

A Two-Dimensional Analysis of Ethical Language

A.D. Glaser

A thesis submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Master of Arts in the Department of Philosophy.

Chapel Hill
2005

Approved by

Advisor: Geoffrey Sayre-McCord

Reader: Jesse Prinz

Reader: Keith Simmons

ABSTRACT

A.D. GLASER: A Two-Dimensional Analysis of Ethical Language
(Under the direction of Geoffrey Sayre-McCord)

Contemporary metaethics has developed under the implicit presupposition that there is a tension between the normativity and objectivity of ethical discourse. Mackie denied the existence of moral facts on the grounds that this tension renders ethical properties metaphysically queer. I refute Mackie's queerness argument by appealing to prevailing semantic theories. According to two-dimensional semantics, there are two intensions associated with meaning. The primary intension can be used to capture the psychological component that many traditional theorists believe is essential to meaning. The secondary intension accommodates the wide, or external, content of our terms. I argue that primary and secondary intensions can account for the normativity and objectivity of moral language, respectively.

To the kind, intelligent, self-confident, and self-expressed
lesbian, gay, bisexual, transgender, queer, and straight teenagers
with whom I have the privilege to work
and to whom I often defer moral judgment.

TABLE OF CONTENTS

I. Two-Dimensional Semantics.....	1
II. Refuting the Queerness Objection.....	10
III. Remaining Issues.....	28
Bibliography	33

I. Two-Dimensional Semantics

Perhaps the most pronounced divide in contemporary metaethics lies between analyses that give a central place to the putative action-guiding nature of ethical language, and those that prioritize its apparent objectivity. The former tend to cash out the meanings of ethical words subjectively, in terms of our emotions or motivational frameworks, while the latter tend to reduce ethical properties to descriptive or non-normative, natural (objective) properties of the world.

Due precisely to the tension between these two views, Mackie (1975) famously denied the existence of moral properties. He argued that the putative normativity and objectivity of ethical properties cannot be reconciled, and concluded that ethical properties are too metaphysically 'queer' to exist. In this paper I argue that prevailing semantic theories allow us to formulate an analysis of ethical language that accommodates its normativity and objectivity without rendering ethical properties metaphysically queer.

According to the referential theory of meaning, words and sentences mean what they do by virtue of what they denote. For instance, the term 'dog' means dog because it refers to dogs, while 'Charlie Chaplin' derives its meaning from the fact that it refers to Charlie Chaplin. The referential theory of meaning is threatened by a number of serious objections. First, it becomes clear that meaning includes more than reference when we consider that terms can differ in meaning and yet refer to the same object. For instance, 'Superman' and 'Clark Kent' have different meanings but refer to one and the same thing. Similarly, the

terms 'renate' and 'cordate' pick out the same set of objects, (assuming that all and only creatures with hearts possess kidneys), but differ in meaning.

This is a crude version of the referential theory of meaning, and more sophisticated versions have been formulated that resist the objections considered. One such version says that terms mean what they do by virtue of their reference across possible worlds. On this view, 'renate' and 'cordate' don't have the same meaning because they refer to different objects in different possible worlds. But even this more sophisticated version of the referential theory is up against its own devastating objections. Some terms refer to the same thing in every possible world, but still seem to differ in meaning. For instance, 'triangular figures' and 'trilateral figures' both refer to all and only triangles in every possible world, but their meanings differ. 'Triangular' means 'has three angles', while 'trilateral' means 'has three sides'.

Distinguishing between *intension* and *extension* allows us to formulate theories of meaning which resist the objections considered above. Where the extension of a term is the set of things the term is true of, it is clear that there's more to meaning than extension. We can call this additional part of meaning *intension*. Many believe that this notion of the intension of a term should be cashed out psychologically, and until recently most have agreed that meaning should be understood as intension, as wholly distinct from extension. It has also traditionally been thought that meaning determines extension.

According to these theories, the meaning of a term is not at all constituted by the thing the term denotes. Kripke (1980) and Putnam (1975) have raised decisive objections to these traditional theories. With a series of clever Twin Earth examples, Putnam motivated a claim that flew in the face of traditional theories of meaning. He showed successfully that the

meanings of many terms are in part constituted by external objects in the speaker's environment, by the term's extension. To motivate this claim, he considers two people, Oscar₁ and Oscar₂. Oscar₁ is an Earthling and Oscar₂ is his Doppelgänger on Twin Earth. Twin Earth is another planet in our solar system, exactly like Earth except that on Twin Earth, the wet and watery stuff that fills the oceans and the rivers, etc. is XYZ and not H₂O. We are to suppose that Oscar₁ and Oscar₂ are psychologically identical, and that both have a thought which they express with the words, "There is some water."

According to Putnam, this example raises difficulties for traditional theories of meaning, which presuppose that psychology determines meaning (in the sense of intension) and meaning (intension) determines extension. Since Oscar₁ and Oscar₂ are psychological duplicates, these presuppositions entail that their terms have the same intension. But when Oscar₁ utters 'water', it refers to H₂O and when Oscar₂ utters 'water', it refers to XYZ. If their intensions are the same, intension cannot determine reference since the referents of their terms differ. So the traditional theorist must drop one of his two presuppositions. He must deny either that psychology determines intension or that intension determines extension.

More precisely, Putnam considers two views one may take regarding the meaning of 'water'. First, one may claim that 'water' has the same meaning on Earth and Twin Earth, but on Earth water is H₂O and on Twin Earth it's XYZ. Accordingly, one may take 'water' to mean some lengthy descriptive phrase along the lines of 'the colorless, odorless, potable liquid that fills the lakes and the oceans, etc.', or, for short, 'the watery stuff'. Here, though the terms have the same intension, their extension will vary between possible worlds since *which* stuff is the watery stuff will vary between worlds. On Earth, the watery stuff is H₂O, so 'water' picks out H₂O here. But on Twin Earth, where the watery stuff is XYZ, 'water'

will pick out XYZ. Here the two terms differ in extension without differing in intension. But if intension determined extension, then two terms couldn't differ in extension without there being a corresponding difference in intension. So if we adopt this view, we must drop the presupposition that intension determines extension.

The second option Putnam considers is to deny that psychology determines meaning, insist that water is H₂O on Twin Earth and indeed in all possible worlds, and claim that when Twin Earthlings utter the term 'water' it differs in meaning from our term. This is the route that Putnam takes. He denies that meaning is wholly constituted by internal psychological states individuated by narrow content and insists that many terms have wide content, that their meanings are constituted in part by objects external to the speaker's psychology.

