett would not imply the veracity of Dummett's verificationism—only that semantic externalism would not imply semantic realism because externalism would be consistent with truth's being evidentially constrained. Hence Wright is defending Dummett's negative program.

Responding to Wright, Devitt pushes the point that, no matter how you slice it, it must be admitted that understanding a sentence does not require that one be able to discover whether it is true. Alexander George offers a rejoinder on the Dummettian's behalf:

> Dummett's view does not depend solely on speakers' having a capacity *actually* to recognize the conditions of verification of a sentence when confronted with them…but depends also on speakers' having a capacity *in principle* to recognize these conditions. (1984, p. 525)

Devitt's response to George is that there is no way to construe 'in principle' such that Devitt's original objection is met without rendering realism and anti-realism indistinguishable. The reason is supposed to be that, when 'in principle' is understood broadly enough to deflect Devitt's objection, Dummett's manifestation requirement will not be met since we will have appealed to capacities that are not "presently manifestable" (p. 280). Dummett, however, does not claim that we must be able to manifest knowledge of meaning presently, only that we must

be able to manifest it when circumstances would allow us to do so. There may in fact be practical limitations which we could in principle overcome. We might think of things this way: Dummett only requires that we be able, in principle, to manifest our understanding of any expression we do understand. There is nothing keeping a speaker in 1750 from being able, *in principle,* to determine that 'water' as uttered by an Earthling is H2O and 'water' as uttered by a Twin Earthling is XYZ.

At this juncture, we would do well to take a look at what Dummett himself has to say on the matter. As Dummett's commentators are well aware, Dummett objected for various reasons to the externalist programs of Putnam and Kripke, in particular to Kripke's use of the notion of rigid designation. The question currently on the table, however, is whether a Dummettian can fend off the threat posed by Twin Earth cases while accepting the basic intuition driving those cases. Dummett seems to think such an independent defense is possible. Consider the following telling passage:

> It is part of the sense of such a word that it stands for a species or a kind of substance, not for something recognizable by external appearance alone…to grasp the use of a word just is to attach a sense to it…[the sense] is what a speaker knows when he understands the word…[the sense] must be capable of being exhibited as a means of determining the reference…when someone knows the sense of a sentence, what he knows is how the truth-value of the sentence is to be recognized, whenever we are in a position to do so (1981, pp. 146-147).

In the first part of the above quotation, Dummett seems to be endorsing something along the lines of what another use theorist, Paul Horwich, has suggested as a way of handling Twin Earth

cases. According to Horwich, the use theorist might accept a basic rule of use for 'water' such that:

> x is water ↔ x has the underlying nature, if any, of the stuff in *our* seas, rivers, lakes and rain. (2005, p. 16)

Thus, Horwich contends, expressions on Earth and Twin Earth may have the same meaning (in virtue of being used in the same way) while having different extensions (p. 16).

Where Horwich talks of meaning, Dummett speaks of sense. Dummett endorses the Fregean distinction between sense and reference, with the twist that, as we have just seen, he takes grasping the sense of an expression to be the same thing as grasping its use. So it would come as no surprise to Dummett that one can grasp the sense of a word without knowing its reference. The very distinction between sense and reference was introduced by Frege to explain how it could be that one might understand a term without knowing to what the term refers, in particular without knowing that the term in question has the same referent as another term which one also understands. In an interview regarding his work on Frege's philosophy, Dummett says:

> Sense is something that the mind can grasp. The sense, together with the way the world is, determines its reference. In the case of a singular term, the reference is an object of some kind. It might be a city or a river or a person…so the reference is something you don't have to know in order to understand the expression whose reference it is. You may not know what is the capital of Bulgaria. You understand the expression 'the capital of Bulgaria' but you don't know what its reference is…reference is something which goes to determine the truth or falsity of what is said, but which you don't have to know in order to understand what is said…knowing the sense of a proper name is simply knowing how to determine whether a particular individual bears that name. (transcript)

Hence we can see that, for Dummett, to understand a sentence is to grasp its sense, and to grasp its sense is to know how to use it. So what one must manifest is not knowledge of the referent of a term, but an understanding of what the term means, what its sense is. Accordingly, what is conveyed via our use of a sentence is not information which necessarily includes knowledge of the referent, but rather information about how to use the sentence, i.e., the sense of the sentence.

As for the third externalist problem for Dummett, externalism doesn't give us any special reason to think that we can understand sentences whose truth values we cannot, even in principle, ascertain. Thus, Devitt's concern that the Earthling has no presently manifestable capacity to determine the referent of 'water' is misplaced. The first externalist problem is also misguided. The fact that Earthlings and Twin Earthlings end up referring to different things with the same term is not enough to show that, in some important respect, the meanings of the terms cannot be the same. 'Meaning' is an ambiguous term. Different thinkers mean different things by it. The externalists discussed above establish to their own satisfaction that the meaning 'water' is different on the two planets, and then attempt to make trouble for Dummett on this basis. The fact is, as we have seen, that Dummett would likely not endorse their view of meaning anyway. What Dummett thinks is grasped, conveyed, and understood is the *sense* of an expression. While it may be somewhat out of the ordinary, it is completely open to Dummett to defend a notion of sense such that a term can have a sense associated with it which attaches to different referents in different circumstances. This is, in fact, the same kind of thing that we find with Kaplanian character. Furthermore, it remains open to the Dummettian to construe use broadly so that the uses of 'water' on Earth and Twin Earth are actually different.[6]

[6] Thanks to David Braun for extensive discussion of these matters.

The second externalist problem is more worrisome. Dummett's approach to the realism issue has us ask what kinds of truth conditions our sentences could possibly have. His way into this problem is via understanding—certain facts about what we are able to understand might lead us to believe that our sentences have either realist or anti-realist truth conditions. So the idea that understanding a sentence can't be the same as grasping that sentence's truth conditions is indeed a threat to Dummett's overall approach.

Dummett's views about the relationship between meaning, truth, and understanding are foundational to his views on realism and anti-realism. Understanding is grasping a sense, and grasping a sense is knowing how an expression is to be used. We have repeatedly seen that Dummett construes use in such a way that it is inherently tied to our capacities for epistemic appraisal. Indeed, if use is not construed that way, Dummett's arguments against semantic realism do not even begin to establish their intended conclusion. If use is construed that way, Dummett has at least made a case against semantic realism. The trouble is that, much more likely than not, Dummett is simply wrong about how language works. In defending his manifestation requirement, it was argued that, surely, if meaning is essentially conveyable, then it is reasonable to require that knowledge of meaning be manifestable in use. At the same time, it is probably just false that manifestation of knowledge of meaning requires demonstrating an ability to recognize anything at all. Likewise, it is probably not the case that we learn language only by learning to respond appropriately to specific evidential stimuli. Finally, it is most likely not the case that the only way we are able to give meaning to our words is by responding to certain epistemic situations with certain utterances.

So Dummett is likely wrong about language. Still, we can try to imagine how a Dummettian might respond to the second externalist problem for Dummett—the only alleged problem which poses a real threat. Recall the problem:

> 2. Understanding a sentence cannot be the same as grasping its truth conditions, because in light of Twin Earth cases someone might understand 'water' without knowing its referent and hence without knowing how the term contributes to the truth conditions of whole sentences containing it.

Here is a suggestion: Perhaps a Dummettian could understand capacities for epistemic appraisal in such a way that what is understood, grasped, and manifested in use is essentially tied to the particular stuff that caused the sensory evidence in the first place.[7] The thought is that through the use of language, perhaps we demonstrate an epistemic capacity to determine not just the nominal essence of a referent such as 'water' but also the real essence. There is something to this. The idea is similar to that of Horwich's mentioned above. If this idea were correct, what we would be showing others when we used words in particular ways would be more than the fact that we associated those words with a particular evidential situation narrowly construed as, for instance, the tactile sensation of wetness and the visual sensation of clearness. We would be showing others that we associated the term 'water' with *water itself,* whatever that turned out to be. Care must be taken at this point. We would not thereby be showing others that we associated the concept of water with the concept of H2O, even though H2O is what water is. We might not know that by using the word 'water' we referred to H2O. At the same time, we might be showing others that we have the epistemic capacity to determine the truth value of sentences

[7] Thanks to David Braun for suggesting this thought concerning epistemic capacities. I am indebted in what follows to discussions with him. I hope I have not misrepresented his suggestion.

containing the term ‘water’ as we use it. Such a view would be consistent with Dummett’s ideas. But does it overcome the objection we are considering?

In order for the view just described to meet the problem presented by Twin Earth cases, the view would have to somehow allow that understanding a sentence amounted to grasping its truth conditions. The difficulty posed by Twin Earth is supposed to be that we might fully know how to use a sentence and hence, according to Dummett, understand it, without knowing what the conditions for its truth really are. It is important to see what the issue is not. It is not that we might understand a sentence without knowing what its truth value *actually is.* The problem is supposed to be that we might understand what a sentence means, in Dummettian terms—grasp its sense, i.e., know how to use it appropriately, without knowing what it would take for it to be true. So in the case of an Earthling who goes to Twin Earth, drinks a glass of XYZ, and then says ‘I drank water today’, that Earthling says something false. The Earthling could know what it would take for the sentence to be true. He must simply know that he must be drinking the same kind of stuff that he normally drinks when he drinks water on Earth. And, indeed, if you told the Earthling that water on Earth had a different microstructural essence from what people on Twin Earth referred to with the term ‘water’, the Earthling would realize that what he thought to be true was in fact false—he did not drink water today. He drank something else. So, then, it seems that the Earthling does know what it would take for his utterance to be true. And it also seems that he might manifest such a grasp of the truth conditions of the sentence he uttered just as Dummett suggested in the above quotation when he said, “It is part of the sense of such a word that it stands for a species or a kind of substance, not for something recognizable by external appearance alone” (1981, p. 146).

So, I would suggest, Dummett does not need to be too concerned about Twin Earth cases, though they do help to bring out some under-explained aspects of the somewhat common meaning-as-use view of language. I have seriously suggested that Dummett is simply wrong about language, and I stand by that suggestion. There is no serious pressure on one who rejects Dummett's Wittgensteinian understanding of language to accept Dummett's anti-realist conclusions.

All of that said, we can still ask about the relevance of Dummett's reasoning today. Note that Dummett has not only provided arguments against semantic realism and for semantic anti-realism, he also gives us certain ways of understanding what realism and anti-realism are, and makes certain claims about the implications of an anti-realist view. For instance, Dummett tells us that the best way to think about realism and anti-realism generally is in terms of truth conditions. There may be value in this claim even if Dummett's arguments pertaining to semantic realism and anti-realism simply do not go through. Dummett also thinks that the semantic anti-realist ought to be compelled to give up the principle of bivalence and, with it, classical logic. This is an important claim even if we accept that Dummett's anti-realist arguments fail. For there may be other arguments capable of establishing the fact that our sentences, or at least some of them, have anti-realist truth conditions. As long as that is a possibility, we will want to know whether an anti-realist conclusion in some area should lead us to revise our reasoning in that area. It is to this topic that we turn in the next chapter.

Chapter 4: Semantic Anti-Realism and the Revision of Classical Logic

1. Introduction

The point of this chapter is to follow up on some ideas of Dummett's that we began to develop in the previous chapter. In that chapter, we saw that Dummett's positive arguments for semantic anti-realism were unconvincing. The main reason was that his argument required acceptance of a use-theory of meaning that practically no realist would want to endorse. This makes Dummett's reasoning very nearly, if not precisely, question-begging; it shouldn't convince any realists to become anti-realists. All the same, it remains the case that a satisfactory argument for semantic anti-realism might be given, though we will not attempt to give one here. Instead, we will explore some of the consequences of accepting Dummett's framing of the realism issue in terms of evidence-transcendent and evidentially constrained truth conditions. In particular, we are going to be concerned with the reasons why a Dummettian anti-realist, who takes our sentences to have anti-realist truth conditions and endorses an evidentially constrained notion of truth, should accept intuitionistic rather than classical logic. In the end we will conclude that Dummett was never in possession of an undecidability concept that could do all the work he needed it to do. That said, the project of generalizing intuitionism from mathematics to other areas remains viable and interesting.

There are two threads to follow here. The first concerns a point briefly raised in the last chapter about the connection between anti-realist truth conditions and effective decidability. Why would Dummett's arguments, which must make use not of pure unknowability but rather the more restricted notion of lack of an effective decision procedure, lead us to accept the conclusion that sentences about a given subject matter have truth conditions such that we must be capable, in principle, of knowing the truth values of those sentences? Another way of putting

this question is: What is it about the relevant notion of undecidability which makes it the case that the application of bivalence to undecidables commits us to an evidence-transcendent notion of truth for those undecidables? The other thread involves Dummett's purported argument that his semantic anti-realist cannot accept classical logic for the area about which he takes an anti-realist attitude. And just what is an anti-realist attitude? For our purposes it will be taking the view that the sentences in a given area of inquiry have anti-realist truth conditions, which amounts to the same thing as the view that an epistemically constrained notion of truth applies in that area. Some have asked whether a Dummettian anti-realist is under any special pressure to reject classical logic in favor of, in particular, intuitionistic logic. The debates surrounding these points are, as one might expect, rather convoluted. It turns out that these two threads will need to be untangled somewhat as we go along.

> **thread one:** What is the relationship between evidence-transcendent truth and the application of (determinate) bivalence to sentences lacking an effective decision procedure?
>
> **thread two:** Must one who rejects evidence-transcendent truth and endorses instead evidentially constrained truth also reject (determinate) bivalence and with it classical logic?

Because they contain overlapping notions, it is tempting to think that these questions could, or even should, be answered together. Unfortunately, the debate over whether the anti-realist need be a logical revisionist is obscured by a lack of clarity on just how the notions involved in our questions are related. We will attempt to keep the threads as separate as possible while being as clear as we can about the relevant ideas.

2. Independence and Transendence

We noted in the last chapter that admitting that an undecidable has a determinate truth value requires admitting that its truth is evidence-transcendent only in the sense that we do not know whether, were we to attempt to discover its truth value indefinitely, we would eventually make the discovery. It was also noted that this sense of 'evidence-transcendent' does not fit our own definition of evidence-transcendence. Recall the definnitions of evidence-transcendent and evidentially constrained truth:

> **evidence-transcendent truth:** Evidence-transcendent truth applies to a class of sentences iff those sentences can have a truth value independently of our ability to ascertain which one they have.
>
> **evidentially constrained truth:** Evidentially constrained truth applies to a class of sentences iff those sentences cannot have a truth value independently of our ability to ascertain which one they have.

Our definitions of evidence-transcendent truth and evidentially constrained truth accord with the standard explications of the realist and anti-realist ideas that, respectively, a sentence either has a determinate truth value independently of our ability to come to know that value or its possession of a determinate truth value depends in some way on our ability to come to know. At least since Dummett, this general way of conceiving the two positions has been widespread. Consider the following statement of Loar's: "A realist about certain statements holds their truth or falsity to be independent of our ability to verify or to falsify them" (1987, p. 81). Consider as well this statement of Jenkins':

> Realism about a subject matter is often defined as the view that truths or facts or states of affairs relating to that subject matter are mind-independent—that is, that they obtain

> independently of their being known, knowable, conceivable, or related in some other way to our mental lives. (2005, p. 19).

Jenkins' work figured in chapter 2, where we wrestled with the notion of mind-independence. At that time we said that semantic, as opposed to metaphysical, characterizations of the realism dispute were just as much about the question of mind-independence as the metaphysical characterizations, since our epistemic capacities are aspects of our minds. It is important to see that *independence* is the central realist notion in our current investigation as well, as the quotations from Loar and Jenkins illustrate, and as I shall explain.

One of the main contentions of chapter 3 was that undecidability in the context of Dummett's arguments should be understood as lack of an effective decision procedure rather than as complete unknowability. Dummett also, of course, spoke often of evidence-transcendence. It is unclear whether he intended the truth of a sentence to be evidence transcendent just in case the sentence had a determinate truth value despite being undecidable. Unfortunately, in the debate that has emerged in the wake of Dummett's work, the confusion has only deepened. Now we must clarify the notion of evidence-transcendence. In the literature, terms such as "evidence-transcendent" and "recognition-transcendent" are sometimes used to indicate our notion, that is, the notion that something can have a truth value independently of whether we can now, or could ever, come to know what that value is. At other times, such terms seem to be intended to pick out the idea that something might have a truth value even though we had no effective procedure for determining what that value is. Let's admit that the terminology here is sloppy. Based on the way terms such as 'evidence-transcendent' and 'recognition-transcendent' are actually used, one would have a hard time in many cases telling which notion a

given author intended. This fact should not keep us from making the distinction ourselves and treating each disambiguation appropriately.

Let's distinguish two different understandings of evidence transcendence. The first is at work in our definition of evidence-transcendent truth already given, and we will call it the *independence reading.* The second we will call the *undecidable reading.*

> **evidence-transcendent truth (undecidable):** Evidence-transcendent truth (undecidable) applies to a class of sentences iff those sentences can have a truth value despite our lack of an effective procedure for determining which one they have.

On the first, original reading, the point is just that our ability to know the truth doesn't make a difference to the truth. Whether we could come to know the truth value of a given sentence or not, it would have a determinate truth value anyway, or so says the realist. On this way of viewing things, the independence reading of 'evidence-transcendent' has it that there is at least one possible world where the sentence under consideration has a determinate truth value while in that world we are incapable of discovering its truth value. Granted, there may be worlds where we are in fact capable of ascertaining the sentence's truth value, but that makes no difference either. The sentence would have a determinate truth value either way.

