

WORDS AND DOING THINGS: WHAT MEANINGS CAN DO AND THE ROLE OF
MOTIVATIONAL STATES IN DETERMINING MEANING

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

In Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

David Fielding

May 2019

© 2019 David Fielding

WORDS AND DOING THINGS: WHAT MEANINGS CAN DO AND THE ROLE OF
MOTIVATIONAL STATES IN DETERMINING MEANING

David Fielding, Ph. D.

Cornell University 2019

This dissertation consists of three independent papers that are linked by two threads. The first thread is the question of what sorts of explanatory roles can be played by propositions and semantic values. The first paper, “Doing without Minimal Propositions”, explores this question by looking at the role of semantic values in conversation, in particular examining the claim that only a small range of linguistic expressions need context in order for sentences containing them to have a proposition as their semantic content. I argue to the contrary that the largely context-insensitive propositions posited by such a view are not needed to explain the range of linguistic facts that motivates them.

My second paper, “Problems for Metasemantic Expressivism”, explores the question of whether the semantic values standardly assigned to normative expressions can help explain the special connection that normative thought and discourse seem to have to non-cognitive mental states: states like desire that play a special role in motivating action. I argue against the metasemantic expressivist attempt to do this, which is a version of expressivism that states that normative expressions have their (standard) semantic values in virtue of expressing non-cognitive mental states. Specifically, I claim that there is no way for an expressivist theory to give a satisfactory explanation of in virtue of what normative predicates have their extensions.

My final paper, “Referential Normativity and the Metaphysical Commitments of

Normative Realism”, is linked to my second one by the other thread: the question of whether non-cognitive mental states play a central role in grounding the content of normative language and thought. The issue examined in this paper is whether every concept with a certain normative role, which connects uses of that concept to non-cognitive mental states, must refer to the same thing, and if so whether this fact could result in a version of normative realism that does without controversial metaphysical commitments to irreducible normative properties or objects. I argue that the realist motivations that lead theorists to posit this explanatory role for non-cognitive mental states could not be satisfied by a theory that avoids controversial metaphysical commitments.

BIOGRAPHICAL SKETCH

David Fielding grew up in College Park, Maryland. He graduated from Reed College in 2011 with a B.A. in Philosophy and Mathematics before entering the Cornell Philosophy Ph.D. program the following Fall.

ACKNOWLEDGEMENTS

In addition to the more direct contributions indicated below, this dissertation grew from philosophical dispositions and abilities that developed under the influence of too many people to name. From my undergraduate major in philosophy at Reed College, the influences that stick out most clearly in my mind are my professors Mark Hinchliff, Mark Bedau, and William Taschek, as well as my fellow major Andrew Ho. At Cornell, the professors and students that stand out the most are those I interacted with most in my early-year courses, reading groups, and informal discussions. Those professors are Will Starr, Nico Silins, Matti Eklund, Ted Sider, Sarah Murray, Harold Hodes, Nicholas Sturgeon, Julia North, Derk Pereboom, Tad Brennan, Andrew Chignell, and Gail Fine. And the students are Eric Epstein, Eric Rowe, David Kovacs, Philippe Lemoine, Thomas Foerster, Fran Fairbairn, Chad McIntosh, Lucia Munguia, Avi Appel, and especially my cohort mates: Yuna Won, Brandon Conley, Katie Mathie Smith, Kim Brewer, and Marta Heckel.

Comments from my committee—Will Starr, Harold Hodes, Sarah Murray, and Matti Eklund—played an essential role in the development of all three essays included here. In addition to those comments, the first paper benefited from comments at two Cornell graduate student workshops and one meeting of the Cornell semantics reading group. This paper originally grew out of a year worth of discussions on issues related to context sensitivity with Will Starr. The second paper benefited from comments at the Cornell graduate student workshop, the Cornell dissertation writing group, and the placement seminar, especially those from Julia Markovits, who ran the seminar. This paper developed first from an expressivism reading group run by Yuna Won and then further over the course of weekly discussions of Michael Ridge's work with both her and Will Starr. The third paper developed from a reading group on Matti Eklund's book *Choosing Normative Concepts*, whose participants were Yuna Won, Eric Epstein,

and Thomas Foerster. Discussions of Williams' "Normative Reference Magnets" in a reading group run by Vivek Mathew also helped me arrive at the ideas within it.