This means that the meaning of the term is individuated (at least) in part by facts about the world, facts concerning the world outside the speaker's mind. The meaning of Oscar's term 'water' is individuated at least in part by the stuff that is the watery stuff in Oscar's world. This conflicts sharply with the traditional view, which held that what a term means can be fully explicated by appeal to internal psychological states alone. As Putnam put it, "meaning just ain't in the head!"

On Putnam's view, for Oscar₁, the term 'water' means something different from what it means for Oscar₂, despite their psychological similarities. According to him, when Oscar₁ utters the word 'water', it means (in at least one sense of 'means') H₂O, while when Oscar₂ utters the term 'water', it means XYZ. Though Oscar₁ and Oscar₂ are psychological duplicates, their terms differ in meaning, and this difference is explained by the fact that the terms have wide content. The meaning of the term is not given by internal psychological properties alone, but includes external facts about the speaker's environment.

Kripke, too, argues effectively that this is the correct view to take about the meaning of 'water'. In Kripkean terms, this is the view that 'water' is a rigid designator. A term is a rigid designator iff it refers to the same thing in every possible world where that thing exists. Kripke argues that 'water' is rigid by considering a number of counterfactuals. He imagines a substance that superficially resembles water, a substance that possesses all of the conspicuous properties by which we identify water, but that is not in fact H₂O. He writes:

...would we say that some water wasn't H₂O? I think not. We would say instead that just as there is a fool's gold there could be a fool's water; a substance which, though having the properties by which we originally identified water, would not in fact be water.¹

An important consequence of Kripke's view is a distinction between a priority and necessity. That water is H₂O is necessary, true in all possible worlds, despite its being an empirical (a posteriori) discovery. 'Water', as we use the term, designates H₂O in every possible world, regardless of which things in that world superficially resemble water. Thus it is necessarily true that water is H₂O. Nonetheless, this is not an a priori truth, since the identification of water with H₂O required empirical investigation.

What is important at this point is not that we agree with Kripke and Putnam that 'water' is a rigid designator. Those who identify meanings with intensions will deny this claim. (And in fact, even Putnam, who construes intensions as descriptions, agrees that intensions are a component of the meanings of our terms.) They may insist that the watery stuff on Twin Earth is water, even though it isn't H₂O, just as the air there may not have the same chemical composition as our air.

For now, I wish to refrain from entangling myself in this dispute. I will proceed as if Kripke and Putnam are correct that 'water' is rigid, but nothing crucial hinges on this. What

¹ 1980, p. 128.

is important for my purposes is that we understand the difference between the two views. We can understand meaning in terms of intension (where what exactly an intension is remains to be spelled out) distinct from extension, or we can include extension in our theory of meaning, insisting that what a term denotes is a part of what it means.

Martin Davies, Lloyd Humberstone, David Chalmers, and others² have recently worked out a theory of semantics that accommodates both the intension and the extension of our terms. On their two-dimensional framework, two notions of necessity are distinguished, each corresponding to a component of a term's meaning. The first notion of necessity is truth in all possible worlds *considered as actual*. That is, something is necessarily true in this first sense iff it is true regardless of how the actual world turns out, or if it is true in whichever possible world we take be the actual world.

In this sense, it will be a necessary truth that water is the watery stuff, the clear, potable stuff that fills the oceans and the rivers, etc. If we consider *as actual* a world where the watery stuff is XYZ, then in that world water is XYZ. This is because 'water', as we use the term, picks out whatever it is that is watery in the actual world. So if we suppose that XYZ is the watery stuff in the actual world, then 'water' will pick out XYZ. If, in the actual world, the watery stuff is not XYZ, but something else, say, kryptonite, then 'water' will pick out kryptonite. In the actual world, 'water' picks out the watery stuff, regardless of which world is actual, or regardless of how the actual world turns out (in particular, regardless of *which* stuff is the watery stuff in the actual world.)

This notion of necessity corresponds to what Chalmers calls the *primary intension* of a term, and this is how we can understand what I have been calling the term's *intension*. Equipped with this new notion of necessity, we are now able to give a clearer explanation of

² See Davies and Humberstone (1980) and Chalmers (1996). Also see Jackson (1998) and Stalnaker (1978).

what primary intensions are. A term's primary intension may be thought of as a function from possible worlds *considered as actual* to extension. The primary intension tells us what counts as water counterfactually. The primary intension of 'water' is, roughly, 'the watery stuff'. It is not necessary that we think of intensions as descriptions, but the description helps to flesh out the character of the relevant function.

Primary intensions are important. First, they give us a way to capture the psychological component of meaning. The traditional theorists against whom Putnam argues thought that meanings ought to be identified with intensions, understood in psychological terms. Primary intensions accommodate the psychological component that these traditional theorists believe is essential. They also capture what we know a priori about our terms, what we know about our terms independently of how the actual world turns out.³ For instance, the primary intension of 'water' captures what we knew about the term's meaning before we knew that water was H₂O. Third, primary intensions serve an important role in theories of meaning since they explain how it is that our terms come to have the reference they do in the actual world. Primary intensions were familiar to the original baptizers of our terms. Finally, a term's primary intension is what we grasp when we come to understand a term. As Putnam suggests, in order to count as having acquired the term 'tiger', for instance, one must know, roughly, how to recognize tigers in the actual world; we must know that tigers are large striped cats. Putnam writes:

Suppose our...speaker points to a snowball and asks, "is that a tiger?" Clearly there isn't much point in talking tigers with *him*. Significant communication requires that people know something of what they are talking about...speakers are required to

³ Whether primary intensions are knowable a priori is currently up for debate. This does not matter much for the thesis I advance, since I employ primary intensions in order to accommodate the *psychological* component of meaning. Doubts that primary intensions are knowable a priori are generally driven by the idea that primary intensions include something *other than* psychology. Since what concerns us is the psychological component, we can set those doubts aside.

know something about (stereotypical) tigers in order to count as having acquired the word “tiger”...⁴

It is important to recognize, though, that in some sense water is merely contingently watery, and tigers are merely contingently large striped cats. There may be albino tigers without stripes, just as there may be non-watery water. The second notion of necessity distinguished on this two-dimensional framework is the traditional notion of counterfactual necessity with which most of us are already familiar. This is the notion that Kripke appeals to when he claims that water is necessarily H₂O.