It is worth noting that there are nested modalities here—some modal operators fall under the scope of others. This should be no problem, as long as we are careful about which formulas are true (or not) at which worlds, and as long as we are careful about our specification of the accessibility relation in each case. There is no room to spell out the technical specifics, but here is a representative example, expressed informally: On the independence reading of evidence-transcendence, the formula Fa has evidence-transcendent truth in the actual world, w_1, if and only if there is a possible world w_2, accessible from w_1, for which the following two things

obtain: Fa has a truth value at w_2 and there is no world w_i, accessible from w_2, such that at w_i it is true that we know the truth value of Fa. There are different kinds of accessibility relations that may be relevant depending on the context at hand. For instance, we may wish to limit the relation according to whether the sort of possibility we have in mind is metaphysical, logical, or nomic. I do not wish to make any assumptions here about which kind of possibility matters.

The anti-realist notion of evidentially constrained truth can be spelled out in similar yet simpler terms: The formula Fa has a truth value in an evidentially constrained way in the actual world, w_1, if and only if there is at least one possible world w_2, accessible from w_1, at which we know the truth value of Fa.

Notice that in the case of evidentially constrained truth we do not require nested modalities. In this way, the concept is simpler. At the same time, we have done nothing to specify what the world w_2 is like. So we have done nothing to specify what the constraint on truth is in any particular case. This should come as no surprise, because the notion of truth may be constrained in numerous different ways. Perhaps the reader will be disinclined to count the result of evidential constraint as a truth concept at all. But to be so disinclined is just to lack any inclination to endorse Dummettian anti-realism.

So a sentence's having an evidence-transcendent truth value requires the existence of a possible world where the sentence has a determinate truth value even though we are incapable of coming to know what that value is. Note that the truth of a sentence could be evidence-transcendent in the undecidable sense though there were no worlds where that sentence had a determinate truth value while we can't come to know it. For recall that the lack of an effective decision procedure only tells us that we don't know that we would eventually land upon an answer if we tried again and again. Such a state of affairs is entirely consistent with our being

capable of coming to know the answer. So the undecidable reading of evidence-transcendence certainly doesn't entail the independence reading. And the independence reading doesn't entail the undecidable reading either. This is most easily seen once it is recognized that the independence reading deals only in possible worlds while the undecidable reading additionally involves the notion of time. The independence reading is silent on temporality and has no temporal implications beyond that if we are unable, at some world, to come to know the truth value of a sentence, then we are not able to come to know its truth value in any amount of time. Consider now a sentence *S* to which evidence-transcendent truth actually applies.[8] Then there is a world w accessible from the actual world such that *S* is true at *w* and at *w* we are incapable of learning that S. Hence at w we have no effective procedure for discovering whether S is true. It might seem as though the independence reading of evidence-transcendence does entail the undecidability reading since at w we are not only incapable of discovering S's truth value but we also have no effective procedure for making such a discovery. The reason this line of thought is not enough to establish that the undecidability reading follows from the independence reading is that we began by saying that evidence-transcendent truth applied to S *in the actual world.* Whatever the status of our abilities and procedures at w may be, we might possess an effective procedure for coming to know S *in the actual world.*

Recall from chapter 2 Jenkins' distinction between modal and essential independence. It is now time to explore her distinction further. Concerning essential independence, Jenkins says that the issue is whether some fact about our minds is part of the "essence of *something's being the case*" (p. 200). Concerning modal independence, the issue is supposed to be whether some fact about our minds is necessary for something to be the case. Jenkins notes that, "many

[8] Thanks to David Braun for pushing me on this point.

familiar claims about possible recognition-transcendence…seem to be claims of modal independence" (p. 200). For our purposes, we can frame things in terms of sentences and translate 'being the case' as 'being true'. Then, on the *essential independence* reading, we would have the realist claiming that it is no part of what it is for a sentence to be true that we have any ability to come to know that it is true. Correspondingly, the anti-realist would claim that part of what it is for a sentence to be true is for us to be capable of coming to know that it is. On the *modal independence* reading, however, we get our definitions of realist and anti-realist truth and truth conditions from the previous chapter. On this reading, the realist says that a sentence can have a truth value even if we are not capable of coming to know what that value is, while the anti-realist says that if a sentence has a truth value then we are capable of coming to know what truth value it has.

These two distinct ways of framing the realist and anti-realist positions in terms of independence may not be equivalent, but I would like to suggest that they are both acceptable as far as a Dummettian is concerned. Further theoretical commitments may impel one who understands the realism debate in Dummett's way to choose essential independence over modal independence, or vice versa, but there is nothing in the Dummettian understanding of the issues to force a decision either way. That said, the modal notion is quite useful, and as noted we have been working with the modal notion. I would also like to suggest that modal considerations are indispensable when attempting to establish whether truth is essentially knowable, whether it is part of what it is for something to be true that we be able to come to know that it is. Demonstrating that it is possible for a sentence to be true while we could not come to know that it is would be sufficient to show that it is no part of truth's essence that it be knowable. At the same time, someone could accept that it is impossible for a sentence to be true without us being

able to know that it is while rejecting the further claim that part of what it is for a sentence to be true is that it be knowable. That further claim is a further piece of metaphysics we are under no special obligation to accept. We might, for instance, reject the idea that anything has an essence at all.

Jenkins would not approve of our use of modal independence in framing the realism issue. She provides at least four arguments intended to prioritize the importance of essential independence in the realism debate. I will defend my prioritization of modality, strictly understood, as opposed to essence by attacking each of her arguments in turn. One of the main claims that Jenkins makes against giving modal independence a significant theoretical role is that

> Essential independence is a view about the very nature of something's being the case: a view about *what it is* for that thing to be the case. Modal independence, on the other hand, merely concerns the possibility of that thing's being the case while something else is *not* the case. (2005, p. 201).

Jenkins' claim here is that, "the question of essential independence is…more philosophically fundamental than the question of modal independence" (2005, p. 205). She concludes that realism and anti-realism ought to be framed in terms of essential independence since the dispute between them is supposed to be "a deep debate about the nature of reality" (p. 205).

Jenkins' reasoning is questionable on at least two counts. First of all, it is by no means apparent that deep debates about the nature of reality cannot be decided by considerations involving concepts that are less than fundamental. Jenkins would appear to agree, as when she says, "while modal independence claims may turn out to be interestingly related to realism, they do not constitute the core of the realist position" (2005, p. 205). But if the debate can be settled without recourse to our most fundamental concepts, it is hard to see why the fact that the realism

dispute is supposed to be a deep debate should lead us to think that the views at issue aren't best conceived in less than fundamental terms.

Kit Fine (1994) can be credited with bringing to the philosophical forefront the deep difficulties involved in analyzing the notion of essence in modal terms. Jenkins herself draws attention to Fine's work and notes that there are similarities between his discussion and hers (Jenkins 2005, p. 208 n1). In a later work (2001), Fine argues that the realism debate in a given area is best understood as the debate as to whether the facts in that area are grounded in other facts or are themselves fundamental.

Jenkins could, of course, endorse a view akin to Fine's and simply say that the issue of realism is just the issue of fundamentality—in this case the question as to what the most fundamental thing is, i.e. the thing that grounds the reality of everything else. The issue here is that Jenkins talks about something being more philosophically fundamental while the phenomenon of essence she describes has to do with something being metaphysically fundamental. It may be the case that the essence of something has more to do with metaphysical fundamentality than the modal nature of something, where essence and modality are understood in the way of Jenkins and Fine. This was just Fine's point to begin with. But it does not follow that the question as to what something's essence *is* must be more philosophically fundamental than the question as to what its modal nature is. To assume as much is to assume that questions of metaphysics are prior to questions of logic or language. But as is by now quite familiar, there is a deep question about how to construe the realism debate. The Dummettian tradition stems from a commitment to dealing with realism at the level of language, not metaphysics.

Jenkins says no more about the idea than she does in the lines I have quoted. In the larger passage above, she makes much of the notion that essential independence is about "the very

nature of something's being the case," while modal independence "merely concerns the possibility of that thing's being the case while something else is *not* the case" (p. 201). This appears to be her main reason for thinking that modal independence is less philosophically fundamental than essential independence.

Admittedly, modal independence claims will make reference to "something else," but it's important to see that essential independence claims also involve other things' being the case. That's the whole point of the *independence* part of the view. Essential independence involves independence from other things' being the case inasmuch as those things are *not* involved in *what it is for something to be the case.* For instance, a realist of the appropriate kind would say that what it is for something to be the case in mathematics has nothing to do with what is the case concerning our mental lives. Granted, modal independence doesn't tell us about the "very nature" of anything, but modality is certainly one of the most fundamental concepts in our repertoire. I conclude that Jenkins fails to adequately motivate her concerns about modal independence and philosophical fundamentality.

Jenkins' most serious objection to the framing of realism issues in terms of modal independence concerns the question as to whether a constructivist about mathematics should admit mathematical truths in possible worlds where there are no beings like us to perform the requisite constructions:

> Consider a constructivist philosopher of mathematics who is asked to imagine a possible world w with no people in it, and say whether 2 + 2 = 4 at w. She may argue that she is at liberty to answer 'yes,' provided she maintains that its being the case at w that 2 + 2 = 4 amounts to nothing more than the fact that *she,* in assessing w, subjects it to the mathematical structure which she has constructed. (Jenkins 2005, p. 203).

In this case, Jenkins argues, we have a fact of arithmetic which is *modally* independent of us since it obtains despite the fact that, in the possible world in question, there are no beings like us at all and hence our minds are no particular way at that world. Jenkins further claims that a constructivist could happily endorse this way of thinking of things. The upshot is supposed to be that arithmetic is conceivably essentially dependent upon our minds given that facts of arithmetic are the result of our constructing mathematical reality while those facts are nonetheless modally mind-independent. Jenkins concludes that modal independence is the wrong tool for thinking about the mathematical realism issue (p. 203). She compares this sort of situation with common thoughts concerning anti-realism about things like beauty. An anti-realist about beauty might think that facts about beauty are modally independent of our minds in the following sense: In a world with no people, sunsets might be beautiful inasmuch as we, when imagining those worlds, find the sunsets we imagine there beautiful. But this would not commit us to realism, despite being an instance of beauty's modal independence from us, because the fact about the sunsets being beautiful is still the result of *our finding them beautiful* (p. 203).

The main thrust of Jenkins' remarks is the contention that the realism debate is not usefully framed in terms of modal dependence and independence. She contends that modal independence is not sufficient to establish realism since it is consistent with anti-realism (of the mathematical variety in the example just sketched, for instance). At the same time, she argues that modal dependence is not sufficient to establish anti-realism. For it may be, she says, that, as a matter of fact, all of the truths in a given area are knowable by us though it is in no way part of the reality about which those facts are concerned that we be able to know them. Her example is, again, arithmetic. She outlines a view she calls "optimistic realism" which holds that

arithmetical facts are what they are independently of us though it is impossible that there be an arithmetical fact which we could not prove (p. 204).

Jenkins' points here get at some very difficult and important issues pertaining to our understanding of realism and anti-realism. She claims that most anti-realists actually view *essential* dependence as central to their views, citing examples from Berkeley to Kant to Crispin Wright (pp. 204-205). This sort of observation is certainly accurate with regard to the mathematical constructivists, who first came to certain conclusions about how mathematical reality is constructed by us and from there drew further conclusions about the epistemic consequences of their views. Along these same lines, Jenkins claims that in reality philosophers typically form opinions about the nature of the reality that concerns them first and then take particular modal facts about our ability to know to follow as a consequence (p. 205).

Whether and how questions of modality rather than questions of essence in particular are relevant to the realism debate cannot be worked out in full detail in the present work. Regardless, we ought to muster some sort of answer to Jenkins, as her arguments would appear to threaten at their very foundation the conceptions of realism and anti-realism with which we have been working. In any case, there is reason to think that Jenkins has been unfair in her assessment of the use of modal dependence and modal independence. Consider her original example intended to establish that mathematical anti-realism in the form of constructivism is consistent with mathematical facts being modally independent of our minds. An initial doubt is one that she herself brings up, though she does not take it to be conclusive. Her example was supposed to be one where no people exist to construct the arithmetical fact in question though the constructivist could argue that the fact still obtains at that world because we impose mathematical structure on the world when we assess it for the truth or falsity of, for instance, 2 +

2 = 4. Concerning this case, Jenkins says, "the constructivist might find this too counterintuitive, and prefer to say that 2+2=4 is *not* in fact the case at [world] w because nobody exists at w to perform the requisite mathematical construction" (p. 203). In my view, that is more or less how a constructivist should respond. More likely, a constructivist will say that it makes no sense to speak of what is the case at that world since we are not there to either construct mathematical reality or not. Jenkins suggests something along these lines when she says, "perhaps our intuitions in this area can be sufficiently respected by claiming that *if* we inhabited such worlds as w, *then* 2+2=4 would be the case at those worlds" (p. 203). Jenkins leaves the matter there, but there is more we can say to cast doubt on her main claim. The consideration is simple: Modality, rather than essence, takes center stage again once we admit that the mathematical facts *would* be different, no matter which possible world were under assessment, if facts about our ability to construct mathematical reality *were* different. When we think about things that way, we see that Jenkins' example is ill-conceived—it does not truly capture a case of modal independence in the sense that matters to the mathematical realism debate. What matters is not that we, in the actual world, construct mathematical reality a certain way and then assess worlds according to whether the mathematical facts we constructed obtain there. What matters, rather, is that the mathematical facts *would* be different in possible worlds where we constructed them differently, or failed to construct them at all.

Something similar can be said about Jenkins' example of optimistic realism. She wants us to think that we might be realists about arithmetic in the sense that facts in that domain do not reduce to facts about our mental lives while nevertheless holding that it is impossible for there to be an arithmetical fact that we cannot know. Such a position is certainly coherent. But Jenkins is unfair in her treatment of what this implies for the relevance of modal considerations to the

realism debate generally. For she must admit that, were our mental capacities of a different nature, those same arithmetical facts might cease to be knowable by us. Thus, again, her example is ill-conceived as it does not do justice to the sorts of modal considerations that realists and anti-realists commonly take to be relevant to their debate. We will continue operating with our modal concepts, as Jenkins' arguments that essence ought to be the operative notion in the realism debate remain unconvincing.

3. Intuitionism and Agnosticism

3.1 Background

The Dummettian understanding of anti-realism about anything at all is modeled after a particular kind of anti-realism in mathematics known as constructivism. The mathematical intuitionists were constructivists. The intuitionists took a particular attitude towards mathematics based on their understanding of the metaphysics of mathematical reality. For them, mathematical reality is the result of our constructing it. It is not often enough appreciated that the intuitionists took their view of mathematical reality to justify the acceptance of intuitionistic logic for mathematical statements (Haack 1974). Whereas classicists typically take their logic to be prior to everything else, the intuitionists did not see the endorsement of their logic as flowing from any sort of first principles in the same way. Rather, because mathematical reality depends upon our ability to construct it, there is a special relationship between our minds and that reality such that for something to be the case we must be capable of constructively proving that it is the case. The intuitionistic idea that truth is constrained by proof gives us the standard intuitionistic semantics out of which their logic falls. The intuitionists, then, may be said to put their metaphysics first. Devitt ought to approve.

In his seminal paper "The Philosophical Basis of Intuitionistic Logic," Dummett turns the intuitionistic model on its head (1983). He conceives of the intuitionist as endorsing a certain meaning theory for mathematical statements such that what matters to meaning are not evidentially unconstrained truth conditions, but instead evidentially constrained *proof* conditions. Dummett notes that this justification for intuitionistic logic is rooted in the theory of meaning rather than in any particular metaphysical view of mathematical reality. Thus, he concludes, intuitionistic anti-realism could be extended to any area of inquiry where an anti-realist theory of meaning were appropriate. For the rest of his life he defended the view that a realist theory of meaning, according to which our sentences have evidentially unconstrained truth conditions, could not work for any region of discourse containing undecidable sentences.

Regardless of which justification for the intuitionistic outlook one prefers, there are a few things about that outlook which must be highlighted in order to appreciate the complications involved in extending intuitionism to non-mathematical areas. Perhaps the most important of these is agnosticism—remaining neutral on certain matters, citing the fact that we simply do not know what to think about them. Another point to highlight is the view that truth amounts to provability, while falsehood is equated with the impossibility of proof. Specifically, for the intuitionists, something is impossible for us to prove if we are in possession of a proof showing that the assumption that it is provable (i.e., true) leads to contradiction. In the extension of intuitionism to the contingent, empirical realm, mention of contradiction gets put to one side and focus centers on the impossibility of proof. Proof itself gets generalized as some form of ideal justification.