My dissertation also couldn't have come to light without the constant help of Pamela Hanna and Dorothy Vanderbilt in guiding me through the torrent of paperwork that its writing entailed.

And while writing this dissertation, my life was non-philosophically enriched by too many people to name, including by most of those mentioned above. This category also includes Ashley Schneider, Andrew Ho, Joe Wasserman, Martín Moore, Alex Esposito, Rob Wesley, Ben Sales, my parents Ellen and Richard Fielding, my siblings Peter, Mariaea, and Catherine Fielding, and my nephew Matthew Johnson.

I am deeply thankful to everyone mentioned above, as well as to everyone that I have mistakenly failed to mention.

TABLE OF CONTENTS

Biographical Sketch	iii
Acknowledgments	iv
Chapter 1: Doing without Minimal Propositions	1
Chapter 2: Problems for Metasemantic Expressivism	40
Chapter 3: Referential Normativity and the Metaphysical Commitments of Normative Realism	68
References	101

CHAPTER 1

DOING WITHOUT MINIMAL PROPOSITIONS

1. Introduction

A central question in philosophy of language concerns how much of communicated meaning comes from the conventional meaning of the sentence uttered versus how much comes from context. Imagine that purposeless dust storms carve a string of shapes into a rock with the appearance: 'That is a spaceship.' If we treat this string as being an instance of the English sentence it looks like, then we will view it as having a meaning, but there is a sense in which that meaning will be incomplete. The sentence will not express a proposition—a claim about the world capable of being either true or false—because in it the word 'that' fails to refer to an object. On more prosaic occasions, the intentions of speakers to talk about a specific object and facts about what objects are salient in a conversation will fix a referent for 'that', but no such reference-fixers are available here. We can conclude that sentences like 'That is a spaceship' fail to express propositions without the help of context.

A question immediately arises of how widespread this phenomenon is. Do only sentences with obviously context-sensitive words like 'that' or 'I' need a context of use to determine the proposition they express? In more subtle ways, many other terms can be used to communicate different things in different contexts. Depending on what activity is salient, Brian will need to be prepared for different things to make an utterance of 'Brian is ready' true. But that fact doesn't on its own entail that 'Brian is ready' fails to express any proposition without the help of context. Perhaps, in addition to the the more specific meanings it conveys in conversation, 'Brian is ready' has a context-invariant linguistic meaning such that it is true just in case Brian is ready for any

activity whatsoever. This is the view of so-called “Semantic Minimalists”, who claim that—aside from sentences containing a small list of clearly context-sensitive expressions like 'that'—every sentence expresses a context-invariant “minimal proposition”. Opposing the Minimalists are *Linguistic Pragmatists*.

For a given expression, the definitional claims of Minimalism and Pragmatism conflict in a straightforward way. According to a Minimalist account of the expression, sentences containing it express a context-invariant proposition known as a *minimal proposition*.¹ According to a Pragmatist one, speakers can only express propositions when uttering sentences involving the expression through the help of context. This paper will examine this debate, focusing on the arguments of the Minimalists, Herman Cappelen and Ernie Lepore. I will conclude that there is much less to distinguish the views of Pragmatists from Cappelen and Lepore's than they suppose, and that the arguments that Cappelen and Lepore level against Pragmatists do not succeed.