As Kripke has made clear, since the watery stuff in the actual world is H₂O, whether a substance counts as water counterfactually depends on whether that substance is H₂O. Further, according to this second notion of necessity, it is *not* necessary that water is the watery stuff (despite its being a priori) since the watery stuff in various possible worlds may or may not be H₂O. In possible worlds where the watery stuff is XYZ and not H₂O, the watery stuff is not water. And in a possible world where H₂O is a dark black solid found only in deposits beneath the Earth’s crust, the water is not watery.

On Chalmers's view, this second notion of necessity corresponds to a term’s *secondary* intension. The secondary intension is a function from (counterfactually) possible worlds to extension. Secondary intensions are also important for theories of meaning. Kripke and Putnam have argued that some terms have wide content, that their meanings are constituted in part by objects external to the speaker’s mind. The notion of secondary intension provides a useful framework for taking stock of the wide content of our terms.

Two-dimensional semantics is useful since it allows us to accommodate both components of meaning considered above. Some identify meanings with primary intensions, cashing out

⁴ 1975, p. 168.

the notion psychologically, or in terms of some description or concept that speakers associate with their terms. On the other hand, like Kripke and Putnam, the referential theorist insists that our terms have wide content, that extension is part of meaning. Two-dimensional semantics accommodates what is important to both of these views. By incorporating both primary and secondary intensions, two-dimensional semantics occupies a sort of middle-ground between conflicting theories of meaning. Chalmers writes:

Both the primary and the secondary intensions can be thought of as candidates for the “meaning” of a concept. I think there is no point choosing one of these to qualify as *the* meaning; the term “meaning” here is largely an honorific. We might as well think of the primary and secondary intensions as the *a priori* and *a posteriori* aspects of meaning, respectively.⁵

According to Chalmers’s version of two-dimensional semantics, it is a conceptual truth (true by virtue of meaning) that water is watery. For this is necessarily true in the first sense of necessity, in the sense of being true in every possible world considered as actual. Water is the watery stuff (the odorless, colorless liquid that fills the rivers and lakes, etc.) in our world, regardless of how that world turns out, regardless of what that stuff actually turns out to be. So it is knowable *a priori* that water is watery. But, for Chalmers, it is also a conceptual truth that water is H₂O. For, given that the watery stuff in the actual world is H₂O, ‘water’ picks out H₂O in every counterfactual world. He writes:

...there are two varieties of conceptual truth, depending on whether we equate the “meanings” with primary or secondary intensions, paralleling the two varieties of necessary truth... “Water is watery stuff” is conceptually true and necessarily true in the first sense; and “Water is H₂O” is conceptually true and necessarily true in the second.⁶

Note that both of these conceptual truths hold, even though H₂O is not watery in every (counterfactually) possible world.

⁵ 1996, p. 62.

⁶ 1996, p. 66.

My purpose here is not to defend two-dimensional semantics, but to show how it can be used to formulate a theory of the meanings of ethical terms that preserves the normativity and objectivity of ethical language without upsetting a naturalistic ontology. Construing both primary and secondary intensions as (in some sense) conceptually linked to our terms allows us to proceed with this endeavor in the most straightforward way. The main argument I will reject (Mackie's argument from queerness) contends that ethical language has conflicting conceptual entailments. I think we can use the two notions of conceptual truth provided by the two-dimensional framework to show that the conceptual entailments of ethical language do not conflict. They don't conflict because they each correspond to different notions of conceptual entailment.

But I do not think that the force of the argument I advance depends on the truth of two-dimensional semantics. I am not in principle opposed to deviations from the picture of meaning I have elaborated, so long as it is properly understood how those deviations bear on the analysis of ethical language I propose. If either primary or secondary intensions are not, in fact, conceptually linked to our terms, then the response to Mackie I advance will have to be adjusted to accommodate that fact. I hope to make this all clearer in the next section of this paper.

II. Refuting the Queerness Objection

Two considerations compete for centrality in analyses of ethical language. The first is that ethical claims seem to be intrinsically normative. The claim that Φ -ing is right seems in some sense to point towards Φ -ing. Ethical claims seem to have practical import, to tell us what to do. The second consideration is the putative objectivity of ethical language. Ethical claims purport to describe the world and to be true or false independently of our theorizing about them. For Mackie (1977), the normativity and objectivity of moral claims wedges them awkwardly between our decisions regarding what to do and our beliefs about the way things are. Mackie claims that the normativity and objectivity of ethical language renders ethical properties metaphysically queer, and concludes that there are no moral facts.

It seems there are two central features of morality that contemporary metaethicists are concerned to accommodate. The first is the important role that our ethical beliefs play in our decisions regarding what to do. Ethical beliefs seem to have a practical component; they seem to bear importantly on how we choose to live. This is what is meant by the *normativity* of ethical language; ethical beliefs seem to have implications for what one is *to do*. Also, there is a strong tie between ethical beliefs and the emotions and motivations of moral agents. If I have very strong moral beliefs about Φ -ing, it seems unlikely that I will be emotionally indifferent towards Φ -ing.⁷

⁷ The normativity of ethical language and its intrinsic link with our motivational states are not unrelated. Attempts to explain the normativity of ethical language often appeal to an intrinsic link with our emotions and motivational frameworks. I believe that, ultimately, an account of this sort will succeed. But I will not develop such an account here. I will argue for an intrinsic link between moral beliefs and motivation and hope that the normativity of ethical language will follow fairly straightforwardly from there.

The second feature of morality that has captured the attention of metaethicists is that ethical claims seem to be *objective*, to be true or false independently of our theorizing about them. Ethical claims purport to describe how the world is. The claim that Φ -ing is wrong purports to ascribe a property to Φ -ing, to tell us something about how Φ -ing is. Also, we seem to be able to employ ethical language in deductively valid arguments. That Φ -ing is always wrong seems to deductively entail that my Φ -ing yesterday was wrong. Moreover, there is an important relationship between ethical claims and non-normative features of the world. Ethical questions seem to have correct answers, and those answers seem to be made correct by how the world is non-normatively. There is an important and deep connection between the fact that Φ -ing is, say, an act of wanton cruelty and the fact that Φ -ing is wrong. These considerations attest to the putative objectivity of ethical language.