The intuitionistic semantics in terms of proof conditions leads, most notably, to the rejection of two coveted stand-bys of classical logic: The law of excluded middle and the rule of

double negation elimination. It is not hard to see why the intuitionist would reject them. In accordance with the intuitionistic rendering of truth and falsity, the intuitionistic negation operator functions to tell us that a proof of the negated item is not possible. The law of excluded middle in intuitionism amounts to the claim that P is either provable or impossible to prove. It certainly seems conceivable as an epistemic state of affairs, however, that we might lack good reason to think that P is provable while at the same time having no good reason to think that it is impossible to prove. So the intuitionist is left without adequate grounds to assert either disjunct and hence without adequate grounds to assert the disjunction itself. This is the agnostic attitude. Similarly, intuitionistic double negation elimination would amount to the claim that, given the impossibility of proving that it is impossible to prove P, we should conclude that it is possible to prove P. Again, conceivably, we might be in a situation where we have no good reason to think that P is provable despite having proven that it is impossible to prove that it is impossible to prove P (~ ~ P). Hence the intuitionist rejects the classical inference rule (for good general treatments of intuitionistic logic and the constructivist philosophy which motivates it, see Dummett 1977, chapter 4 of Khlentzos 2004, and chapters 2 and 3 of Kapsner 2015).

The account of logical revision just given likely makes the reader wonder how anyone could question the fact that a semantic anti-realist, who endorses an epistemically constrained notion of truth, must reject classical logic. Focus on excluded middle. In the above account, we made reference to a state of affairs such that there was no telling either way whether P was provable or provably impossible to prove. It is in such circumstances that an intuitionist would refuse to assert the disjunction on the grounds that there is not sufficient reason to assert either of the disjuncts, given the intuitionistic semantics. The major difficulty here comes in when the intuitionist attempts to explain what it is about this situation which makes it so that we should

withhold judgment on an instance of P or not-P. In this way, we are brought back to the analysis of the notion of undecidability that features so prominently in the intuitionist's arguments.

In our discussion of decidability and undecidability from the previous chapter, we emphasized that intuitionism has a difficulty with undecidability as strict unknowability. As it turns out, an undecidable sentence is one about which the intuitionist takes an agnostic attitude. The intuitionist cannot, as we saw, conceive of undecidability as an inability to know that a sentence is true combined with an inability to know that it is false. For this leads almost immediately to contradiction, given the intuitionistic understanding of truth, falsity, assertion, and negation. Similar reasoning makes it so that the intuitionist cannot say of a particular sentence that it will never be decided (Cf. Williamson 1994). When presented with an undecidable, the intuitionist chooses to make no decision as to what the truth value is. Perhaps more importantly, without a suitable guarantee that we would ever discover the sentence's truth value if we tried hard enough to get one, the intuitionist does not even admit that the sentence has a determinate truth value.

Recall that what matters is not just whether an effective procedure for determining a sentence's truth value exists, but whether we are in possession of such a procedure. Accordingly, what matters is not just whether a sentence is effectively decidable, but whether we know that it is effectively decidable (Cf. Rasmussen and Ravnkilde 1982, p. 388). The Dummettian anti-realist claims that there is a special relationship between truth and our ability to know the truth. If effective decision procedures are supposed to be relevant to Dummettian anti-realism, we should do some work to better understand the bearing of such procedures on our ability to know. More clarity is needed concerning what the intuitionistic anti-realist attitude is as well as the motivation for that attitude.

Sanford Shieh, on his way to a definition of undecidability that a generalized intuitionist can use, suggests simply that "a sentence is Undecidable just in case it is not presently decided, *and*, we don't have an Effective Procedure for it" (1997, p. 61). He then goes on to tell us that,

> Each Decidable sentence is such that *we know that we can either recognize it as true or recognize it as false,* either because we have already given a proof or a refutation of it, or because we have a Procedure. Moreover, it is plausible that an Undecidable sentence is such that *we don't know that we can either recognize it as true or recognize it as false.* (p. 61, emphasis in original)

He then lands upon the following definition of undecidability:

> **undecidable$_S$:** A statement is *Undecidable* just in case it is not Decidable; i.e., we don't know, or have adequate reason to think, that either we can recognize it as true or we can recognize it as false. (p. 61)

Notice that by now, explicit mention of effective procedures has fallen away. Indeed, one of Shieh's goals is to establish that we don't need the strict notion of an effective decision procedure in order to make sense of Dummettian anti-realism. Concerning the anti-realist's refusal to apply bivalence to undecidables, Shieh contends that, "the conclusion which follows from accepting that anti-realism applies to Undecidable sentences is that we *don't know* that they satisfy the principle of bivalence" (p. 63). This is why the intuitionist refuses to endorse the unrestricted validity of the law of excluded middle. The idea is just that we don't know that either P or not P for all P. And as Shieh points out, the knowledge operator is crucial here (p. 64). For the intuitionist, just like the classicist, cannot assert the negation of the law of excluded middle without running into contradiction. The intuitionist, rather, asserts the negation of the

claim that we *know* the law of excluded middle. The intuitionistic line is that we do not know that P or not P is true, but we also do not know that P or not P is false.

Shieh emphasizes "our ignorance about the extent of our capacity to recognize the truth and falsity of Undecidable sentences" (p. 63). Crispin Wright pursues a similar line as he works to generalize intuitionism in Dummett's vein. Wright offers a description of an epistemic conundrum he calls "being in a quandary":

> We do not know, do not know how we might come to know, and can produce no reason for thinking that there is any way of coming to know what to say or think, or who has the better of a difference of opinion. (2001, p. 71).

Wright's ambitions for his conception of quandary are quite high. He believes he can use the notion to formulate an intuitionistic answer to a number of issues. Among these is the problem of vagueness. Wright's strategy is to analyze indeterminacy as an epistemic matter. He then claims to find this kind of indeterminacy in cases of vagueness as well as in cases of undecidable sentences, not to mention matters of personal taste. The end product is a model for extending intuitionism from mathematics to other areas. I will present the model first, and turn to its application to vagueness when considering some objections to the model later. I begin by situating the Quandary work in the context of Wright's earlier work.

3.2 Wright's Basic Revisionary Argument, Round 1

In his *Truth and Objectivity* (1992), Crispin Wright presents an argument for why someone who endorsed an epistemically constrained notion of truth ought to revise classical logic. In the relevant passage, Wright has just argued that an anti-realist who endorses Hilary Putnam's conception of epistemically constrained truth in terms of idealized warranted assertibility is going to run into trouble in cases where neither the sentence in question nor its negation would

be justified under "ideal epistemic circumstances" (Wright 1992, p. 39). The difficulty stems from the acceptance of something Wright calls the Negation Equivalence (p. 40):

(NE): "It is not the case that P" is true if and only if it is not the case that "P" is true.

Wright then argues that similar considerations cause trouble for all kinds of evidentially constrained notions of truth. To illustrate, he has us consider what is perhaps the most basic kind of evidential constraint we could place on truth, something he calls Evidential Constraint (p. 41):

(EC_1): If P is true, then evidence is available that it is so.[9]

He then claims that the semantic anti-realist should, if anyone should, admit the possibility that, for certain statements, there might be no evidence at all for either the given statement or its negation. It turns out that Wright has in mind the idea that, "not every issue can be guaranteed to be decidable, even under epistemically ideal circumstances" (p. 41). We will take issue with Wright's conceptions of undecidability in due time. For now, let's follow his reasoning to see where it goes. He points out that, since no evidence is available for the sentence P under consideration, EC makes it so that we must conclude that "it is not the case that P is true" (p. 41). I have quoted Wright's exact words to bring out the facts that 1) he does not say we need to conclude that P is false and 2) interestingly, he doesn't bother to put the relevant strings of characters in quotation marks like he does in his definition of NE. Regarding the first point, as long as we have not accepted bivalence, P's not being true and its being false are not equivalent. Regarding the second, there seems to be no relevant difference so we will move on.

Given that it is not the case that P is true, NE then leads us to the conclusion that not-P is true. And so, Wright claims, the semantic anti-realist appears forced to admit that a lack of

[9] The subscript '1' serves to distinguish this version from EC from other versions Wright gives in other works. In general, throughout the present work I try to keep to the notations of the authors discussed. The motivation is to make it easier to reference their idiosyncratic concepts, definitions, etc.

justification for a given sentence translates into the presence of justification for its negation. As Wright explains, the trouble for the anti-realist is that she can't simultaneously endorse NE, EC, and the occurrence of the special notion of undecidability described above according to which there is nothing to push us in the direction of a decision as to whether P or not-P (p. 41). And the possibility of just such a state of affairs, where "some statements may be such that no evidence bearing upon them is available either way," Wright thinks is needed, along with epistemically constrained truth, to "explain the semantic anti-realist's characteristic refusal to allow the unrestricted validity of the principle of Bivalence" (p. 41). The trouble is that the admission seems to be directly at odds with the requirement, stemming from the combination of EC and NE, that the anti-realist view a lack of evidence for some P as the presence of evidence for not-P.

Wright thinks that the result is only problematic if the anti-realist also endorses the law of excluded middle. In fact, he takes the inconsistency of EC, NE, the possibility of so-called "neutral" states of information which he loosely associates with some sense of undecidability, and the law of excluded middle, to give the semantic anti-realist sufficient reason to reject the law of excluded middle. He reasons that NE and EC certainly force the acceptance of (p. 42):

(A): If no evidence is available for P, then evidence is available for its negation.

In classical logic, A would indeed be equivalent to (p. 42):

(C): Either evidence is available for P or evidence is available for its negation.

But in order to run the equivalence, we need the following instance of the law of excluded middle, which the anti-realist is under no particular pressure to accept (p. 42):

(B): Either evidence is available for P or it is not.

Wright claims that it is B which is inconsistent with his undecidability concept that it might be the case that no evidence is available either way for P or its negation. Since B is an instance of

the law of excluded middle, Wright concludes that the semantic anti-realist ought to revise classical logic. In Wright's words,

> The thesis of EC, that truth is essentially evidentially constrained, *must* enjoin a revision of classical logic, one way or another, for all discourses where there is no guarantee that evidence is available, at least in principle, to decide between each statement of the discourse concerned and its negation. (p. 43)

We see again Wright's pesky formulation of undecidability. Note that his notion doesn't appear to be utter unknowability. And his notion does not on its surface seem to have much to do with undecidability understood as lack of an effective decision procedure, though as we will see, Wright seems to think that it does.

3.3 Wright's Basic Revisionary Argument, Round 2

Much has been written on the argument of Wright's which we have just canvassed. We will return to a discussion of the relevant literature. For now, we should move on to explore his formula for being an intuitionist about anything at all. Earlier we brought up Wright's notion of being in a "Quandary." He formulates this concept in an attempt to provide the proper justification for his notion of undecidability, which as we have just seen he uses in conjunction with the acceptance of an epistemically constrained notion of truth and his Negation Equivalence to justify the revision of classical logic.

Wright thinks that the Negation Equivalence falls out of our standard notion of truth. He thinks that it is one of the central features which must be had by any candidate truth concept. Thus he wants to keep it when he formulates his own epistemic truth predicate (which we are not discussing here). In other work, however, Wright streamlines his argument in such a way that the Negation Equivalence does not come up. Wright sees the different versions of the argument

as essentially the same and calls it the "Basic Revisionary Argument" (2001, p. 65). To get this argument going, he first gives us a formulation of an undecidability concept,

(NKD): $\sim K \forall P(FeasK(P) \vee FeasK(\sim P))$

where Kφ means that φ is known and FeasK(P) means, in essence, that P is feasibly knowable ('NKD' appears to stand for something along the lines of "not known to be decidable"). (Wright 2001, p. 65). Wright's notation is somewhat awkward—one wants to see a possibility operator where we are talking about what is know*able* rather than what is known, and Wright uses 'K' in one place non-modally and in the next place modally. It will help us to remain faithful to what he wrote if we use his own notation. Others writing on the same issues use different notation to express concepts similar to Wright's.

This time, Wright's epistemic constraint on truth is (p. 66):

(EC_2): $P \rightarrow FeasK(P)$

Now, if we assume the law of excluded middle,

(LEM): $P \vee \sim P$

we get something inconsistent with NKD, specifically, that either P is feasibly knowable or its negation is: (FeasK(P) v FeasK(~P)). Wright claims that the anti-realist ought in these circumstances to give up LEM since EC_2 is the core of the semantical anti-realist view and NKD is, he claims, "incontrovertible" (2001, p. 66).

We ought to question Wright's justification for thinking that the anti-realist is beholden to NKD or anything like it. The reason he offers is, "does [NKD] not merely acknowledge that, relative to extant means of decision, not all statements are decidable?" (p. 66). On the very next page, he admits that there is difficulty here. In order for his revisionary argument to have any force, the anti-realist's epistemic constraint on truth needs to be consistent not only with NKD,

but also with the justification for endorsing NKD in the first place (Wright 2001, p. 67). Much of the discussion in the literature on why and whether a Dummettian anti-realist needs to endorse intuitionistic logic centers on the question as to whether the anti-realist's epistemic constraint on truth can be made compatible with some pertinent notion of undecidability. As we began to see in the previous chapter, insufficient care has been exercised in handling the notion of undecidability, with it sometimes being defined in terms of a lack of knowability, sometimes a lack of feasible knowability, sometimes a lack of an effective decision procedure, and so on. We have seen that that the anti-realist cannot consistently think that there is a sentence P such that it is impossible to prove P and it is impossible to prove not-P. I submit that coming up with a benign version of the concept of undecidability which does the work the revisionist anti-realist wants it to do is one of the chief difficulties facing the anti-realist.

So, then, what sort of justification does Wright think the anti-realist could accept for NKD? He goes through more than one option. What he is on the lookout for is some "principle of agnosticism" or "modesty" (2001, p. 67). His first suggestion is (p. 67):

> **(AG):** P should be regarded as unknown just in case there is some possibility Q such that if it obtained, it would ensure not-P, and such that we are (warranted in thinking that we are) in no position to exclude Q.

A more perspicuous way of putting AG might be: If we are not in a position to exclude the epistemic possibility of not-P, then we do not know that P. Furthermore, the concepts of epistemic modesty and agnosticism are more aptly associated with something like NKD rather than Wright's AG, which only tells us when we should think we don't know something generally, not specifically we are in no position to know either P or its negation. All the same, this is Wright's argument, not ours, so we will use his terminology.

Wright rejects AG as the proper motivation for NKD on the grounds that, in order to justify NKD, it must be epistemically possible that $\sim\forall P(FeasK(P) \vee FeasK(\sim P))$, which is not epistemically possible since that sort of claim, as we have seen again and again by now, gives rise to contradiction if we admit that truth is epistemically constrained (in this case by EC_2) (Wright 2001, p. 68). Hence AG cannot help someone who accepts EC_2 justify NKD. Jon Cogburn, in a different context and using a different formalism, points out that intuitionistically there is no contradiction derivable from $\sim\forall P(FeasK(P) \vee FeasK(\sim P))$ and the epistemic constraint (Cogburn 2002, p. 236). This is because, in intuitionistic logic, an existential negative is not derivable from a negative universal. So it would appear that Wright could have admitted AG as the justification for NKD and claimed that since the assumption of classical logic gives us a contradiction, we should just reject classical logic. But we will push forward.

It is sometimes said (Cf. Salerno 2000) that whatever reasoning is used to establish the rejection of classical logic and the acceptance of intuitionistic logic, it must be intuitionistically acceptable. There are times when it matters whether an intuitionist would accept a particular line of reasoning. For instance, in an effort to convince someone to reject classical logic, it would do no good to show that particular assumptions lead classically to contradiction when those same assumptions also lead intuitionistically to contradiction. The problem then would be with the assumptions, or perhaps with both classical logic and intuitionistic logic. We see the foregoing kind of point made in Salerno (2000, p. 214). Salerno goes on, however, to argue that whatever reasoning is used to adjudicate between intuitionistic and classical logics, it must be intuitionistically acceptable. His reasons are that the two sides would not be able "to communicate at all" if they could not agree on the logic of the meta-language in which the debate were being hashed out, and further that "such inferential restrictions will not load the dice

in favor of the revisionist" (2000, p. 221). Since all intuitionistic inferences are classically acceptable, but not vice versa, the intuitionist and the classicist could agree on intuitionistic reasoning in the meta-language. Salerno's point here is fair enough, but it is perhaps too fair. If we classically reason ourselves into a predicament, one way out of the predicament may be to reject the reasoning that led us there. For if, as Tennant has put it, "the aim is to perform a *reductio* of the realist's own principle of bivalence…[then] if strictly classical moves help one to do this, one could say that the realist is simply digging his own grave" (Tennant 1997, p. 186).

Keep in mind that Wright is looking for an anti-realistically admissible justification for NKD so as to warrant the refusal to endorse LEM. Philosophers engaged in debating this issue typically speak of *rejecting* LEM, but they are also typically quite cognizant of the fact that the argument for the revision of classical logic is not an argument for the conclusion that LEM is false, i.e., ~ (P v ~P), which of course leads to contradiction in both intuitionistic and classical logic (Cf. Wright 2001, p. 67). The fact has been stated so many times the reader is likely weary of hearing it, but there is no avoiding such repetition as the fact constitutes the most salient difficulty in formulating a coherent anti-realism based on the idea that truth is evidentially constrained. At any rate, the point bears repeating that the intuitionist is agnostic toward LEM. What we are after here is Wright's particular explication of the proper sort of agnostic attitude.

In order for Wright's explication in terms of what he calls Quandary to apply to the Dummettian realism debate, one must endorse a particular understanding of what is at issue between Dummett's realist and Dummett's anti-realist. The issue, recall, is how to treat undecidable sentences. The realist endorses bivalence for undecidables while the anti-realist does not. It is important to be clear about what this means. As Rasmussen and Ravnkilde are astute to point out, it is not the thought that there are two possible truth values which matters

here, but rather the having of a truth value, any truth value, *determinately* (1982, p. 388). This sort of view was advanced in the previous chapter. Wright conceives of the issue in the same way. What he notices is that the issue of determinacy comes up also in the vagueness debate. Hence he thinks he can give a version of intuitionistic anti-realism which would make sense of why, perhaps, in cases of vagueness as well as in the more standard cases of undecidables, one might want to revise classical logic. So here he is attempting to give a certain epistemic gloss on indeterminacy.

One way of thinking of indeterminacy which Wright rejects is that conception according to which being indeterminately true is equated with either lacking a truth value altogether or having some other possible truth value besides truth and falsity. His main complaint about the conception just given is that it attempts to analyze indeterminacy as the determinate possession of some third feature. He argues that such a state of affairs doesn't accommodate our intuitions about actual cases of indeterminacy, where, he says,

> we are baffled to choose between conflicting verdicts about which polar verdict applies, rather than as cases which we recognise as enjoying a status inconsistent with both…the indeterminacy will be initially manifest not in (relatively confident) verdicts of indeterminacy but in (hesitant) differences of opinion…about a polar verdict, which we have no idea how to settle—and which, therefore, we do not recognise as wrong. (2001, p. 70).

The second version of AG which Wright offers is of little concern to us, as he ends up rejecting it for reasons not worth laying out here. We have almost all of the context we need in order to understand the version of AG Wright eventually endorses and takes to justify NKD. We just need to fit Wright's notion of Quandary into all of this. According to what he calls the

Quandary View, "the root characterization of indeterminacy will be by reference to *ignorance"* (2001, p. 71). Thus we get the idea that P is indeterminately the case just when we don't know that P, we don't know how to find out whether P, and we don't currently possess a reason to think that there is a way to come to know that P. In such a situation, we are in a Quandary.

More precisely, he says, "A proposition P presents a quandary for a thinker T just when the following conditions are met:

(i) T does not know whether or not P

(ii) T does not know any way of knowing whetheror not P

(iii) T does not know that there is any way of knowing whether or not P

(iv) T does not know that it is (metaphysically) possible to know whether or not P.

(2001, p. 92).

Importantly, he does not allow in cases of Quandary that i – iv be satisfied by (p. 92):

(v) T knows that it is impossible to know whether or not P.

For, as we know, for an anti-realist v leads to contradiction. Rather, he tells us, "a quandary is uncertain through and through" (p. 92). Thus we have Wright's initial characterization of the sort of agnostic attitude appropriate for intuitionism. In the context of his idea of Quandary, Wright offers a new version of AG. This time, with the view in mind as to how to treat LEM in particular, and to avoid the problem that Quandary doesn't seem to help with explaining an agnostic attitude toward atomic (non-compound) statements (2001, pp. 71-72), he frames things in terms of compounds (p. 76):

(AG+): [a compound statement] is known only if there is an assurance that a suitably matching distribution of evidence for (or against) its (relevant) constituents may feasibly be acquired.

Wright's language here is confusing. He uses such phrases as "pattern of evidence" and "pattern of knowledge" to describe the requirement that, in order to be justified in believing a compound statement, we must be appropriately justified in believing its constituents (p. 77). The requirement on knowledge given by AG+ is not clearly different from considerations we laid out earlier on when explaining how an intuitionistic proof-conditional semantics tells us, in rather straightforward terms, that we are in a position to know a disjunction such as (P v ~P) just in case we are in a position to know one of its disjuncts. Wright takes his AG+ to give us the right justification to "properly profess ignorance of such a compound statement in any case where one has no reason to offer why an appropriate pattern of knowledge should be thought achievable" (p. 77) Concerning LEM, he says,

> when the truth of the ingredient statements is taken to involve evidential constraint, then [knowledge that ~ ~ (P v ~P)] does not in general amount to a reason to think that the appropriate kind of evidence for one disjunct or the other must in principle be available in any particular case…where we do not know what to say, do not know how we might find out, and can produce no reason for thinking that there is a way of finding out or even that finding out is metaphysically possible. (2001, pp. 77-78).

Let us return to Wright's generalized notion of undecidability, NKD, and his more recent analysis of evidential constraint, EC_2. This all began when Wright started looking for an agnostic principle that would justify NKD without running into contradiction in association with EC_2. Ultimately, he claims that the semantic anti-realist, whose central claim is that truth is evidentially constrained, should be compelled to revise classical logic in favor of intuitionistic logic whenever she finds herself in a case of Quandary. In cases of Quandary, Wright claims, NKD is justified. Whatever reasons are offered for the anti-realist's EC_2, if we are in a

Quandary, then according to Wright we should endorse NKD for the subject matter in question and, hence, reject LEM via the Basic Revisionary Argument given earlier.

Wright argues that cases of Quandary are real, and that their existence makes it the case that sentences about the area where Quandary applies will be undecidable in the specific sense of NKD. Wright formulates his conception of Quandary in the context of the problem of vagueness. He believes he has adequately described the epistemic situation we find ourselves in with regard to vagueness in such a way that an intuitionistic answer to the vagueness problem is justified and coherent. That epistemic situation is what he calls Quandary. He then claims that intuitionism is motivated anytime we find ourselves in Quandary, which may happen in mathematics as with standard mathematical intuitionism, but also in the contingent, empirical realm. Indeed, what has been lacking amidst efforts to extend the intuitionistic model of anti-realism from mathematics to other areas is precisely what Wright provides—a description of the way in which we might be said to be in the same position with regard to contingent, empirical statements as the intuitionists believed themselves to be in with regard to the statements of mathematics. In order to more fully understand what Wright thinks he has done, and to consider some important objections to his view, we will need to say something about the problem of vagueness.

3.4 Vagueness, Indeterminacy, Intuitionism

Imagine that we have a pile of grains of sand. We know that it is a pile of sand because it fits all of the criteria for being a pile of sand: It is composed of many individual sand grains left to rest on top of one another in such a way that the collection of grains, loosely stacked, rises numerous feet up from the ground which is at its base. It would appear that we can tell that this is surely a pile of sand. It is a pile of sand in a way that a mere scattering of a few sand grains on the

ground fails to be a pile of sand. At the very least, we know that one sand grain by itself does not count as a pile. Imagine now that we take our pile of sand and begin to pluck grains from it one at a time. After each plucking we can ask again, "Does this count as a pile of sand?" It seems, at first, absurd to suggest that there is any cutoff point, that the removal of any one single grain would make all the difference between a pile and a non-pile. But still, it would also seem that we know that there must be such a cutoff because, after all, via this process of removing single grains we do eventually get to a point where we know that we do not have a pile of sand anymore. So, then, which grain of sand is the cutoff? A further issue, harped upon by Wright, is that if there is a cutoff, then it would appear that what counts as a pile is actually precise, which conflicts with the intuition that we were dealing with a vague predicate to begin with.

The problem of vagueness is difficult—we will not attempt to solve it here, though we will be assessing Wright's purported solution. One solution to the problem is to say that there is no fact of the matter as to where the cutoff is. Such a solution would involve rejecting the principle of bivalence for sentences containing vague predicates—some such sentences, those involving borderline cases where it is unclear whether the vague predicate properly applies, could not be said to be determinately either true or false. Another solution, endorsed by Williamson (1994), is to say that sentences with vague predicates do have determinate truth values though we are unable to determine what those values are. All of this should sound familiar, because the same sorts of issues come up with vagueness as come up in discussions of Dummettian realism and anti-realism. So it should perhaps come as no surprise that an intuitionistic way of dealing with the problem of vagueness has been attempted. Perhaps the first suggestion of such an attempt comes from Putnam: "An idea that has occurred to me, but that I cannot claim to have thought through, is the following: treat vague predicates (e.g., 'bald') just

as undecidable predicates are treated in intuitionistic logic" (1983, p. 285). He gives a brief sketch of how the intuitionistic answer goes. The problem is framed as follows (p. 285):

(1) (n) (If a man with n hairs is bald, then a man with n+1 hairs is bald)

(2) A man with 0 hairs is bald

(3) (n) (A man with n hairs is bald)

Putnam says (3) is clearly false, and (2) is clearly true, so we must reject (1). But if we negate (1), in classical logic it follows that there *is* some value for n such that someone with n hairs counts as bald and someone with n+1 hairs does not count as bald. That existential claim is what Wright calls "the *unpalatable existential"* (2001, p. 64), and it's supposed to be unpalatable because it seems to say just that the predicate in question is not vague at all, but in fact precise (Wright 2001, p. 63). Putnam's trick here is to point out that in *intuitionistic* logic we are not forced to admit that there is a precise cutoff point because it does not follow from the denial of the universal claim (1) that there is any particular n satisfying the negation of the conditional, and furthermore to negate a conditional in intuitionistic logic is not to provide oneself with the antecedent of the conditional and the negation of the consequent (Putnam 1983, p. 285).

Stephen Read and Crispin Wright (1985) argue that Putnam's trick doesn't work in the simple manner Putnam suspected it did. They begin with the thought that the applicability of a truly vague predicate cannot be said to have a precise cutoff point. From that thought they reason that we must deny the unpalatable existential (though Wright was not yet calling it that), which in classical logic will give us (1), which will as above force the conclusion that every man is bald (Read and Wright 1985, p. 56). In response to the Putnamian solution that we reject classical logic in favor of intuitionistic logic, Read and Wright show that we also get absurdity

on intuitionistic logic. For if we deny the unpalatable existential, this leads intuitionistically to (p. 57):

(n) [~BM(n+1) → ~BM(n)]

where 'BM(n)' means that a man with n number of hairs is bald. Beginning with a clear case of a non-bald man, the conditional just stated will eventually lead us to the claim that someone with no hair is not bald (p. 57). Read and Wright conclude that intuitionistic logic cannot help us avoid the sorites paradox which they view as generated from the basic thought that if a predicate is vague, then we must deny the existence of a precise cutoff for its application conditions (p. 58).

Years later, Wright takes up the banner of an intuitionistic solution to the problem of vagueness himself. Notice that Williamson, who argues that a precise cutoff does exist though we cannot know where it is, would not find the unpalatable existential unpalatable. This is because Williamson can still make sense of vagueness by construing it as an epistemic, rather than a metaphysical phenomenon. Ironically, Wright is inspired by such an approach, that of the so-called epistemicists, to understand vagueness epistemically but in the vein of the intuitionists. Hence we get his Quandary view according to which indeterminacy in general, vagueness being just one example, should be understood epistemically rather than metaphysically. Wright's solution to the difficulty presented by Read and himself is now to say that, "the vagueness of *F* should be held to consist not in the falsity of the unpalatable existential claim but precisely in its association with quandary" (Wright 2001, p. 78). His idea is that we are in a quandary as to where the precise cutoff point is because among the borderline cases where the cutoff, if there is one, must be, "we do not know whether to endorse them, do not know how we might find out,

and can produce no reason for thinking that there is, or even could be, a way of finding out [whether the predicate in question applies]" (p. 79).

Wright, then, has basically come around to the view that the issue he and Read found with Putnam's intuitionistic solution to the sorites paradox itself has an intuitionistic solution. The solution is to refuse to endorse the unpalatable existential and refuse to endorse its denial. Then we do not find ourselves saying that there is a sharp cutoff to the application of a vague predicate, which in the absence of the epistemicists' understanding of vagueness seems to contradict the predicate's being vague in the first place, and we also do not find ourselves asserting that it is *false* that there is a sharp cutoff, so we avoid the absurdity brought on by the sorites. Here we have the familiar intuitionistic agnosticism captured in the characteristic refusal to admit that refusing to endorse a claim requires endorsing its falsity.

Wright proposes the idea of being in a Quandary as a particular way in which a claim might be undecidable while maintaining consistency with an epistemic constraint on truth and justifying the rejection of classical logic. Intuitionistic logic would then be available in cases of vagueness precisely because borderline cases, at least according to Wright, present Quandaries.

Sven Rosenkranz objects to Wright's treatment of vagueness. Rosenkranz argues that while the Quandary conception of undecidability may be consistent with constraining truth epistemically, justification for thinking that one of them is the case undermines justification that the other one is the case. He says,

> The resources needed to justify an epistemic constraint on truth for statements of a given class are likely to deprive us of the resources needed to show that it is unknown that those statements are either knowably true or knowably false, and vice versa. (2003, p. 462).

Presumably Rosenkranz thinks that this is a problem for Wright's way of dealing with vagueness because we lose motivation for the intuitionistic logic at the heart of Wright's treatment.

One of Wright's responses to Rosenkranz is just to point out that we don't need an evidential constraint on truth in order to motivate a refusal to endorse bivalence. Wright claims that "*any* sufficient case, EC-driven or not, for agnosticism about Bivalence over the predications in a Sorites series is a reason for agnosticism about the existence of a sharp cut-off," and he thinks that borderline cases involving vague predicates by themselves give us some reason not to endorse bivalence (2003, p. 472). In our current context, however, there is reason to be worried about Rosenkranz's objection since it would appear to undermine Wright's Basic Revisionary Argument and hence the entire Quandary-centered generalization of the intuitionistic brand of anti-realism.

Rosenkranz's reason for thinking Wright to be in trouble is that the epistemic constraint on truth may very well tell us of two items in a series that if one of them is F, then we are capable of coming to know that it is, i.e.,

(a): $Fx \rightarrow \Diamond KFx$

and that if its successor in the series is not F, then we are capable of coming to know that as well, i.e.,

(b): $\sim Fx' \rightarrow \Diamond K \sim Fx'$

but these together do not imply

(c): $\exists x(Fx \wedge \sim Fx') \rightarrow \exists x \Diamond K(Fx \wedge \sim Fx')$

for, Rosenkranz wants us to think, we need some assurance that our ability to know of some x that it is F does not somehow preclude our ability to come to know of x' that it is not F (Rosenkranz 2003, p. 454. Let us ignore that he has introduced quantification—on p. 455 he

mentions that the quantified sentence would be "true a priori" based on the schema). And it seems, for one thing, that borderline cases are supposed to be such that we lack that kind of assurance since, after all, they are supposed to put us in a Quandary with regard to the application of F (Rosenkranz 2003, p. 456).

Wright's answer to Rosenkranz is to claim that he could just as well make due with this conditional instead, which does follow from (a) and (b):

(W): $\exists x(Fx \wedge \sim Fx') \rightarrow \exists x(\Diamond K\, Fx \wedge \Diamond K \sim Fx')$

Wright says that he has here avoided to his own satisfaction the difficulty brought up by Rosenkranz by avoiding any mention of knowledge of a conjunction (Wright 2003, p. 470).

Rosenkranz's objection would seem to cut deeper than Wright is willing to admit. Consider borderline applications of vague predicates. Given Wright's formulation of Quandary, we may indeed be in one. Yet in cases such as these where, as Wright continuously emphasizes, we do not even know that it is metaphysically possible to know, what reason could we possibly have for thinking that any facts about whether a vague predicate properly applies to a borderline case must be such that, if they obtain then we are capable of knowing that they do? That, I take it, is what Rosenkranz's worry boils down to. But I think it is answerable, both in vague cases and other purported instances of indeterminacy understood as Quandary. The point is just that the intuitionist, being agnostic, is unwilling to admit that we *are* capable of knowing, for any borderline case, that the predicate in question properly applies. If I am right about what Rosenkranz's concern is, he has simply thought that any justification for an epistemic constraint must come along with justification that we are able to know whether P or not-P for the domain in question. The epistemic constraint, however, is conditional upon P's being the case, and the intuitionist, if he has his, will here respond that he has not committed himself to P's either being

the case or not to begin with, and so is under no pressure to think himself justified in believing that he can come to know whether or not P. Lacking such justification, he can safely take himself to be justified he is in a Quandary while being justified in endorsing the epistemic constraint.

We might take issue with Wright's claim that being in a Quandary requires being in a position such that we do not have a reason to think that knowledge in that area is "even *metaphysically possible*" (2001, p. 75). He doesn't really say anything about what he means by that phrase. I would suggest that the phrase threatens to turn into a metaphysical phenomenon the indeterminacy he is attempting to analyze as an epistemic one. What does it take for knowledge to be metaphysically impossible? And for whom? Usually it matters that the possibility in question be aimed specifically at beings like us. If Wright intends something along the lines of "we do not know that *any* being could ever come to know P because we do not know that evidence for P ever could exist," then he would seem to be saddled with the result that the indeterminacy is actually metaphysical. For surely some being, perhaps divine, could come to know any fact which obtained.

Wright adds the metaphysical possibility clause in order to deal with an objection he himself raises. He says,

> Suppose I do not know, and do not know how I might know, and can produce no reason for thinking that there is any way of coming to know that P; likewise for not-P…still, might I not have all those three levels of ignorance and still know that it is the case *either that P is knowable or that its negation is?*...all that follows…is that we are, as it were, thrice unwarranted in holding either disjunct. (2001, p. 74).

He goes on to say that there is an ambiguity involved in the claim that we do not know whether A or B: We might mean that we know one of them is true but we don't know which, or we might mean that we do not know that one of either A or B is true. He calls the former interpretation "weak" and the latter interpretation "strong" (p. 74). He then frames the objection currently under consideration as claiming that Quandary as so far specified only entails that we don't know whether P is knowable or not-P is knowable weakly interpreted, whereas the Basic Revisionary Argument requires that we don't know the disjunction strongly interpreted. Wright's attempted solution is to add the clause about knowledge being metaphysically possible. He claims that inasmuch as we do not know that knowledge pertaining to a given P is possible in that sense, we lack justification for endorsing the decidability concept under consideration:

> When a difference of opinion about a borderline case occurs, one who feels that she has no basis to take sides should not stop short of acknowledging that she has no basis to think that anything amounting to knowledge about the case is metaphysically provided for. And if that is right, then there cannot be any residual ground for regarding [FeasKP v FeasK~P] as warranted. (p. 75).