The central thesis of this paper is that we can account for the both the normativity and objectivity of ethical terms by appeal to their primary and secondary intensions, respectively. This is intended primarily as a response to Mackie's argument from queerness. There is no problematic tension between the objectivity and normativity of ethical claims, since each corresponds to a distinct component of meaning. I believe that a two-dimensional framework facilitates a successful response to Mackie, allowing that ethical claims are both intrinsically normative and intrinsically objective. A two-dimensional picture of ethical language thus respects the motivations for the leading views in metaethics, views that have previously been supposed to be at odds.

Those who take the normativity of ethical language seriously often posit an internal or necessary link between ethical beliefs and the emotions and motivational frameworks of moral agents. Those who do so are called *internalists*. Internalists believe that there is a

necessary link between believing, say, that Φ -ing is right and being moved to Φ , or at least feeling 'tugged' towards Φ -ing. Of course, not all of our ethical beliefs are motivational. We don't always act in accord with them. But the idea is that believing, say, that Φ -ing is right at least involves (logically, or necessarily, entails) feeling (or at least being disposed to feel) some sort of a pull or a tug towards Φ -ing. Internalism does not require that that tug invariably result in action.

Internalists believe that this link between morality and motivation deserves a central place in analyses of our ethical language. Many accommodate the intrinsic link between morality and motivation by cashing out the meaning of ethical language entirely in terms of our emotions and motivational frameworks. For instance, Ayer (1936) and Stevenson (1937) hold that ethical claims are essentially expressions of emotion. If to say that Φ -ing is right is, for example, simply to assert or express one's approval of Φ -ing, it is clear how rightness is necessarily linked to our attitudes and emotions. By positing an intrinsic link between moral beliefs and motivation, the internalist has an easy time explaining why ethical beliefs are normative, why they play the role that they do in our decisions regarding what to choose. Ethical language plays an important role in our decision-making processes because the emotional states to which it is necessarily linked play an important role in those processes.

I believe that much work remains to be done in psychology before the internalist can say for certain which psychological states are relevant to the meanings of ethical language. Research in emotions is so new that at this point I believe it is difficult to provide anything but a very rough approximation. Moreover, the language we have for describing the vast range of emotions we feel regarding choice is pitifully deficient. What is the difference between approving of something and liking it? Are there important distinctions to be drawn

among the vast range of emotions that count as approval? Before these types of differences are articulated and a precise and thorough account of the relevant emotional states is developed, the best the internalist can do is to gesture towards the kinds of states that might be involved in understanding the meaning of ethical language. As of now, we are very far from providing an account of moral emotions that is even close to being sufficiently detailed.

Nonetheless, I believe that internalists are onto something important about ethical language. To understand what the internalists have gotten right about moral language, consider again two-dimensional semantics. The primary intension of 'water' (the watery stuff) allows us to accommodate the phenomenological or psychological states that we associate with the meaning of our term 'water'. The primary intension also tells us what is a priori about the term 'water', what we knew about the meaning of 'water' before we knew what water *was*. It is also important to understanding how the term comes to refer to what it does in the actual world. The primary intension tells us the conspicuous properties of water, which enable us to know what stuff in the actual world counts as water. Finally, if two-dimensional semantics is correct, then one must be familiar with these properties to count as having grasped the meaning of the term.

In precisely the same way, I believe that the psychological states appealed to by the internalist play an important role in analyzing the meaning of ethical language. Consider rightness. There is a lack of consensus about which things are included in the extension of the term 'right'. Even among philosophers, there is deep disagreement. According to Mill, the extension of the term 'right' includes all and only those acts that maximize happiness. According to Kant, it includes all and only those acts that are universalizable. If we are to count Kant and Mill (and others) as meaning the same thing by 'right' then there must be

more to the meaning of the term than extension. We are uncertain about what rightness consists in, but are nonetheless in a position to say something about what the term means. By appealing to the conspicuous features of rightness (say, its tendency to elicit approval, or whatever the best formulation of internalism tells us is conceptually linked to believing that something is right), I think we can flesh out the primary intension of 'rightness'.

On this view, (something like) 'rightness tends to elicit approval' has the same modal status as 'water is watery'. These will both be a priori truths, true by virtue of meaning, necessarily true in the sense of being true in every possible world considered as actual; in particular, regardless of which stuff in the actual world tends to elicit approval. As in the case of 'water', I believe the primary intension of 'rightness', characterized by the conspicuous features of rightness, the impact it has on our emotions and motivational frameworks, tells us what we know a priori about rightness. It tells us what we know about rightness despite our uncertainty about what rightness *is*.

Why construe the primary intension of rightness in this way? Why believe that the primary intension of rightness is the emotional and motivational states that the internalist holds dear? Primary intensions are known a priori because they tell us what our terms apply to independently of how the actual world turns out. Primary intensions are functions from various ways the actual world might be, to extension. They tell us what our terms pick out *if* the actual world is a certain way, what the term applies to counteractually.

By observing what counts as water in the actual world, we can see what we'd count as water counteractually. In the present case, we look to see if a thing has the properties we associate with water (roughly, being wet and watery); those are the properties by which we judge whether something counts as water in the actual world. We call the stuff that meets that

description 'water', and we deny that things that fail to meet that description, big solid orange things for instance, are water. And it is this fact, the fact that we use the term 'water' to pick out the watery stuff, which determines what 'water' comes to refer to in the actual world. We use 'water' to pick out the watery stuff in the actual world, and in the actual world that stuff happens to be H₂O. That's why 'water' refers to H₂O.

According to Putnam, the traditional meaning theorist thinks that meanings are essentially psychological entities. Thinking of primary intensions as functions from counterfactual possibilities to extension gives us a way to accommodate the psychological component of meaning. The primary intension of 'water' can be characterized by what it's like phenomenologically for us to interact with water. The primary intension of 'water' can be fleshed out with some description along the lines of 'clear, odorless, potable, wet and watery liquid' because that is what water is like to us.

Much like the traditional meaning theorist, the traditional internalist thinks that the meanings of ethical language should be understood psychologically. I think that in both cases, primary intensions can account for what's central to the view. Internalists believe that the meaning of ethical language has a psychological component, that ethical language is intrinsically linked to our emotions and motivational frameworks. Just as we may use primary intensions to capture the psychological component of meaning, which is important to the traditional theorist, so we may use primary intensions to capture the psychological component of the meanings of ethical language, which the internalist believes are essential. Just as we may use primary intensions to capture the psychological component of the meaning of 'water' (the phenomenological impact of water on us, its being wet and watery) so we may use it to capture the psychological component of 'good' (the phenomenological

impact of goodness on us, its tendency to elicit approval). And in both cases, we capture the psychological component in specifying what our terms pick out counterfactually.