Wright takes his addition of the metaphysically possible knowledge proviso to establish that we do not know the decidability conjunction to hold in the strong sense. Based on what he says, however, it is difficult to see how he has done any more than to insist upon the strong interpretation. What else could it mean, after all, to say that we have no basis for thinking that "anything amounting to knowledge about the case is metaphysically provided for"? Such insistence is at least consistent with, and even motivated by, an intuitionistic attitude toward the matter in question. Given good reason to take such an agnostic attitude toward a subject matter, one might be quite justified in endorsing intuitionistic anti-realism in that area. What must be

emphasized is that the agnostic attitude is likely to be justified by metaphysical considerations about the nature of the subject matter at hand, as it was for the mathematical intuitionists.

4. The Logic of Logical Revision

4.1 Wright and Wrong

Joe Salerno (2000) argues that both Dummett and Wright fail to give good arguments for why the Dummettian anti-realist ought to revise classical logic. In the same paper, Salerno puts forward his own argument for logical revision based on anti-realist premises. Jon Cogburn, in his (2002), offers criticism of Salerno's interpretation of Wright, but then gives an objection to both Salerno's argument and Wright's argument. Cogburn also, in his (2005), criticizes Neil Tennant's (1997) interpretation of Dummett's argument for logical revision as well as Salerno's interpretation of Dummett's argument for the same conclusion. Cogburn then in the same paper explains how he thinks we should interpret Dummett. In what follows, I will survey and critically respond to this dialectic as I see it.

> Early in his paper, Salerno asks whence the pressure to revise classical logic comes: Does the obligation arise in light of a[n] anti-realist *refutation* of a classical principle? Or does it arise because once we embrace anti-realism some classical principles turn out merely to lack a certain kind of justification or privileged epistemic status that logicians standardly claim for them? (2000, p. 212).

By now it should be obvious that, in response to the first question, the reason for revision doesn't follow from some sort of refutation of the law of excluded middle. Down that way lies contradiction. In light of Wright's discussion, it would seem that the answer to the second question is yes, the semantic anti-realist lacks the requisite justification for the laws of excluded middle and double negation elimination.

Central to all of the aforementioned arguments for logical revision is some version or another of what I have been calling an *undecidability concept*. Undecidability, to repeat, is a version of unknowability. We said in the previous chapter that a sentence's undecidability amounts to a particular kind of restriction on our ability to come to know the sentence's truth value. More accurately, when we know that a sentence is undecidable on the effective procedure understanding of that notion, we know something about what we *don't* know: We don't know that we will come upon an answer in a finite amount of time if we continue to apply the methods of discovery at our disposal. Following Shieh and Wright, we should understand the anti-realist as concluding that we do not have adequate justification that the sentence in question has a determinate truth value at all.

In the present debate over whether someone who endorses an epistemic conception of truth need also jettison classical logic, the concept of undecidability is commonly treated as though its definition, and even its implications, were straightforward. Given a rigorous definition of effective decidability, its definition may be straightforward enough. Even then, however, the implications for the debate are, I would suggest, far from clear. Things are much less clear when the disputants in the debate seem to understand undecidability as everything from strict unknowability to "evidence-transcendence" (whatever they mean by that) to lack of an effective decision procedure. What the various authors seem to be looking for is a reason, given an epistemic constraint on truth and the existence of sentences which are, in some sense of the term, undecidable, why classical logic must be rejected. Most often the rejection of classical logic comes in the form of a rejection of bivalence, though Wright's argument for revision is a notable exception, as he argues directly for the rejection of the law of excluded middle rather than the principle of bivalence which typically justifies that law.

Especially in the case of Tennant, Dummett's arguments for why we cannot accept a realist notion of truth for undecidables are not sufficiently distinguished from whatever Dummett's, or anyone's, argument is for why endorsing an anti-realist notion of truth should lead us to reject classical logic. An undecidability concept is certainly central to Dummett's argument for the rejection of realist truth conditions for undecidables. I've argued that the precise notion ought to be that of lacking an effective decision procedure rather than strict unknowability. Dummett does not single out the question of logical revision as separate from his inquiry into whether a realist theory of meaning is sustainable. Rather, he seems to see the two questions as being answerable at once. All the same, the questions can be addressed separately. As long as we are going to address them (more or less) separately, it is crucial that we are clear about how the notion of undecidability is functioning in each case. It is not obvious that the same concept is supposed to be at work in both cases.

As Salerno sets the stage for his own argument in part by criticizing Wright's, let us return to Wright's Basic Revisionary Argument. That argument says that the law of excluded middle, in conjunction with his version of epistemically constrained truth (EC_2) and his undecidability concept (NKD), gives us a contradiction. Abstracting away from the specific notions, we can understand the structure of his argument in a general way like this, where LEM is the law of excluded middle, UND is the relevant undecidability concept, and EC is an epistemic constraint on truth:

(LEM + UND + EC) → CONTRADICTION

The upshot is supposed to be that the semantic anti-realist, who of course must admit that truth is epistemically constrained and, according to Wright, must admit that there are undecidables on pain of incoherence, should reject the classical law. In an effort to regiment the various concepts

at play in this at times confused and convoluted debate, I will continue to use 'EC' and 'UND' to refer to the more general notions of, respectively, epistemically constrained truth and undecidability, regardless of how the specifics are fleshed out in particular cases.

Salerno interprets Wright, more or less correctly, as arguing that EC and LEM entail a certain *decidability*. This is because, with LEM in hand, we may create an instance of it such that $\Diamond P\varphi \vee \sim\Diamond P\varphi$. Take 'P' to mean "is mathematically verified" (Salerno 2000, p. 213) or, perhaps better, simply *proven*. The instance comes simply from substituting $\Diamond P\varphi$ for φ in LEM ($\varphi \vee \sim\varphi$). We may then, via EC and Wright's Negation Equivalence, derive

Dec_S: $\Diamond P\varphi \vee \Diamond P \sim\varphi$

The reasoning that leads to Dec_S is Wright's, detailed above on p. 22. So far Salerno's construal of Wright is consistent with what Wright says, except that, as we have seen, Wright uses notions such as *being feasibly knowable* and *being supported by evidence* as his decidability concepts rather than *being proven.* At the time of Wright's *Truth and Objectivity,* the text off of which Salerno is working, the best justification Wright offered for his undecidability concept was that it was *a priori* unknowable "that the scales of in principle available evidence must tilt, sooner or later, one way or the other, between each statement and its negation" (Wright, 1992, p. 42). In other words, according to Wright, we must admit, *a priori* even, that it is possible that we will never come to know the truth value of a given P. Our discussion of Wright involved his more recent attempt to justify an undecidability concept (NKD) with a "principle of agnosticism" (AG).

Salerno here interprets Wright as saying that the decidability which follows from LEM and EC (and the Negation Equivalence which Salerno chooses to ignore) is unacceptable since, "*for all we really know the decidability of the discourse could be false"* (Salerno 2000, p. 216,

emphasis in original). This interpretation is consistent with what Wright says in the later paper about his NKD, the undecidability concept which says that we don't know that we can feasibly know P or feasibly know not-P for every relevant P. Recall that Wright struggled to justify NKD with some kind of agnostic principle, the first candidate AG telling us that we couldn't be said to know something unless we were in a position to exclude the possibility that its negation were true. Salerno's discussion of Wright most likely influenced Wright's later thinking, as he does cite Salerno in the later paper. Salerno certainly seems to have been on to Wright's original candidate AG, as when he interprets the earlier work as saying that "upon perusal of our current set of things known, we find nothing that explicitly contradicts the negation of the decidability thesis" (Salerno 2000, p. 216).

Salerno frames Wright's argument in terms of what is known. On Salerno's interpretation, Wright argues that *knowledge of* decidability would be inconsistent with the epistemic possibility that the negation of decidability were true. Earlier, Salerno simply told us that the decidability thesis followed from LEM and EC. Knowledge of such decidability was never mentioned. In order to get knowledge into the picture, Salerno recasts the premises LEM and EC as *knowledge of* LEM and *knowledge of* EC, explaining in a footnote that knowledge of decidability follows if we assume, as he thinks we should, "the non-contentious epistemic principle that allows us to claim knowledge of the known consequences of known premises" (p. 225n12).

Let us now note that the negation of Salerno's Dec_S would be $\sim (\Diamond P\varphi \vee \Diamond P \sim\varphi)$. Recall that, for an intuitionistic anti-realist, if it is impossible to prove φ and it is impossible to prove ~φ, then it follows that $\sim(\varphi \vee \sim\varphi)$, which then gives us $(\varphi \wedge \sim \varphi)$: Contradiction. Similarly, Salerno's problem with his version of Wright's argument is that the epistemic constraint on truth

is both classically and, more importantly, intuitionistically, inconsistent with undecidability understood as the negation of Dec_S. Thus, Salerno concludes, the anti-realist cannot say that it is epistemically possible that Dec_S is false, for Dec_S is *strictly impossible* given the anti-realist constraint on truth.

If we return to Wright at this point, we notice that in his later paper Wright offers the very consideration just mentioned as a reason why his AG:

> **(AG):** P should be regarded as unknown just in case there is some possibility Q such that if it obtained, it would ensure not-P, and such that we are (warranted in thinking that we are) in no position to exclude Q.

could not be the correct kind of motivation for NKD, which the reader will recall is his version of the undecidability concept:

> **(NKD):** $\sim K \forall P(FeasK(P) \vee FeasK(\sim P))$

Wright is then at pains to produce a better motivation for NKD. In our current context, such a move counts as an attempt to overcome the difficulty highlighted by Salerno. Wright, in that later paper, is trying to find motivation for the claim that we do not know the decidability thesis, which we may as well formulate in Salerno's way:

> **Dec_S:** $\Diamond P\varphi \vee \Diamond P \sim\varphi$

He is seeking such motivation because he wants to generate a contradiction with the claim that we *do* know the decidability thesis, a claim which follows from the law of excluded middle:

> **LEM:** $\varphi \vee \sim\varphi$

and the epistemic constraint on truth:

> **EC:** $\varphi \rightarrow \Diamond P\varphi$

The result, of course, is supposed to be reason to reject LEM. It would appear that Wright has an answer to Salerno along the following lines: The anti-realist has a reason to deny that we know the decidability thesis relevant to Salerno's paper, Dec_S, but this reason is not that it is epistemically possible that the negation of that thesis is true. Wright's view is that we have reason neither to deny the decidability thesis nor to endorse it when we are presented with the sort of epistemic situation he has labeled Quandary—the intuitionistic view on the matter is, recall, "agnostic through and through." And Wright's new motivation for the claim that we do not know the decidability thesis is that we do *not* have reason to believe that evidence for either of the disjuncts may "feasibly be acquired." Thus Wright replaced his AG with AG+:

> **(AG+):** [a compound statement] is known only if there is an assurance that a suitably matching distribution of evidence for (or against) its (relevant) constituents may feasibly be acquired.

Salerno gives us his own argument for logical revision. Like Shieh before him and Wright after him, he realizes that it cannot be the denial of the decidability thesis which matters here. And like those authors, he hits upon the idea that it must be, rather, the denial that we *know* the decidability thesis to be true. Salerno calls the resulting version of the undecidability concept "Epistemic Modesty" (p. 219). It results simply from placing a knowledge operator and then a negation sign outside of the decidability thesis (I have dropped his universal quantifier because it is unnecessary):

> **EM:** $\sim K (\Diamond P\varphi \vee \Diamond P \sim\varphi)$

Salerno thinks that realist and anti-realist alike would accept EM since it "simply amounts to the humble recognition that we have not yet confirmed that each understood mathematical claim or its negation is humanly provable in the long run" (p. 219). One will recall that Wright sees his

own undecidability concept as quite modest and acceptable to all, though he does endeavor to find support for the idea in the form of his AG+.

Aside from EM, Salerno also puts knowledge operators in front of LEM and EC so that we get the following, where I again drop the unnecessary quantifiers, but use Salerno's nomenclature (p. 219):

Exclusively Classical Thesis: $K(\varphi \vee \sim\varphi)$

Dummettian Anti-Realism: $K(\varphi \rightarrow \Diamond P\varphi)$

Salerno's argument then proceeds much like Wright's. From the classical thesis and the anti-realist thesis it follows (given Salerno's epistemic entailment principle from before) that we know that either φ is provable or ~φ is provable: $K(\Diamond P\varphi \vee \Diamond P\sim\varphi)$ (Salerno, p. 219). This claim is in direct contradiction with EM. Given that the anti-realist is committed to the epistemic constraint, and that both sides should, according to Salerno, accept his undecidability concept (the modesty principle), what has to go is the classical thesis. Salerno concludes that the anti-realist is justified in asserting the negation of the classical thesis: $\sim K(\varphi \vee \sim\varphi)$ (p. 219). He points out that his argument does not fall prey to the difficulties found with Wright. The epistemic constraint and the modesty principle (what I have called the undecidability concept) do not lead to contradiction on intuitionistic logic, though they do on classical logic. And as mentioned before, adding a knowledge operator to the decidability thesis and negating that instead of the thesis itself avoids contradiction on an intuitionistic semantics. Salerno claims that he has avoided Wright's problems by making use of quantified propositional logic (p. 219). It is unclear why he thinks that. What seems more important is that he has introduced a knowledge operator.

Cogburn takes issue with Salerno's interpretation of Wright. We have already noted that Salerno's interpretation of Wright, and hence his criticism of the same, is somewhat out of date as Wright has at least attempted to update his argument in light of the sorts of concerns Salerno raises. All the same, Cogburn's criticisms are worth looking into. Where Salerno argues that Wright's argument fails because it depends on the anti-realist endorsing an epistemic possibility that is unavailable to him, that is, the negation of the decidability thesis Dec_S, Cogburn argues that this is reason to think that Salerno has gotten Wright's argument all wrong. We ourselves have been over the difficulties associated with accepting, even as an epistemic possibility, the negation of things such as Salerno's decidability thesis Dec_S, which would give us the by now infamous $\sim(\Diamond P\varphi \vee \Diamond P \sim\varphi)$ which amounts to a contradiction on the intuitionistic semantics. We've also been over Wright's more recent thoughts on the matter and seen that Wright has retreated to ground which is firmer, at least at first glance, regardless of what he may or may not have been thinking when he wrote *Truth and Objectivity* (1992).

The first problem Cogburn has with Salerno's interpretation of Wright is just that no intuitionist would think that the negation of a claim were an epistemic possibility just based on the conviction that the claim itself is unwarranted (Cogburn 2002, p. 235). Cogburn's complaint is somewhat odd given that what Salerno is supposed to be offering us is an argument for rejecting classical logic. All Salerno really needs is for it to be reasonable for a *classical logician* to think that, if we are not warranted in accepting a claim, then the negation of the claim must be an epistemic possibility. And that certainly would be reasonable for a classicist, given that a classicist will endorse LEM. If there were no other difficulty involved, which there is, Salerno's version of Wright's argument would work fine. The non-revisionary (classicist) anti-realist ought, in this case, to see the resultant contradiction and reject one of the assumptions that

led to it. Being unwilling to reject the epistemic constraint, LEM would be tossed. Then we would have a revised logic. Note that it is no problem that one of the assumptions that fueled the argument would now be rejected. That's just how *reductio ad absurdum* works. There is, of course, another difficulty involved, and this has to do with Salerno's objection that the negation of the decidability thesis is not an epistemic possibility for anyone who accepts the epistemic constraint, since known impossibilities cannot be epistemically possible. As we saw, Wright and Salerno end up solving the problem by negating not the decidability thesis but instead the claim that we know the decidability thesis to hold.

So, Cogburn's first objection to Salerno's interpretation of Wright misses its mark. His second objection involves the fact that, in order for the negation of Dec_S to result in a known contradiction for the *intuitionist*, and therefore not be an epistemic possibility, there must be an existential quantifier to the right of the negation sign. If there were only a universal quantifier, as Salerno has it, and as Cogburn points out Wright must have intended, there would be no contradiction on intuitionistic logic (Cogburn 2002, p. 236). Intuitionistically, we cannot derive an existential negative from a negative universal. All the negative universal claim tells us is that it is impossible to prove that all the elements of the domain of quantification have the feature in question. It would not follow that it is then possible to prove that at least one member of the domain does not have the feature. And hence there would be no contradiction.

Cogburn is certainly correct, but in the end this is no real difficulty. Again, the goal here is to *argue for* intuitionism. At this point in the dialectic, it doesn't matter if a given reductio fails on the assumption of intuitionistic logic. As before, if we have consistency given intuitionism and inconsistency given classicism, the real reductio is on classical logic (see p. above). Salerno himself makes too much of the idea that the reasoning involved in a decision to

reject classical logic must be intuitionistically sound. He is correct to say that the operative notion of undecidability must be acceptable to an anti-realist. But he does not need every inference on the route to revision to work for both the classicist and the intuitionist. To repeat: We may be led to revise our logic based solely on classical considerations.

The version of Wright's argument Cogburn prefers, and against which he provides his own objections, is based on Wright's earlier presentation of the Basic Revisionary Argument. Recall that Wright's earlier (1992) presentation makes use of the concept of *evidence being available for P*, while the later (2002) presentation makes use of *P being feasibly knowable.* Cogburn focuses on Wright's C (see above), which says, "either evidence is available for P or evidence is available for its negation." He points out, correctly as far as I am concerned, that Wright doesn't say much about why he thinks that C is problematic (Cogburn 2002, p. 235). He then quotes Wright as saying, essentially, that C is inconsistent with the idea that we have no good reason to think that we will eventually come to know either P or not-P (Cogburn 2002, p. 235).

As a matter of fact, Wright does, though not on the page Cogburn cites, give some reason to think C problematic. We saw as much earlier when we quoted Wright as saying that the issue was that the intuitionist's refusal to endorse bivalence was traceable to the belief that, "some statements may be such that no evidence bearing upon them is available either way," (citation, from above). And here we have yet another undecidability concept. In light of the issues with which we have been dealing, it would appear that the intuitionist cannot accept that no evidence either way is even possible, lest we get into contradiction again. But to say that our current state of information simply leaves it undecided as to whether P or not-P is uninteresting and provides no impetus for an anti-realist view of any sort.

At any rate, Cogburn promises to make C palatable to an anti-realist. He rightly notes that the goal here is to extend intuitionism from mathematics to the contingent realm. Thus he situates his argument in the context of warrants which are *defeasible.* Cogburn also adduces some distinctions of A.J. Ayer's between "in principle" and "in practice" verifiability as well as "weak and strong" verifiability (Cogburn 2002, p. 237). The in-principle/in-practice distinction is just the same as the one we saw in our previous chapter's discussion of strict finitism. We can verify something in practice just in case our actual capacities in this world allow us to verify it. We can verify something in principle, however, if some finite extension of our actual capacities would allow us to verify it, thus overcoming limitations of space and very large but finite stretches of time (Cogburn 2002, p. 237). The weak/strong distinction involves levels of certainty. Something is weakly verifiable if and only if it "can be determined to be more probable, other things being equal," while it is strongly verifiable if and only if it "can be determined to be true with certainty" (p. 238). Cogburn further argues for a modal distinction to account for the fact that if something is false in the actual world, then we can't come to know it since knowledge requires truth. The distinction is unnecessary, however, since what we ought to be concerned with is whether we can come to know what truth value the sentence in question has.

Let us return then to Wright's C, the claim that either evidence is available for P or evidence is available for its negation. Coburn thinks he has made it acceptable to an anti-realist because it would only be implausible on the assumption that the disjunction is understood exclusively rather than inclusively (p. 240), an assumption which he claims only makes sense if we regard the relevant verifiability predicate applied to P (in Wright's case, evidence being available for P) as a sufficient condition for the truth of P (Cogburn 2002, p. 240). He then

claims that regarding the verifiability predicate that way requires us to assume that the verifiability in question is strong, not merely weak (p. 240). What Cogburn seems to be after is the idea that in the contingent, non-mathematical realm, weak verification is the appropriate notion. And he goes on, in his criticism of Salerno, to say that, "there is nothing controversial about claiming to know that every sentence or its negation is weakly, possibly verifiable in principle" (p. 241). Thus Cogburn thinks that Salerno's Epistemic Modesty principle is in fact not plausible, interpreted appropriately for empirical, contingent claims with defeasible warrants (p. 241). In fact, Cogburn pushes his point to the conclusion that, in general, "the more reasonable it is to claim that it is currently known that Dummettian anti-realism is true [i.e. that truth is evidentially constrained], the less reasonable it is to claim that epistemic modesty is true" (p. 242). The details of Cogburn's insightful analysis aside, he appears to be worrying about something along the lines of what Rosenkranz was concerned about earlier: How can we be justified both in thinking that we know that if something is the case then we can know that it is and that we don't know what to say either way as to whether a given P is the case or not? We would seem to be so ignorant as to not know enough about our subject matter to know that the epistemic constraint held. Cogburn's worry is illusory for the same reasons it was during our discussion of Rosenkranz. The intuitionist is, as Wright continuously harps, agnostic all the way down. The epistemic constraint only says that if P is the case, then we can know that it is. But the epistemic constraint does not commit the intuitionistic anti-realist to saying that P either is the case or it isn't. The intuitionist is happy to say that we simply *do not know* whether P or not-P for a given P, which is consistent with the idea that if P is the case, then we can know that it is.

4.2 Tennant, Salerno, and Cogburn on Dummett on Logical Revision

The dialectic among Tennant, Salerno, and Cogurn in their attempts to pin down Dummett's argument for why the semantic anti-realist ought to reject classical logic is quite tortuous. I warn the reader now that my conclusion will be much simpler than the discussion that is about to ensue. For I ultimately conclude that none of either Tennant, Salerno, or Cogburn manage to come up with an argument for logical revision which is convincingly Dummett's. Dummett's real reason for thinking that classical logic must be revised seems to have been simply that the anti-realist semantics in terms of proof made it so.

Salerno gives what he takes to be Dummett's argument for logical revision. The argument is quite simple. It just says that the combination of LEM with an epistemic constraint on truth such that

$$\forall \varphi\ (\varphi \rightarrow \Diamond P\varphi)$$

and an undecidability concept in the form of

$$\exists \varphi\ (\sim\Diamond P\varphi\ \wedge \sim\Diamond P \sim\varphi)$$

leads to a contradiction which should force us to reject LEM. Salerno's issue with the argument as presented is that contradiction can be obtained, even intuitionistically, from the epistemic constraint and the undecidability concept all on their own. The rejection of LEM, it would appear, is a non-sequitur here, as LEM was not even required to produce contradiction (Salerno 2001, p. 214). My problem with Salerno's interpretation is that his rendering of the relevant undecidability concept is itself intuitionistically incoherent—it deals with strict unknowability. From an intuitionistic standpoint, such a concept is inconsistent with everything, since it entails contradiction all by itself.

Tennant also offers his own take on Dummett's reasons for thinking his anti-realist should reject classical logic. Tennant sees Dummett as running a reductio on the principle of bivalence. In a nutshell, Tennant's version of Tennant's argument is incredibly simple: Bivalence implies the existence of an undecidable sentence, but the epistemic constraint on truth implies that no sentence is undecidable. Contradiction. Thus we should reject bivalence (Tennant 1997, p. 180). My main criticism of Tennant is that he does not provide any good reason for thinking that Dummett took the assumption of bivalence to entail the existence of any undecidable sentences. Dummett, rather, took the existence of undecidables to be a fact agreed upon by both sides of the realism debate. In his exposition of Dummett, Tennant produces an entire paragraph of reasoning which he attributes to Dummett, but the paragraph is clearly a paraphrase at best, and Tennant provides no page number to account for where Dummett said such a thing (Tennant 1997, p. 180). He merely tells us that Dummett makes the argument in his "The Philosophical Basis of Intuitionistic Logic." In the relevant paragraph, Tennant says, "assume…for the sake of *reductio,* that bivalence holds…then there will be an undecidable sentence" (p. 180). As far as I can tell, this is Tennant's reasoning, not Dummett's. Tennant does quote Dummett's paper later on, but he adds a bracketed remark to indicate that Dummett is talking about bivalence and excludes relevant pieces of the actual passage (Tennant 1997, p. 182). I find no evidence in the passage itself to suggest that Dummett thought he was performing a reductio on bivalence. Tennant concludes that Dummett's attempted reductio of bivalence doesn't work because Dummett provides no real argument for why the existence of an undecidable sentence would follow from the assumption of bivalence (Tennant 1997, p. 194). I conclude that Tennant is simply wrong about what Dummett was trying to say.

Cogburn, for his part, objects to both Salerno's and Tennant's attempts to reconstruct Dummett's argument for logical revision. His main complaint is that Dummett could not possibly have meant to advance either of the arguments those authors attribute to him, for they are so obviously inadequate (Cogburn 2005). Cogburn gives us his own interpretation of Dummett's purported argument for logical revision. Notice that Salerno's rendering of undecidability involved an existential quantifier, as did Tennant's. Cogburn's does as well, though the difference he adds is that it is merely *possible* that there is a sentence such that we are unable to know it and we are unable to know its negation. And he takes Dummett's argument to be based on the idea (what he labels "Dummett's Insight") that the possibility of the existence of such a sentence is implied by our standard truth conditional, model-theoretic semantics (Cogburn 2005, p. 24). Notice that we are still dealing with that unfortunate conception of undecidability which is anathema to an intuitionist, the one that says an undecidable is a sentence such that we cannot know it and we cannot know its negation. Cogburn attributes to Dummett, in essence, the argument that the possibility of the existence of an undecidable sentence is inconsistent with the epistemic constraint on truth, since that constraint entails the knowability principle. Cogburn's Dummett rejects classical semantics to avoid contradiction (Cogburn 2005, p. 27). Cogburn praises his version of the argument for avoiding Salerno's and Tennant's criticisms of Dummett. He then adds,

> Unlike the arguments Salerno and Tennant discern, the main conclusion is not that we should eschew classical inference, but rather that we should eschew the use of classical model theory in the theory of meaning. It is only because of Dummett's further belief that classical inference needs the justification of classical model theory that his argument can be applied against classical inference. (p. 28).

Cogburn cites some passages of Dummett's which would certainly seem to support the idea that Dummett thought classical semantics inadequate. The main issue I have with Cogburn's interpretation is that he makes too much of what Dummett says about a limited number of cases. In particular, Cogburn takes Dummett's comments about the difficulty involved in quantifying over an infinite domain to explain Dummett's entire reason for rejecting classical truth-conditional semantics (Cogburn 2005, p. 24). In reality, I would suggest, in the relevant passages of Dummett's he was only concerned to illustrate the difficulty an intuitionistically-inclined theorist would have with the classical semantics. Furthermore, Cogburn himself falls into the trap of formulating the undecidability concept as a strict inability to know. Such a construal of the concept makes it inadequate for the purpose of explaining why the intuitionistic anti-realist would refuse to endorse the notion that an undecidable sentence has a determinate truth value—on the construal in question the anti-realist can't even admit that there are undecidable sentences at all, much less have a view about what to do with them.

I conclude that neither Salerno, nor Tennant, nor Cogburn gives a compelling treatment of the reasons why Dummett thought his anti-realist ought to abandon classical logic in favor of intuitionistic logic.

5. Some Conclusions

There has been much ado about the question, "Need the semantic anti-realist revise classical logic?" In my view, much of this ado has been about much less than it has commonly been thought. The semantic anti-realist, whose defining feature is that she endorses an epistemic constraint on truth such that all truths are knowable, certainly ought to revise her logic accordingly. Classical logic is typically taken for granted, and not unjustifiably so. It is the logic that, by all outward appearances, we appear to reason with in our everyday lives, even in those

moments when we sit down to do the sort of philosophy that leads us to question whether the logic that we take for granted is the correct one. So much was the gist of Quine's "Truth by Convention," in which he told us that we ought not take seriously Carnap's logical pluralism since classical logic was inescapable, it always being the case that our natural meta-language, so to speak, was bound to be classical.

Dummett is often understood as a global anti-realist, that is, one who thinks that semantic realism fails in every area of inquiry. Dummett certainly suggests as much. A more reasonable interpretation of Dummett's viewpoint, however, is that semantic realism must fail for any region of discourse where the sentences about the subject at hand are undecidable, where undecidability is understood in an appropriate way. Realism must fail in those areas, he thinks, for the reasons belabored in the previous chapter—given semantic realism, there would be no explaining the facts about our learning, expressing, and endowing of meaning based on our use of the language. Crispin Wright in particular has carved out space for local realisms and anti-realisms. While Dummett's arguments appear to have global force because he considers mainly facts about our language independent of subject matter, Wright has brought it to our attention that, perhaps, anti-realism is appropriate not just where the sentences about our subject matter are undecidable, but in particular when we think that an epistemic notion of truth is appropriate for undecidables. Dummett would likely agree.

The ado concerning the revision of classical logic has been about so little because, to begin with, there is very little space in which a non-revisionary anti-realist might operate. More importantly, however, the revision of classical logic follows quite quickly from the revision of classical semantics. It takes a certain state of denial to even begin to question why the semantic

anti-realist should want to revise her logic. Crispin Wright was, at one stage, full of such denial. During that period, he remarked,

> It is true that if somebody accepts Bivalence for a class of statements for whose truth-values we cannot in every case guarantee a means of decision, then he is at least committed to holding that we cannot guarantee that truth everywhere coincides with decidable truth. But, unless he accepts the transition from 'we cannot guarantee that P' to 'it is a possibility that not-P', he has not thereby committed himself to the possibility of verification-transcendent truth. This transition is intuitionistically suspect: of any mathematical statement which is not effectively decidable it would be intuitionistically correct to say that we cannot guarantee the existence of means of verifying or falsifying it, but it is not, in view of the intuitionists' account of negation, acceptable as an intuitionistic possibility that there should simply *be* no means of verifying or falsifying the statement in question. There is therefore a doubt whether endorsement of Bivalence for other than effectively decidable statements is of itself an admission of the possibility of verification-transcendent truth. (1993, p. 434)

The requirement of logical revision on behalf of the semantic anti-realist has seen two periods of intense questioning. The first has to do with the passage of Wright's just quoted, which is from a paper he wrote asking whether the Dummettian anti-realist is under any particular pressure to revise classical logic. The major doubt is captured in the passage itself. The second period occurs, perhaps unsurprisingly, just after Wright publishes his *Truth and Objectivity* (1992). In that work he no longer doubts that an anti-realist in Dummett's vein must revise her logic—rather, he offers an *argument in favor* of that view. We have seen the argument. Indeed, we have seen it under more than one interpretation, including two of

Wright's. In the wake of that work, Tennant, Salerno, and Cogburn began to ask what Dummett's argument for logical revision could be. It became a serious question, at least as far as they were concerned, whether the anti-realist was under any particular pressure to adopt an intuitionistic logic, or at least to reject the classical one. It would appear that Wright had made everyone doubt whether Dummett's anti-realist needed to endorse generalized intuitionism after all, first by questioning the move from semantic anti-realism to logical revisionism in the first place, and then by offering an explicit argument for why the semantic anti-realist is forced into revision.

For anyone used to Dummett's line of thinking, it will seem quite natural, once we have endorsed an epistemic constraint on truth, to replace the standard, realist, truth-conditional semantics with a semantics based on proof rather than truth, in the intuitionist way. And once we endorse that semantics, logical revision would appear to follow straightforwardly. What line, then, is Wright pursuing when he seeks another justification for the rejection of classical logic? On this matter he says in his later work,

> Persuasion that truth is essentially—or locally—evidentially constrained might thus lead to dissatisfaction with classical *semantics*—and hence with the classical justification for classical logic. But why should that enjoin dissatisfaction with the logic itself? There would seem to be an assumption at work that classical logic needs its classical justification. But maybe it might be justified in some other way. Or maybe it needs no semantical justification at all. (Wright 2002, p. 65).

Among the attempts to either discern in Dummett a specific argument for why his anti-realist ought to revise classical logic or to give an original argument for that same conclusion, some undecidability concept or another always plays a central role. In the arguments for logical

revision advanced by Wright and Salerno, as well as those advanced on Dummett's behalf by Salerno and Tennant, an undecidability concept is supposed to interact with some expression of a distinctly classical law, that of excluded middle, as well as an epistemic constraint on truth, to deliver a contradiction which is supposed to impel the anti-realist to reject classical logic. The argument Cogburn imputes to Dummett is somewhat different. The problem with it is that it only explains why Dummett would think logical revision necessary in a limited number of cases, those where we must quantify over an infinite domain. A major task of our current discussion has been to determine what the relevant undecidability concept must be and how it is best viewed as operating in the debate. Renditions of the concept according to which what is relevant is whether we *can* come to know a sentence's truth value run up against the difficulty that the semantic anti-realist cannot accept this form of undecidability as a real possibility since, along with the epistemic constraint on truth, we are easily found to be in contradiction. The typical solution to this problem, which is an appropriate one in the context of intuitionism, involves framing things in terms of what is known. This way, the undecidability concept gets cashed out as the claim that we *do not know* that it is either possible to know φ or possible to know ~φ. We saw Wright, Shieh, and Salerno take this sort of tack.

What, then, of effective procedures? Lack of an effective decision procedure is one form of an undecidability concept. That we don't know that either P is knowable or not-P is knowable is another one. Wright, more than anyone else, has attempted to flesh out the reasons why we might want to claim that we don't know that disjunction. Interestingly, nearly all of the authors discussed in the present work appear to think that the second concept, the one about what we don't know, is what importantly follows from the first concept, the one about effective procedures. We should question that sort of thought. To be fair, the two concepts are

compatible with one another. It might be the case both that we lack an effective procedure for deciding some P while at the same time we do not know ($\Diamond K\varphi \vee \Diamond K \sim\varphi$). The classicist is comfortable admitting that some things are utterly unknowable. Thus the classicist can accept what the intuitionist cannot: $\sim (\Diamond K\varphi \vee \Diamond K \sim\varphi)$. So the classicist can know what the intuitionist cannot, though they can both also fail to know.

At the same time, lack of an effective decision procedure does not *entail* that we do not know ($\Diamond K\varphi \vee \Diamond K \sim\varphi$). For we might *know*, for instance, the left disjunct which says that it is possible to know φ, and yet still not have a procedure for coming to know φ which is *effective* in the sense of being guaranteed to deliver an answer in a finite amount of time. To admit as much is simply to admit that we *could* stumble upon knowledge of φ, perhaps through some non-effective method.