To know the primary intension of goodness or rightness, we must know what would count as good or right counterfactually, across possible worlds considered as actual. To discover this, we should look at what in the actual world gets called 'good' or 'right' by competent speakers. And here we may appeal to psychology. What are the conspicuous phenomenological features of things, by which we judge whether a thing is good or right in the actual world? What relevant feature is shared by those things we believe are good? On what basis do we come to believe that they are good? That will tell us what we'd count as good counterfactually.

Here, I think the internalist provides the answer. The traditional externalist will insist that our emotional responses are not a part of our terms' primary intensions, but this view is untenable, since apart from our emotional responses, it seems that there isn't *any* feature that underlies ethical ascriptions by competent speakers. I use 'right' roughly to refer to things that promote happiness, religious zealots use it to refer to what the bible commands, and the terrorists who guided airplanes into the Twin Towers on 9/11 use it to refer to something altogether different. But these ascriptions aren't just arbitrary. It isn't merely accidental that we use the same term. What our ascriptions have in common is roughly *how we feel* about the things we use our moral terms to refer to.

We are far from being able to say precisely what the relevant emotions are, but it should suffice to point out that we all have positive sentiments about the things we call right and good. I have those sentiments about human happiness, religious zealots have them about following biblical commandments, and so on. We should look to the internalist to provide a

more fine-tuned account of which emotional or phenomenological states are the relevant ones. I think the correct account of the primary intensions of ethical language will incorporate information from psychological research on emotions that is not yet available.

In light of the parallel between primary intensions and the psychological states that internalists believe are conceptually linked to ethical language, it makes sense why internalism is most plausibly formulated as a thesis about ethical *beliefs*. Though some versions of internalism posit a conceptual link between ethical *properties* and motivation, these versions are not very convincing. Not many people think that the fact that some action is right conceptually entails that we feel moved to do it. But if internalism is more modestly formulated as the thesis that ethical *beliefs* conceptually entail something about the believer's motivational states, then this conceptual link can be understood simply in terms of a *prioricity*. Recall that primary intensions are knowable a priori since they obtain independently of how the actual world turns out.

Moreover, once we capture the primary intension of 'rightness', we will have an explanation for what makes the correct normative theory true. Just as the primary intension of 'water' explains how it is that 'water' comes to refer to H₂O in the actual world, so the primary intension of 'rightness', construed in the way I suggest, will explain how it is that 'rightness' comes to refer to maximizing happiness, or universalizability, or whatever rightness turns out to be. Finally, if we construe things in this way, then for one to count as having grasped the term 'rightness', he will have to be familiar with the emotional tug that accompanies beliefs about rightness. One will have to be impacted motivationally in the right sorts of ways by one's moral beliefs in order to count as having those beliefs at all. This is just what the internalist wants.

Of course, not everyone accepts internalism. David Brink (1989) has challenged internalism by insisting that the amoralist – one who accepts that she has moral obligations but feels wholly unmoved by them – is a conceptual possibility. Brink believes that it is possible for someone to believe, say, that it is right to Φ , and yet remain wholly indifferent towards Φ -ing. In response to Brink, internalists have denied that such an agent would genuinely possess the belief that Φ -ing is right. They insist that in her vocabulary, the term ‘right’ does not mean what it does in the traditional moral sense, but has another meaning.

Even if we reject two-dimensional semantics, it is useful for framing the debate between Brink and the internalists. If we characterize the primary intension of ‘right’ in terms of the emotional and motivational states that internalists claim are intrinsically linked to ethical language, we can come to see how the debate between Brink and the internalists hinges on which theory of meaning we adopt. Whether the amoralist means what we do by ‘right’ depends on whether primary intensions are a component of meaning.

If Chalmers is correct and they are, then Brink is mistaken. ‘Right’ means something different for the amoralist, just as ‘green’ means something different to someone who is congenitally blind, and ‘water’ means something different when it picks out H_2O in worlds where the H_2O is not watery. These differences in meaning arise if we construe primary intensions as essentially phenomenological. Since blind people and sighted people associated different phenomenological features with their color terms, their color language differs in meaning. If, on the other hand, primary intensions are not a part of meaning, and two speakers mean the same thing so long as their language converges on the world, then Brink is correct: the amoralist is indeed a conceptual possibility.

Brink believes that ethical claims are objective. He believes that they are truth-apt, and that their truth or falsity obtains independently of our thoughts and feelings about them. This view is motivated by the kinds of considerations discussed at the outset of this section of the paper. Moral questions seem to be the kinds of things we can answer correctly or incorrectly, and the correct answers seem to be available through argumentation. Further, it seems that the truth or falsity of moral claims is something that is determined by how the world is, not by how we take it to be. What we believe to be true morally doesn't seem to make it so.

Those who are concerned to accommodate the putative objectivity of ethical language often attempt to do so by reducing normative properties to descriptive or non-normative objective properties of the world. Let's call them externalists. By reducing the ethical to the descriptive, externalism explains the link between natural facts – that some actions cause suffering, that some creatures have like interests – and the moral facts, thereby accommodating the objectivity of ethical language. By claiming, like Mill, that rightness is, by definition, maximizing happiness, or like Kant, that rightness is universalizability, externalists explain how ethical claims get to be true or false and why their truth or falsity is determined by how the world is. For Mill, torturing babies for fun is wrong because torturing babies for fun fails to maximize happiness, and that's just what wrongness *is*, by definition. If ethical claims logically reduce to non-normative claims, then ethical claims get to be true or false in the same non-problematic way that descriptive claims do.

I believe that we ought to take externalists to be providing the secondary intensions of our ethical terms. Suppose that the best normative theory tells us that 'rightness' picks out maximizing happiness. If we construe maximizing happiness as the secondary intension of rightness, then 'rightness is maximizing happiness' is necessarily true, true by virtue of the

meanings of the terms, true in every (counterfactually) possible world. Construing the secondary intensions of moral vocabulary in terms of non-normative features of the world explains the objectivity of morality. It is a consequence of such an analysis that ethical claims are truth-apt and their truth or falsity is fully determined independently of our thoughts about them, by how the world is non-normatively.

Secondary intensions are functions from counterfactually possible worlds to extension. To know the secondary intensions of moral vocabulary, we must know what counts as moral counterfactually. Suppose that in another (counterfactually) possible world there are no moral judges or agents. There are no people talking moral talk, forming moral beliefs, or experiencing moral emotions. Suppose instead that the only humans there are two toddlers, chained to a tree, cold and wet, sitting in their own filth, being barely kept alive with an intravenous feeding tube, and screaming in terror as they are repeatedly shocked with high voltages of electricity.

I think that the suffering those toddlers experience is bad, and I think its badness obtains independently of anyone's judgment that this is so. Were I (or a whole population of adult humans, for that matter) in that world, and wired so as to be totally insensitive to the suffering of others, the suffering the toddlers experience would nonetheless be bad. This, of course, is not a conclusive argument for construing the secondary intensions of moral vocabulary as response-independent properties, but it speaks in favor of such a construal. I believe that successful externalist arguments have been formulated elsewhere to show that what counts as right or wrong counterfactually depends on non-normative facts about the world, and not on our counterfactual emotional responses to those facts.⁸

⁸ See Sayre-McCord (1997).

Construing primary intensions as the motivational states the internalist believes are necessarily linked to ethical language accommodates the a priori link between moral belief and motivation. Construing secondary intensions as the descriptive features that are necessarily linked to moral properties explains the objectivity of morality. Together, I think these two claims account for what's most central to our moral talk. They do so without giving rise to metaphysical worries about queerness, and they resist the most pressing objections faced by traditional internalists and externalists alike.

Nonetheless, there is commonly thought to be a tension between the normativity and objectivity of ethical language. Internalists and externalists tend to agree that ethical claims cannot be both logically reducible to non-normative claims and intrinsically linked to our motivational frameworks, but they disagree over which of these putative features of ethical language to reject.

Mackie (1977) believes that ethical language has *both* features, and concludes that that renders ethical properties metaphysically queer. He claims that there are no ethical facts. Due to the objectivity and normativity of ethical language, belief in ethical facts commits us to the existence of mysterious metaphysical entities, on par with Moorean non-natural properties or Platonic forms. Mackie agrees with externalists and internalists that the normativity and objectivity of ethical language pull in opposite directions. The normativity of ethics suggests that there is a necessary link between moral language and our emotions or motivational frameworks, while the objectivity of moral language seems in principle to leave open the emotions and motivations of the speaker.

The traditional externalist will not be comfortable construing primary intensions as emotional or motivational responses. Similarly, the traditional internalist will probably want

to jump ship once we include non-normative properties in the secondary intensions of our terms. Both sides are squeamish about accepting the arguments of the other. But the externalists' reluctance to accept arguments for internalism, on the one hand, and the internalists' rejection of the objectivity of morality, on the other, are both typically motivated (not always explicitly) by concerns akin to Mackie's worries about metaphysical queerness.

The implicit presupposition that there is a tension between the objectivity and normativity of morality has had a profound impact on developments in contemporary metaethics. Internalists and externalists are thought to be *opponents*. It is commonly accepted that the reducibility of ethical properties to non-normative ones at least poses a threat to the normativity of morality, or the intrinsic link between moral belief and motivation. That is why internalists have denied that ethical properties are reducible to non-normative ones. They often reject the intuition that ethical claims are objective, and insist that moral language can be fully cashed out in terms of our emotions and motivational frameworks. And both internalists and externalists alike take the central burden faced by the traditional internalist to be the justification of the denial of the objectivity of morality.

Ayer (1936) and Stevenson (1937), for instance, would deny that there is anything more to the meanings of our ethical terms than the emotional states they express. But the obvious challenge for them is to account for considerations that lead us to believe that morality is objective. They have to explain why ethical claims purport to describe the world, why ethical questions seem to have right and wrong answers, why ethical language can be mobilized in deductively valid arguments, and why those arguments seem to yield ethical knowledge. Contemporary metaethicists differ on whether they respond to these challenges successfully, but it is agreed that that is where attention is due.

Symmetrically, externalists, who accommodate with ease the objectivity of ethical language, have traditionally sought to refute internalism. They have taken a rejection of internalism to be important to a successful reductive analysis of ethical language. It is commonly agreed that a successful defense of externalism requires a refutation of internalism. But since they reject internalism and insist that ethical claims are conceptually reducible to non-normative ones, externalists have a hard time accommodating the normativity of ethical language. This opens up a new problem for the externalist, the problem of explaining why we ought to be moral. If ethical facts are not intrinsically normative, it isn't clear why we should act in light of them. The motivational force of ethical obligations is notoriously difficult for the externalist to accommodate.

In response to this difficulty, many externalists have resorted to denying the intrinsic normativity of morality altogether. Some claim that it just so happens that we feel motivated by good things, that there is no deeper, semantic connection between morality and motivation. For them, most of us just so happen to like good things in much the same way that we just so happen to like chocolate. They ignore the role that our emotions play in fixing the reference of our moral terms in the first place.

Arguments have already been given for why rejecting either the normativity or the objectivity of morality would be a mistake. *Both* aspects of moral talk are important to analyses of moral language, and both kinds of arguments (those that establish an a priori link between moral belief and motivation, *and* those that establish a counterfactually necessary link between moral properties and descriptive properties) are accepted by significant proportions of the philosophical community. But while the philosophical community as a

whole has given both kinds of arguments the support they are due, very few individuals accept *both* arguments.

If we have reason to believe that there are objective moral facts, facts that are intrinsically normative (and we do), a two-dimensional construal of ethical language allows us to posit the existence of such facts without upsetting a naturalistic ontology. I believe that successful arguments can and have been given to motivate both the objectivity and normativity of ethical language. Most people who deny one or the other of these claims don't believe that ethical objectivity and normativity can fit together, and so they resort to denying one or the other. And when they do so, the central burden they face is to justify their denial of either the objectivity or normativity of ethical language.

The view I advance relieves them of this burden. We can have it both ways; ethical language can be both normative and objective, and we don't then have to follow Mackie in concluding that there are no ethical facts. The construal I propose allows us to locate ethical facts (objective normativity) non-problematically within a naturalistic framework, as not altogether distinct from our own emotional states and descriptive features of the world. This accommodates the objectivity and normativity of ethical language, something I think we should want to do, and something I think we would be more inclined to try to do, if we didn't believe, under Mackie's influence, that it would force us to admit metaphysically queer properties into our ontology.