We would clearly require of a *good* decision procedure that it be *reliable,* where a *reliable procedure* is one that is, at the very least, more likely to deliver a correct result than an incorrect one. But must a reliable procedure be effective in the sense of delivering a result in a finite amount of time? How reliable would a machine be that was guaranteed to deliver a correct result whenever it did deliver a result, but which was not guaranteed to ever deliver a result? This would surely be a misapplication of our usual notion of reliability—if it were not we could call a car reliable which was guaranteed to deliver you to your destination on the condition that it started even though we could not guarantee that it would ever start. Still, I take it that reliable decision procedures are those which make it probable that a correct result will be delivered when one is delivered at all, while effective procedures are those for which there is a guarantee that repeated application of the procedure will deliver a result in finite time.

Must effective procedures be reliable? Here is one problem with the extension of the intuitionistic model to the contingent realm where, as Cogburn emphasized, our warrants are defeasible. In such a realm we no longer talk of proof but of warranted assertibility. We do not require *certainty,* merely *decision.* At some point we may become *warranted* in asserting/believing something though we might still be wrong about it. We can think of a decision procedure this way in contingent domains, or at least domains where we take justification to be defeasible. In logic and mathematics, we are supposed to be working with the idea that as long as we are in fact in possession of a proof, then that which is proved must be true. It is different with justified assertion/belief: We might be justified in asserting or believing while nevertheless we are wrong; discovery that we are wrong does not make it the case that we were not justified to begin with the way that discovery that the conclusion of a proof is not true forces us to conclude that the purported proof was never a proof at all.

Once we have entered the empirical realm, it is not entirely clear what significance a guarantee to deliver a decision in a finite amount of time really has. One fact stands out: The possession of an effective decision procedure for P no longer seems to imply without reservation that it is possible to know P. All the same, the possession of the procedure and our faithful use of it would seem to rule out, even for a classicist, the *impossibility* of coming to know the truth value of P. I want to suggest that any justification that anyone thought was to be had for the undecidability concept that we do not know ($\Diamond K\varphi \vee \Diamond K \sim\varphi$) is not forthcoming from considerations pertaining to effective decidability. The undecidability concept regarding what is or isn't possible to know, rather, seems to be justified just in virtue of the work that it does in the overall theoretical apparatus. And as we saw with Wright, it would appear that justification for the intuitionistic outlook itself ultimately must come from consideration of the nature of the

subject matter at hand, i.e., its metaphysics. In any case, for all the effort to devise an argument for why the anti-realist ought to reject classical logic, everything seems to turn on what justification we might have for thinking that truth is epistemically constrained to begin with.

Let us return now to the two question threads we set out to pursue at the beginning of this chapter:

> **thread one:** What is the relationship between evidence-transcendent truth and the application of (determinate) bivalence to sentences lacking an effective decision procedure?
>
> **thread two:** Must one who rejects evidence-transcendent truth, and endorses instead evidentially constrained truth, also reject (determinate) bivalence and with it classical logic?

Focus on the second question first. A Dummettian, and indeed a mathematical intuitionist, would say that we are perfectly justified in accepting that decidable sentences have determinate truth values, even bivalent ones. For decidables, we can accept classical reasoning. There is no trouble, from a Dummettian standpoint, with believing that decidable sentences obey the law of excluded middle and the rule of double negation elimination. When it comes to undecidable sentences, however, their undecidability is supposed to make it the case that we cannot accept that they have determinate truth values, certainly not bivalent ones, and additionally that we should not think that classical logic governs them. The reasons for this were essentially laid out above. Once we accept a semantics based on provable truth rather than evidence-transcendent truth, we lose justification for thinking that the law of excluded middle which says that either P is the case or ~P is the case holds. For we have no good reason to think that either P will be provable or its negation will. Similarly, we lose justification for the rule of double negation

elimination. For in that case we have no good reason to think that a proof that it is impossible to prove that it is impossible to prove P entails that it is possible to prove P itself. So, I would say, the answer to the second question is yes.

Answering the first question is more difficult. Dummett speaks often of bivalence because of its association with the classical logic which the intuitionist rejects. But what matters to the notion of evidence-transcendent truth is not bivalence so much as it is determinate truth. The truth value of a sentence is evidence-transcendent if the sentence can have a determinate truth value independently of our ability to come to know what that value is. Now imagine that a sentence for which we lack an effective decision procedure has a determinate truth value. Does that sentence possess evidence-transcendent truth? Not necessarily, at least not in the sense of evidence-transcendent truth we have been working with, and which I have claimed is operative in all sorts of discussions surrounding Dummett's work, not to mention Dummett's work itself. For the notion of evidence-transcendent truth has to do with what we are capable of knowing strictly speaking, while the truth value of a sentence which we cannot effectively decide merely transcends our ability to effectively decide what the value is. If it is a sentence whose truth value we are capable of coming to know even though we possess no effective procedure for coming to know it, then it is a sentence to which epistemically constrained truth may apply despite the fact that it is not effectively decidable. In that sense, an undecidable might possess a determinate truth value despite the fact that the truth value in question is epistemically constrained. Hence, in that sense, if epistemically constrained truth is what matters for anti-realism, figuring out whether undecidables could possess determinate truth values is not a good way to determine whether we should be anti-realist about the subject matter of those sentences.

If undecidability is simply unknowability, then of course epistemically constrained truth cannot apply to an undecidable. We have emphasized the fact that both realists and anti-realists accept the existence of undecidables while the anti-realist cannot accept that there are unknowables strictly speaking. So whatever notion of undecidability matters, it is not strict unknowability. Now we have seen that if the proper notion of undecidability is lack of an effective decision procedure, it would appear that we lack the resources to establish whether truth is evidence-transcendent or evidentially constrained. If undecidability were in fact unknowability, then the possession of a determinate truth value by an undecidable would be sufficient for establishing that the truth of the sentence is evidence-transcendent. This seems to be what Wright has in mind in the passage quoted above, which I reproduce here in part:

> It is true that if somebody accepts Bivalence for a class of statements for whose truth-values we cannot in every case guarantee a means of decision, then he is at least committed to holding that we cannot guarantee that truth everywhere coincides with decidable truth. But, unless he accepts the transition from 'we cannot guarantee that P' to 'it is a possibility that not-P', he has not thereby committed himself to the possibility of verification-transcendent truth.

As we saw, Wright takes issue with the inference he mentions, saying that an intuitionist need not accept it as valid. His main concern is with the question as to whether the anti-realist must reject classical logic. So he is interested in showing that the application of bivalence to undecidable sentences doesn't require that one endorse the realist's notion of evidence-transcendent truth. I am making a similar point, this time in order to raise doubt about the effectiveness of Dummett's reasoning.

To get our heads around the relevance of undecidability to Dummett's reasoning, it will help to express Dummett's thought conceptually, in broad strokes. Consider the following propositions:

1. There are *undecidables*.

2. If a realist theory of meaning based on evidence-transcendent truth conditions is correct, *undecidables* are such that we cannot learn what they mean, express what they mean, or even give them the meanings they are supposed to have.

3. The *undecidables* are also such that our ability to come to know whether they are true or false is limited in the way specified by the concept of *undecidability*.

I have used asterisks to indicate that some concept of undecidability is in play without being specific about which concept is at work. So let us now ask which undecidability concept *must* be at work in each case. Here are the different undecidability concepts we have uncovered from the literature on Dummettian anti-realism, where P is an undecidable:

A. We cannot know P and we cannot know ~P.

B. We do not know that we can know P and we do not know that we can know ~P.

C. We do not have an effective procedure for determining whether P is true or ~P is true.

Wright's Quandary concept is presented as an undecidability notion, but ultimately Wright is arguing that something like B is an appropriate conclusion when we find ourselves in a state of Quandary. Furthermore, it is his rendering of B that he takes to do the work in his argument for logical revision. Indeed, we might even say that Wright takes B to be the important result that follows from any relevant notion of undecidability in such a way that the intuitionistic idea generalizes from mathematics to other areas.

It often seems to be taken for granted in the literature that B follows from C. We have seen that C in fact does not entail B. We might know that we can come to know P while also knowing that we do not have an effective procedure for coming to know P. And yet Wright argues both that being in a Quandary is what justifies B and that the mathematical intuitionists were in a Quandary with regard to certain statements simply in virtue of lacking an effective decision procedure for them. If C justifies Quandary, and Quandary justifies B, then C should justify B. But it does not, as we have seen based simply upon consideration of the possibilities left open by the lack of an effective procedure. Any argument for the extension of intuitionism to areas other than mathematics which relies upon B following from C just cannot work.

We would do well to ask why Wright thought that we must be in a Quandary regarding sentences satisfying C. Perhaps it is supposed to be the case in mathematics but not in other areas that lack of an effective decision procedure puts us in a state of Quandary. The modal considerations we brought to bear on the issue were quite general, and should apply to mathematics as much as to other areas. One might think that the intuitionist's epistemology of constructive proof puts them in a Quandary with regard to sentences which are not effectively decidable. But this does not seem to help either. We can replace the concept of knowability with the concept of constructive provability and the results with regard to the difference between B and C will be the same.

Returning now to our list of undecidability concepts, note that A is intuitionistically unacceptable, and is therefore unacceptable to the semantic anti-realist on the assumption that an intuitionistic logic follows from the adoption of an epistemic constraint on truth. We will assume that it does. So A cannot be the notion of undecidability at work in 1, since both the realist and anti-realist accept 1.

Let's turn to 3. Any of A-C could be at work in 3 for a realist. The anti-realist, being unable to accept A, is limited to one of either B or C. Either of those notions could be in play in 3, as long as the same notion is operative in both occurrences of the undecidability concept. The point I would like to bring out at this point is that Dummett seems to have thought that, whatever notion of undecidability he had in mind, the same one was operative across 1-3. There is reason to doubt that the same concept is actually relevant to each of 1-3. Consider 2. What is it about the undecidables which is supposed to make it the case that 2 is true? This has never been entirely clear. Let's plug in the undecidability concepts A-C and see what happens.

Given Dummett's belief that meaning is use, along with his thought that use is essentially tied to our epistemic capacities, A, which tells us that neither P nor ~P is knowable, turns out to be a rather convincing candidate for the operative undecidability notion in 2. For it would make some sense, if undecidables were just unknowable, and the conveying of meaning, for instance, involved the display of a capacity to know, that the meaning of an undecidable could not be conveyed because there would be no capacity to know. This often seems to be what Dummett has in mind. But as has been continually stressed, the undecidables cannot just be unknowables because the intuitionist would then not be able to accept that there were such sentences.

Turn now to B. Could B, which says simply that we do not know that we can know P and we do not know that we can know ~P be the operative notion in 2? Keep in mind that B is one of the major ideas advanced by those writing in Dummett's wake as a solution to the kinds of issues we find with A. These thinkers tell us that undecidables are not such that we cannot know them. Instead they are such that we do not know that we can know them. Hence we get intuitionistic agnosticism. The idea of B is supposed to be general enough that it could apply to areas other than mathematics, and indeed B would appear to motivate intuitionistic reasoning if it

does not stem from intuitionistic reasoning. The trouble is that B can't be operative in 2. Where it turned out that A was too strong, B is not strong enough. The fact that we currently don't know whether P or ~P just can't motivate Dummett's thought that, given realist truth conditions, the meanings of undecidables are unconveyable, unlearnable, and in fact not the sorts of meanings we could have given to those sentences in the first place. Consider Dummett's idea that what we manifest in use is a capacity to come to know. For all that has been said, we might still have the ability to know P while at the same time not knowing that we can know P.

What about C? As with B, C is not strong enough to motivate 2. This is because we might lack an effective procedure for coming to know the truth value of a sentence while still possessing the capacity to know the truth value of the sentence. As mentioned, it is often thought that C motivates B, but this is simply not the case. And it would be senseless to insist that what we manifest in use is the possession of an effective decision procedure for a sentence. Most of our sentences lack effective decision procedures.

Dummett wants us to think that if the undecidables have determinate, evidence-transcendent truth values, then we cannot explain how the undecidables could possibly have the meanings they are supposed to have. It is difficult to see what reason there could be to believe that the conditional just stated is true. Whatever reason Dummett thinks there is, he seems to think that it has everything to do with the nature of the concept of undecidability. But that nature has proved incredibly difficult to pin down in a way that supports Dummett's reasoning. We certainly *have* various notions of undecidability. But we do not have one single notion which can play all of the roles Dummett needs it to play. I would suggest that Dummett was seduced into thinking that he had discovered a deep connection between the nature of truth, meaning, and the concept of undecidability largely because he found the concept of undecidability in all of 1-3.

The major difficulty is that once we attempt to analyze the concept and determine exactly how it functions, there seems to be no one concept that can do all of the things Dummett needs it to do.

So, the answer to the first thread is that the notions of evidence-transcendence and undecidability are not related in the ways that Dummett seems to have thought they were. At any rate, they are not related in the way he needs them to be in order for his argument to work. That said, the idea that we do not know that we can know P and we do not know that we can know ~P functions well in a generalized intuitionism about anything at all. It avoids intuitionistic contradiction, unlike the stronger thought that we cannot know P and we cannot know ~P, and also captures the sort of agnostic attitude behind intuitionism. So it may very well be that a generalized intuitionism is viable even though Dummett's argument for such an intuitionism does not work. Dummett attempted to establish that the intuitionistic philosophy could be generalized outside of mathematics via considerations pertaining to the theory of meaning. I have argued that while many of the extant criticisms of Dummett fail, Dummett was never in possession of an undecidability concept that could do the work he needed it to do. In addition, we have learned from our study of Wright's work, which represents the strongest attempt after Dummett to pursue Dummett's goal of generalizing intuitionism, that perhaps the best case that could be made for an intuitionistic attitude toward some subject matter proceeds not from an analysis of language but from a decision about the nature of that subject matter itself. I do not think this shows that the philosophy of language is impotent to motivate a stance on the realism issue. Dummett's argument might not work, but his way of conceiving of the issues was completely legitimate. At the same time, metaphysical considerations might tilt the scales towards either realism or anti-realism. As was suggested in the first chapter, we do not need to

put any particular area of inquiry first. We should allow considerations in the philosophy of language and metaphysics alike to lead us to an answer to the question of realism.

Chapter 5: Conclusion

The Kantian perspective on the relationship between our minds and the world, if not Kant's theory itself, has been immensely popular among anti-realists. This should come as no surprise, as the Kantian perspective in its most general form amounts to a stance on how mind-dependent the world that we know is, and it's supposed to be approximately half mind-independent and half mind-dependent. It was mentioned earlier that Putnam likened his own version of anti-realism to Kant's. I would like to suggest that any anti-realism which takes as its central tenant the claim that truth is evidentially constrained can be understood as a version of the Kantian view. These anti-realisms have it that the so-called "external" world that we seek to investigate is necessarily shaped in part by our capacity to come to know about it. Such anti-realisms reject the idea that we learn absolute truths which could obtain even if we could not, at least in principle, come to know them. Thus what we come to know about are not things as-they-are-in-themselves, but instead things as-we-are-capable of knowing about them. Our epistemic capacities stain the world we seek to learn about.

In the first chapter I attempted to clear the ground for the discussion that was to follow. Devitt would like us to think that the realism issue should not be approached as an issue in the philosophy of language. I argued that his reasons are unconvincing. I hope I have also demonstrated through the example of Dummettian anti-realism that there are perfectly good realism questions about the truth conditions of our sentences. In fact, I have attempted to show that the realism issue can be fruitfully approached either as an issue in metaphysics or language, and furthermore that regardless of how one conceives of the major issue, considerations in both areas can spill over into each other.

The concept of mind-independence looms large in discussions of realism and anti-realism. Against those who would suggest that the concept of mind-independence is not a workable notion, I have offered my own definition of mind-independence in terms of independence from our directed cognitive practices. My definition avoids some potential concerns about how to avoid a commitment to metaphysical naturalism and about how we might be realists about minds themselves while still understanding realism in any area as the view that the facts which obtain in that area obtain mind-independently. Those who prefer to construe the realism issue as a metaphysical issue are under the most pressure to make sense of the notion of mind-independence, but mind-independence is also at play when realism and anti-realism are understood as semantical positions: The realist who believes that our sentences have evidence-transcendent truth conditions thinks that the truth of our sentences is completely mind-independent. The anti-realist who thinks that the truth of our sentences is somehow epistemically constrained believes that the truth of our sentences depends on our minds to the extent that what is true depends on what we are capable of coming to know. Again, we have here the Kantian idea that there is a world out there but what we come to know about it is actually shaped by our ability to come to know about it.

I have endeavored to defend Dummett against attacks which, I have argued, fail. I myself have offered some reason to think that Dummett's arguments for semantic anti-realism do not work. However, I must emphasize that I believe Dummett's overall way of setting up the realism issue as one over the truth conditions of our sentences is perfectly legitimate. Recall that Dummett's anti-realism involves the Fregean idea that to understand the meaning of a sentence is to know the condition for its truth, the Wittgensteinian idea that meaning is use, and the intuitionistic idea that what is true is what is provable. Dummett's anti-realist argument

proceeds under the assumption that to understand a sentence, or to grasp its sense, is to know how to use it. It is because of the connection Dummett perceives between use and epistemic capacities that he finds fault with a realist semantics according to which our sentences have evidence-transcendent truth conditions. Thus, in order for one to be swayed by Dummett's argument for anti-realism, one must endorse a use-theory of meaning like Dummett's. Use-theories are popular in some circles, but they are by no means the standard. Certainly a realist is under no special pressure to endorse a use-theory, and hence Dummett's argument will be ineffectual even if it is not flawed. Dummett has sought to duel the realist, but the realist has gone to the café and Dummett is at the saloon.