Two-dimensional semantics gives us a way to account for the psychological component of meaning, without forcing us to deny that some terms have wide content. That is the central motivation for viewing ethical language within the framework of two-dimensional semantics. It accommodates the psychological component of moral language without forcing us to deny

its conceptual link with the external world; we can preserve the intrinsic link between morality and motivation without denying that ethical properties logically reduce to non-normative ones.

Ethical language gets to be intrinsically linked to our motivational frameworks, since the primary intensions of ethical vocabulary are cashed out in terms of our emotions and motivational frameworks. At the same time, ethical properties may nonetheless be logically reducible to objective features of the world. Such features are captured by the secondary intensions of our ethical terms. The correct normative theory will tell us what rightness *is*, and that will give us the secondary intension of rightness. We can maintain that ethical properties are intrinsically normative, without denying that ethical terms, by virtue of what they mean, pick out objective features of the world.

On this view, it is a conceptual truth both that rightness tends to elicit approval, and that rightness is, say, maximizing happiness. This does not render ethical properties metaphysically queer, nor does it entail that maximizing happiness necessarily tends to elicit approval. Analogously, on Chalmers's two-dimensional framework, it is a conceptual truth both that water is watery and that water is H₂O. But this does not entail that H₂O is necessarily watery, because each of these conceptual truths corresponds to a different component of meaning. The normativity and objectivity of ethical claims are reconciled on the grounds that meaning has two dimensions.

This accommodates everything that Mackie, the externalists, *and* the internalists hold dear. A two-dimensional analysis of ethical language provides one unified theory that respects the motivations behind each of the dominant views in contemporary metaethics,

views which have traditionally been thought to be irreconcilable. The two-dimensional analysis casts the central streamlining issues of contemporary metaethics in a new light.

It also resolves some of the most pervasive difficulties faced by proponents of the most commonly held views in ethical discourse. It answers Moore's (1903) open question objection. On the two-dimensional analysis, moral properties reduce to natural properties. Rightness, for instance, may reduce to maximizing happiness. Nonetheless, it will remain an open question whether rightness is maximizing happiness. It will remain an open question in the sense of not being knowable a priori. The two-dimensional analysis also answers the 'why be moral' question. It explains why we (necessarily) ought to be moral, while allowing for the identification of ethical with natural properties.

Finally, even if two-dimensional semantics is not the correct theory of meaning, it is useful for illuminating the relationship between the central controversies in metaethics and parallel disputes in philosophy of language. Our own motivational states compete with natural features of the external world as candidates for the meanings of ethical language. The controversy can be understood as a symptom of unresolved disputes in philosophy of language. There, primary and secondary intensions compete for centrality in theories of meaning. There is no general consensus over whether two-dimensional semantics is correct, and if not, whether primary or secondary intensions capture our notion of meaning. Hence the parallel dispute in metaethics.

No matter which theory of meaning turns out to be correct, a response to Mackie is at hand. If two-dimensional semantics is correct, then the intrinsic objectivity and normativity of ethical language fall neatly into place. If it's not correct, and meaning consists in primary but not secondary intensions, Mackie's assumptions are wrong. Ethical claims are not

intrinsically objective, but it is natural to believe that they are, since secondary intensions are commonly (mistakenly) thought to be a component of meaning. Likewise, if it is secondary, not primary, intensions that are the meanings of our terms, we can explain away our intuitions that ethical language is intrinsically normative on the grounds that primary intensions are commonly mistaken for meanings.

Recently, upon returning from a vacation at the beach, I was asked by a friend if the weather was good there. I paused for a moment, unsure of how to respond. The weather had been cold and rainy the entire time, but I had enjoyed it very much. Whether the correct answer was “No, and I loved it!” or “Yes, it was cold and rainy,” (or perhaps even, “Sorry, you’ll have to rephrase the question; I don’t share your concept of ‘good weather’,”) depends on which theory of meaning is correct. In particular, it depends on the role of primary and secondary intensions in the correct theory of meaning.

III. Remaining Issues

I have attempted to refute Mackie's argument from queerness by appeal to two-dimensional semantics. But a two-dimensional analysis of natural kind terms may seem to many to be much more plausible than a two-dimensional analysis of ethical language. Ethical terms differ from natural kind terms in important ways, and many may take these differences to undermine the plausibility of a two-dimensional analysis of ethical language. I don't believe that these objections are fatal, but I do believe that the account I have proposed will need to be fine-tuned in order to cope with them. In these final pages, I wish to draw attention to the work that remains to be done before a two-dimensional analysis of ethical concepts can succeed. An adequate discussion of these issues is beyond the scope of this paper; I merely wish to point out what I take to be the most important issues that remain to be addressed.

First, consider the broad range of properties to which moral predicates apply. The two-dimensional analysis seems to require that ethical predicates each pick out a single feature of the objective world. Yet it seems that there are a number of various objective properties that can constitute, say, goodness. This already seems to set ethical language apart from natural kind terms, which have traditionally been thought to pick out single features of the world; 'water' picks out H₂O, 'gold' picks out Au, etc.

And there is a deeper problem. Not only does 'good', say, seem irreducible to a *single* natural property but also it seems that convergence by all speakers even on a single *set* of properties is unlikely. The primary intension of 'water' determines what 'water' refers to in

the actual world. But it seems that the primary intension of e.g. ‘goodness’, captured in terms of our psychological states, cannot determine reference in that way since moral agents are moved into the relevant psychological states by such different sets of properties. The set of properties that elicit approval in me are strikingly different from the set of properties that elicit approval in religious zealots. How does the two-dimensional analysis handle this lack of convergence?

Another closely related problem is that the two-dimensional analysis seems not to allow for the fallibility of ethical beliefs. According to it, the reference of ethical terms (good, say) is determined by their primary intensions (tends to elicit approval). But it seems that the good is not whatever *elicits* approval, but what *ought* to elicit approval. It seems conceptually possible that we tend to approve of the wrong things. But the two-dimensional analysis seems to rule out this possibility a priori.

Regarding the first worry, if it turns out that ‘good’ fails to apply to a single natural property (and it’s not clear that it will), this will not in fact set it apart from natural kind terms. There are several examples of natural kind terms that are known to apply to more than one kind. ‘Jade’ is one such example. Putnam himself points out that ‘jade’ applies to two different minerals, jadeite and nephrite. He insists that if ‘water’ were also found to apply to two different kinds (if both H₂O and XYZ were prevalent on Earth, for instance) then we should say that there are two kinds of water.⁹ ‘Jade’ is typically thought to be an isolated instance of a natural kind term that doesn’t pick out a single kind. But recent developments in philosophy of language attest to the contrary. It has been convincingly argued that many common natural kind terms (among them ‘lily’, ‘mushroom’, and even ‘dog’) don’t in fact

⁹ 1975, p. 160.

refer to singular kinds.¹⁰ We may have work to do to figure out how to handle this (potential) lack of convergence, but we will be in good company.

But what of the other two objections, that convergence by all speakers even on a single *set* of properties is unlikely, and that ‘good’, for instance, picks out what *ought* to elicit approval, not what actually does? I think that both of these objections can be answered by appeal to ideal circumstances. ‘Good’ will not pick out what *actually* elicits approval in us, but what would elicit approval in us in certain circumstances. There is a familiar precedent for such an appeal.

Color terms can easily be construed on a two-dimensional framework. Suppose we construe the primary intensions of color terms as the phenomenological states that we associate with our color terms, and the secondary intensions as the physical surface properties that underlie our color ascriptions. Then we will likewise need to appeal to ideal (normal) circumstances in explaining how primary intensions determine reference. ‘Red’ will not apply to *whatever* causes us to have the relevant phenomenological experiences. It won’t apply to white things under red lighting, for instance. ‘Red’ applies to whatever causes us to experience the relevant phenomenological states *in certain circumstances*. A successful two-dimensional analysis of color terms must appeal to normal circumstances.

I think that a similar move must be made in the ethical case. Ethical language is like color terms and unlike natural kind terms in that a complete two-dimensional analysis of ethical language will appeal to ideal circumstances. I am unable to fully develop this part of the

¹⁰ Dupre (1981) describes several such examples. Take the term ‘lily’. Several species of the lily family (Liliaceae) are commonly referred to as lilies: among them are the Lonely Lily, the Avalanche Lily, and the Desert Lily. But the extension of our term ‘lily’ clearly does not include the entire family. It doesn’t, for instance, include garlic and onions, which are also both members of the family Liliaceae.

account in the space I have, but I take this to be the central challenge remaining for a successful defense of the view. I believe that this challenge is surmountable.

How could an appeal to ideal circumstances be developed? Perhaps we should proceed via a form of reflective equilibrium. We begin by noting the clearest morally relevant cases. We note that torturing babies for fun is wrong, for instance. Then we may look to other events that are relevantly similar to the clearest cases, say, torturing animals for new perfume. From there we will probably be able to unveil the inconsistencies in our attitudes and emotional responses, and to figure out what about our current circumstances leads to those inconsistencies. We may conclude that our responses to e.g. animal experimentation are manipulated by the media and by the people around us, and that we would have different emotional responses if we were more attentive to the facts of the matter. If we proceed via reflective equilibrium, we may come to see how ideal circumstances differ from our current ones.

Many may doubt the plausibility of convergence even under idealized circumstances. Often, these doubts are motivated by the presence of radically divergent moral views within and between cultures.¹¹ I am able to say a few things in response to these concerns, though the issue deserves much more attention than I will give it here. First, I think that insisting that convergence is possible is a much more plausible stance than denying the intrinsic normativity or objectivity of ethical language altogether. Insisting on the possibility of convergence is a better difficulty to be stuck with than allowing that either the normativity or objectivity of ethical claims is merely accidental.

Second, and more importantly, humans are ignorant, self-deceived, and manipulated by illegitimate social pressures. Ideal circumstances are radically different from our current

¹¹ See, e.g., Prinz (forthcoming).

circumstances, and it is inappropriate to rule out the possibility of convergence from here. To be sure, the diversity that exists among ethical views across cultures is astounding. But how many of those ethical systems are founded on false metaphysical views? How many ethical systems are unconsciously manipulated by pervasive social norms that are stifling or harmful? Consider, for instance, that it is not uncommon for queer teenagers in this country to be slicing up their arms and legs with knife blades and hanging themselves from basketball hoops because of a strong social stigma against their sexual identities. Surely we ourselves are not immune to the insidious effects of such strongly entrenched social norms. The possibility that humankind has, underneath, *some* shared aims should not be dismissed from the warped circumstances in which we currently find ourselves.

Thinking of ethics in terms of our own motivations and the non-normative features that move us involves a radical shift from the kind of ethical theorizing that has dominated our species for thousands of years. Throughout the centuries we have sought ethical guidance from sources outside of ourselves. Thus we come to care about things that don't *really* matter to us. The two-dimensional analysis shifts the focal point of ethics to the real world, to the descriptive features relevant to the choices we make, and to the emotional responses we have to the world as we experience it. I will continue to hold out hope that such a shift may effect a convergence of the basic underlying values of humankind.

BIBLIOGRAPHY

- Ayer, A.J. *Language, Truth, and Logic*. 2nd ed. London: Gollancz, 1936.
- Brink, D. *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press, 1989.
- Chalmers, D. *The Conscious Mind*. Oxford: Oxford University Press, 1996.
- Davies, M., and Humberstone, L. "Two Notions of Necessity," *Philosophical Studies*, 38 (1980): 1-30.
- Dupre, J. "Natural Kinds and Biological Taxa," *The Philosophical Review*, 90 (1981): 66-90.
- Jackson, F. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press, 1998.
- Kant, I. 1785. *Foundations of the Metaphysics of Morals*. London: Hutchinson, 1956.
- Kripke, S. *Naming and Necessity*. Cambridge: Harvard University Press, 1980.
- Mackie, J.L. *Ethics: Inventing Right and Wrong*. England: Penguin Books, 1977.
- Mill, J.S. 1861. *Utilitarianism*. Indianapolis: Hackett Publishers, 1979.
- Moore, G.E. *Principia Ethica*. Cambridge: Cambridge University Press, 1903.
- Prinz, J. *The Emotional Construction of Morals*. Forthcoming.
- Putnam, H. "The Meaning of 'Meaning'," in K. Gunderson, ed., *Minnesota Studies in the Philosophy of Science, Vol. 8: Language, Mind, and Knowledge*. Minneapolis: University of Minnesota Press, 1975.
- Sayre-McCord, G. "'Good' on Twin Earth," *Philosophical Issues*, 8 (1997): 267-292.
- Stalnaker, R. "Assertion," in P. Cole, ed., *Syntax and Semantics: Pragmatics*. Vol. 9. New York: Academic Press, 1978.
- Stevenson, C.L. "The Emotive Meaning of Ethical Terms," *Mind*, 46 (1937): 14-31.