Dummett's overall approach to the realism issue as one concerning the kinds of truth conditions our sentences have does not require the use-theory of language that Dummett endorses. So the inefficaciousness of his argument should not lead us to abandon his framing of the realism issue. For example, Putnam also sees the realism issue as, largely, the question as to whether truth is evidence-transcendent or evidentially constrained (Putnam 1978, 1981, 1983). Yet anti-realist Putnam of the late 1970s and the 1980s does not endorse the Dummettian version of a use-theory, and Putnam's model-theoretic arguments for his own version of anti-realism are of quite a different character than Dummett's anti-realist arguments. Indeed, I would suggest, the basic idea that what matters to realism is whether truth is evidentially constrained does not require any particular understanding of how language works, though certain theories of language will likely lead more easily to an anti-realist conclusion than others. As theorists interested in realism issues we are at liberty to choose any linguistic theory we like, as long as something along the lines of truth conditions are involved, and then ask whether our sentences have evidence-transcendent truth conditions.

I have emphasized that once we endorse the view that our sentences have evidentially constrained truth conditions, or if one prefers, "proof conditions," a re-writing of the logic that we use to reason about the area concerning which we are anti-realist must ensue. Thus, it would seem, if anti-realism is understood as the view that truth is evidentially constrained, something alone the lines of intuitionistic logic will come along with anti-realism. Chapter 4 in particular was concerned with nailing down the details of the intuitionistic outlook. A major component of Dummett's grand vision is that the intuitionist or constructivist brand of anti-realism in the philosophy of mathematics can be extended to the contingent, empirical realm. Dummett proceeds by first making the intuitionist view a view about how we give meaning to sentences, convey the meanings of our sentences, and learn the meanings of our sentences in the first place. Then, once he has abstracted away from the subject matter of mathematics and made the issue one about meaning in general, Dummett applies similar reasoning to other areas outside of mathematics. I have already discussed the problem a Dummettian faces with attempting to turn realists who do not endorse a use-theory of meaning. The major issue with Dummett's overall project I have attempted to bring out is that there does not seem to be a concept of undecidability which can do all of the things Dummett needs it to do outside of the mathematical realm.

Turn again to the three undecidability concepts I have isolated from the literature surrounding Dummett's work ('P' in each case being an undecidable):

A. We cannot know P and we cannot know ~P.

B. We do not know that we can know P and we do not know that we can know ~P.

C. We do not have an effective procedure for determining whether P is true or ~P is true.

To review quickly, the problem with A is that the anti-realist cannot think that there are any sentences which satisfy it, and yet the Dummettian realism debate is supposed to be over whether

the undecidables have determinate truth conditions. One cannot be in a dispute about whether a certain kind of sentence has a determinate truth value if one does not believe that there are any sentences of the kind in question. So A cannot be the undecidability concept at work in Dummett's work. The problem with B is that, while it is perfectly acceptable to an anti-realist, it renders Dummett's arguments completely unconvincing. Dummett's overall argument has it that, if undecidables have evidence-transcendent truth conditions, then there is no way that we could endow those undecidable sentences with the meanings they are supposed to have, convey those meanings to others, and learn what those meanings are in the first place. But if an undecidable is just a sentence such that we do not know that we can know it and we do not know that we can know its negation, we might still be perfectly capable of coming to know either P or ~P. And we may, through our behavior, demonstrate that we had such a capacity while at the same time not realizing, or even entertaining the idea, that we could come to know whether P or ~P. B captures well the intuitionistic attitude, but it is too flimsy to function in Dummett's actual arguments for intuitionistic anti-realism. Similar problems plague C. We might have the epistemic capacity to come to know P while lacking an effective decision procedure for it. In addition, the general anti-realist intuition is typically motivated by the thought that we cannot make sense of something being the case unless we are at least in principle capable of coming to know that it is the case. The idea that we must be capable of coming to know that it is the case in a finite amount of time, and in fact that we must actually possess a procedure that will guarantee such a result, becomes something of a red herring outside of the mathematical constructivism where it was first relevant.

That said, I do not mean to imply that no good can come of conceiving of the realism issue as Dummett conceives it, as an issue about truth conditions. And furthermore it also does

not mean that the intuitionistic outlook does not do well outside of mathematics. My point is just that Dummett cannot think of undecidability the way that he needs to while expecting his arguments to be at all convincing. It may just be that an anti-realist outside of mathematics can perfectly coherently endorse an intuitionistic attitude towards a given subject matter, but that the reasons for doing so will not be Dummett's. This is where another one of the conclusions of my dissertation comes in: One might elect to be an intuitionist, who thinks that intuitionistic logic governs some area of inquiry, based largely on metaphysical considerations pertaining to what one thinks the stuff in that area is like. If one is inclined to think that what we learn about is not the stuff-as-it-is-in-itself but rather just the stuff-as-we-can-know-it, then one would have a good reason to think that the truth conditions of our sentences in that area are epistemically constrained, and that intuitionistic logic was appropriate for reasoning about the stuff that we are investigating. Dummettian preoccupation with use-theories of meaning and associated epistemic capacities never enter the picture. Thus, in such a case, we will have construed the realism debate as one over the truth conditions our sentences have, but metaphysical considerations will have pushed us into the semantical view.

At the same time, I take Dummett's approach to the realism issue to provide the right context for understanding how one could think that the realism issue is at bottom a matter of metaphysics while thinking that the answer to the metaphysical question might be pursued via semantics. For Dummett argues, based on linguistic considerations rooted in his use-theory of meaning, that our sentences (in particular the undecidables) have epistemically constrained truth conditions. Thus, he says, we should take an intuitionistic attitude toward the subject matter of those sentences and view our reasoning in that area as ruled by intuitionistic logic. If I am right that, regardless of whether one construes the realism debate as metaphysical or semantical, what

is at issue is mind-independence, then it would appear that we should conclude that the objects which our sentences are about have some kind of mind-dependence inasmuch as the truth of sentences about them is constrained by our ability to come to know whether those sentences are true. The Kantian perspective can help here again. The claim is not that the noumena themselves are mind-dependent, only that the objects insofar as we can come to know them have part of their nature lent to them by our own epistemological natures. The noumena remain out of our mental reach. The phenomena are what we can know about. So, I conclude, we ought not take such a hard-line approach as Devitt does when he tells us not only that realism is a metaphysical issue, not a semantic issue, but also that the way that the realism debate ought to be settled is through a consideration of metaphysics, not semantics. My view is that we can understand the realism debate either as an issue in metaphysics or an issue in semantics, and no matter which one we choose it is perfectly reasonable to allow both metaphysical and semantical considerations to lead us to a conclusion as to whether we ought to be realists or anti-realists.

I claimed in the first chapter that we are under no pressure to put any area of inquiry before any other. What I have in mind is something along the line of what I just sketched—at least when it comes to metaphysics and semantics in the realism debate, we can put either one of them first and work our way over to the other. Other theoretical commitments, no doubt, will entice particular theorists to start in particular places, but I urge that Devitt-style constraints on theorizing are ultimately unnecessarily suffocating. And there is room on my view for some minimal amount of theorizing in one area to occur first while that area does not dominate the overall inquiry. Devitt would seem to have it that we should settle the major metaphysical questions, such as, on his view, the realism question, before we move on to do work in other areas. Perhaps we could give it to him that a little metaphysics should be done first, insofar as,

perhaps, we need to agree that there are noumena out there. But the really dirty work of determining what the objects of our knowledge are like could then proceed via considerations of what sorts of truth conditions our sentences have. Ultimately we would be reaching metaphysical conclusions about the objects of our knowledge through linguistic considerations pertaining to what we must be talking about.

I have attempted throughout this work to remain largely neutral on whether one ought to be a realist or an anti-realist. Part of the reason for this is that my work has been of such wide scope that I have not dipped down into any particular region of discourse to register an opinion as to whether a realist attitude is appropriate there. I do think that one can be realist in one area and anti-realist in another, allowing for degrees of realism and anti-realism according to the degree of mind-independence in each area. My own view is that we necessarily understand the world through concepts of our own making, and that in general what learn about are not things-as-they-are-in-themselves but rather the things-as-we-can-know-them. It has not been my goal to articulate and defend such a view. That is a project for the future. What I hope to have done is to have cleared much of the landscape for such a project and begun to lay the groundwork. In the context of the realism debate, if not elsewhere, our metaphysics and semantics should be allowed to speak to one another. Dummett and his progeny have given us a model for how this might occur. Dummett's dream was that mathematical intuitionism could be extended to other areas. I have argued that the undecidability concept at the center of Dummett's anti-realist arguments was never carefully articulated and that in fact there is no single concept on the market that can do all of Dummett's bidding. This does not mean that there is a problem with mathematical intuitionism, or even a generalized intuitionism. It only means that Dummett's argument for a generalized intuitionistic anti-realism doesn't work. All the same, the

intuitionistic attitude, especially as developed by Crispin Wright, remains a respectable way to be an anti-realist about a great many things. The reasoning that might lead one to such an anti-realism rooted in semantics might nonetheless be metaphysical in character.

Bibliography

Alston, W. (1979). Yes, Virginia, There is a Real World. *Proceedings and Addresses of the American Philosophical Association, 52*(6), 779 – 808.

Alston, W. (1996). *A Realist Conception of Truth.* Ithaca: Cornell University Press.

Appiah, A. (1991). Representations and Realism. *Philosophical Studies, 61*(1), 65 – 74.

Asay, J. (2012). A Truthmaking Account of Realism and Anti-Realism. *Pacific Philosophical Quarterly, 93*, 373 – 394.

Byrne, D. (2004). Gardiner on Anti-Realism: A Defence of Dummett. *Dialogue: Canadian Philosophical Review, 43*(1), 3-23.

Cogburn, J. (2002). Logical Revision Re-Visited: On the Wright/Salerno Case for Intuitionism. *Philosophical Studies, 110,* 231 – 248.

Cogburn, J. (2005). The Logic of Logical Revision Formalizing Dummett's Argument. *Australasian Journal of Philosophy, 83(1),* 15-32.

Craig, E. (1982). Meaning, Use and Privacy. *Mind, New Series, 91*(364)*,* 541-564.

Devitt, M. (1991). Realism without Representation: A Response to Appiah. *Philosophical Studies, 61*(1), 75 – 77.

Devitt, M. (1997). *Realism and Truth.* (2nd ed.). Cambridge: Harvard University Press.

Devitt, M. (1998). Putting Metaphysics First: A Response to James Tomberlin. *Philosophical Perspectives 12,* 499 – 502.

Devitt, M. (2006). Scientific Realism. In P. Greenough and M. Lynch, (Eds.), *Truth and Realism* (100 – 124). New York: Oxford University Press.

Devitt, M. (2010). *Putting Metaphysics First.* New York: Oxford University Press.

Devitt, M. (2012). Hilary and Me: Tracking Down Putnam on the Realism Issue. In M. Baghramian (Ed.), *Reading Putnam* (101-120). Hoboken: Taylor and Francis.

Dummett, M. (1975). Wang's Paradox. *Synthese, 30,* 301-324.

Dummett, M. (1977). *Elements of Intuitionism.* Oxford: Oxford University Press.

Dummett, M. (1978). *Truth and Other Enigmas.* Cambridge: Harvard University Press.

Dummett, M. (1981). *Frege: Philosophy of Language* (2nd ed.). Cambridge: Harvard University Press.

Dummett, M. (1982). Realism. *Synthese, 52,* 55 – 112.

Dummet, M. (1983). The Philosophical Basis of Intuitionistic Logic. In P. Benacerraf and H. Putnam (Eds.), *The Philosophy of Mathematics: Selected Readings* (2nd ed.). Cambridge: Cambridge University Press.

Dummet, M. (1991). *The Logical Basis of Metaphysics.* Cambridge: Harvard University Press.

Dummett, M. (1993). *The Seas of Language*. Oxford: Oxford University Press.

Fine, K. (1994). Essence and Modality. *Philosophical Perspectives, 8,* 1 – 16.

Fine, K. (2001). The Question of Realism. *Philosophers' Imprint, 1*(1), 1 – 30.

Gardiner, M. Q. (2000). *Semantic Challenges to Realism: Dummett and Putnam*. Toronto: University of Toronto Press.

George, A. (1984). On Devitt on Dummett. *The Journal of Philosophy, 81*(9), 516 – 527.

Haack, S. (1974). *Deviant Logic: Some Philosophical Issues.* Cambridge: Cambridge University Press.

Horwich, P. (1990). *Truth.* Cambridge: Basil Blackwell.

Horwich, P. (2005). *Reflections on Meaning.* New York: Oxford University Press.

Hunter, G. (1973) *Metalogic: An Introduction to the Meta-Theory of Standard First-Order Logic.* Oakland: University of California Press.

Jenkins, C. (2005). Realism and Independence. *American Philosophical Quarterly, 42*(3), 199 – 209.

Kapsner, A. (2015). *Logics and Falsifications.* Dordrecht: Springer.

Khlentzos, D. (2004). *Naturalistic Realism and the Antirealist Challenge*. Cambridge: Massachusetts Institute of Technology Press.

Kirkham, R. (1989). What Dummett Says About Truth and Linguistic Competence. *Mind, 98*(390), 201-224.

Loar, B. (1987). Truth Beyond All Verification. In B. Taylor (Ed.), *Michael Dummett: Contributions to Philosophy* (81 – 116). Dordrecht: Springer.

Maddy, P. (1990). *Realism In Mathematics.* New York: Oxford University Press.

McDowell, J. (1981). Anti-Realism and the Epistemology of Understanding. In Parrett and Bouveresse, (Eds). *Meaning and Understanding* (225-248). Berlin: De Gruyter.

McGinn, C. (1980). Truth and Use. In *Reference, Truth, and Reality.* M. Platts, (ed.) Routledge. 19-40.

McGinn, C. (1982). Realist Semantics and Content-Ascription. *Synthese, 52*(1), 113 – 134.

Merricks, T. (2003). *Objects and Persons.* Oxford: Oxford University Press.

Miller, A. (2002). What is the Manifestation Argument? *Pacific Philosophical Quarterly, 83,* 353-383.

Putnam, H. (1978). *Meaning and the Moral Sciences.* London: Routledge and Kegan Paul.

Putnam, H. (1981). *Reason, Truth, and History.* Cambridge: Cambridge University Press.

Putnam, H. (1983). *Realism and Reason.* Cambridge: Cambridge University Press.

Putnam, H. (2012). Comments on Michael Devitt's 'Hilary and Me'. In M. Baghramian (Ed.), *Reading Putnam* (121 – 126). Hoboken: Taylor and Francis.

Raatikainen, P. (2010). The Semantic Realism/Anti-Realism Dispute and Knowledge of Meanings. *The Baltic International Yearbook of Cognition, Logic and Communication, 5*, 1-13.

Rasmussen, S. and Ravnkilde, J. (1982). Realism and Logic. *Synthese, 52*(3), 379 – 437.

Read, S. and Wright, C. (1985). Harrier than Putnam Thought. *Analysis, 45*(1), 56 – 58.

Rogers, W. (1971). *Mathematical Logic and Formalized Theories.* Amsterdam: North-Holland Publishing Company.

Rosen, G. (1994). Objectivity and Modern Idealism: What is the Question? In M. Michael and J. O'Leary-Hawthorne (Eds.), *Philosophy in Mind: The Place of Philosophy in the Study of Mind* (277 – 319).

Rosenkranz, S. (2003). Wright on Vagueness and Agnosticism. *Mind, 112*(447), 449 – 463.

Salerno, J. (2000). Revising the Logic of Logical Revision. *Philosophical Studies, 99*(2)*,* 211-227.

Sayre-McCord. (1988). Introduction: The Many Moral Realisms. In Sayre-McCord (Ed.), *Essays on Moral Realism* (1 – 23) Ithaca: Cornell University Press.

Shafer-Landau, M. (2003). *Moral Realism.* New York: Oxford University Press.

Shieh, S. (1997). Undecidability, Epistemology, and Anti-Realist Intuitionism. *Nordic Journal of Philosophical Logic, 2*(2)*,* 55-67.

Tarski, A., Mostowski, A., Robinson. R. (1953). *Undecidable Theories.* Amsterdam: North-Holland Publishing Company.

Tennant, N. (1997). *The Taming of the True.* New York: Oxford University Press.

Williamson, T. (1994). Never Say Never. *Topoi, 13,* 135-145.

Wright, C. (1992). *Truth and Objectivity.* Cambridge: Harvard University Press.

Wright, C. (1993). *Realism, Meaning and Truth* (2nd ed.). Oxford: Blackwell.

Wright, C. (2001). On Being in a Quandary: Relativism Vagueness Logical Revisionism. *Mind, 110*(437) 45 – 98.

Wright, C. (2003). Rosenkranz on Quandary, Vagueness and Intuitionism. *Mind, 112*(447) 465 – 474.

Yu, J. (2003). *The Structure of Being in Aristotle's Metaphysics.* Dordrecht: Kluwer Academic Publishers.