I will proceed as follows. In section two, I distinguish the notions of semantic value and of what is said in uttering a sentence before going on to give examples of contextual variation in the latter. In section three, I discuss the three main accounts of that variation in the literature—Minimalism, Pragmatism, and Contextualism—explaining how these views differ in terms of the concepts already introduced. Section four then discusses each of the arguments advanced against Pragmatism by Cappelen and Lepore, finding that they all fail. This will involve examining the cognitive role that Cappelen and Lepore posit for minimal propositions. I find that their arguments concerning the role of minimal propositions fail to take into account the role that

¹ Minimalists are not committed to the view that the sentences 'Go to the store.' or 'Where are we?' express propositions. Throughout this paper, references to “sentences” should be read as limited to well formed declarative sentences. Also, as will be made clear in section III, the sentence will not necessarily be completely context-invariant if it contains expressions from the Basic Set of Context Sensitive Expressions in addition to the expression under discussion.

sentence meaning plays for Pragmatists and in particular the similarity of Pragmatist sentence meaning to Minimalist sentence meaning.

2. Contextual Effects on What Is Said

Like other natural languages, English has many expressions whose meanings vary depending on the context in which they are uttered and do so in a way that's clear even prior to linguistic theorizing. 'I', of course, refers to the speaker, and different contexts have different speakers. It's a vague matter exactly which expressions are self-evidently context-sensitive in this manner, but the following list described by Cappelen and Lepore (and all found earlier in Kaplan (1989)) captures the idea well:

The personal pronouns 'I,' 'you,' 'he,' 'she,' 'it' in their various cases and number (e.g., singular, plural, nominative, accusative, genitive forms), the demonstrative pronouns 'that' and 'this' in their various cases and number, the adverbs 'here,' 'there,' 'now,' 'today,' 'yesterday,' 'tomorrow,' 'ago' (as in 'He left two days ago'), 'hence(forth)' (as in 'There will be no talking henceforth'), and the adjectives 'actual' and 'present'...Words and aspects of words that indicate tense also have their reference [determined by context]. And there are also contextuials, which include common nouns like 'enemy,' 'outsider,' 'foreigner,' 'alien,' 'immigrant,' 'friend,' and 'native'. (Cappelen and Lepore 2005, 1).²

They do not offer a perfectly precise account of which expressions will end up in this list, but in addition to generalizing from the above examples they later provide three non-definitive tests for context-sensitivity. These will be discussed later in the section dealing with their use of these tests as arguments for their view. I will follow Cappelen and Lepore in referring to these

2 They express some doubts about the inclusion of contextuials in a footnote to the page and do not again bring them up as examples.

expressions as the “Basic Set of Context Sensitive Expressions.”

It's worth pausing to clarify a few things about the nature of these context sensitive expressions. First, they are semantically context sensitive, meaning that the semantic content of each expression can differ in different contexts.³ Semantic value and content are technical notions, not simply synonymous with the pretheoretical notion of meaning. And it's open, at least prior to such theorizing, that there are ways in which meaning can vary depending on context other than *via* semantically context sensitive expressions. In particular, it's open that an expression can be used in a single sentence to say different things in different contexts without that expression being context sensitive in the relevant sense.

More positively, the semantic value of an expression will be determined by the explanatory role that assigning that value to that expression plays in a larger semantic theory. Semantic theories are designed to explain a number of linguistic facts, such as entailment and consistency relations between sentences, and they do so by specifying semantic values for individual words and rules for combining the semantic values of words depending on their arrangement in a sentence. The semantic value of such a word will be whatever value is assigned by the semantic theory that best explains the linguistic facts such theories are meant to explain. In turn, the semantic value of a sentence will be what results from the process of applying the combination rules of the semantic theory to semantic values of the sentence's parts.

Seth Yalcin (2014) offers a much more detailed account of semantic values along these lines, but further details are not needed now. The above sketch already emphasizes what matters to the paper, which is that semantic values are theoretical entities, and we cannot tell prior to theorizing exactly what relation they will have to more everyday linguistic notions. Everyone has

3 Throughout the paper 'context sensitive' will mean semantically context sensitive unless explicitly stated otherwise.

some idea of what a person said as opposed to merely implied when they uttered a sentence, as shown by how this difference seems to affect judgments about whether someone lied or merely misled (Mahon, 2016). Some theorists have identified what a person said with the semantic value of the sentence they uttered relative to context.⁴ We cannot tell by mere intuition whether this identification is correct.

Although what is said is conceptually distinct from semantic value relative to context, the most straightforward way to argue that expressions from the Basic Set are really context sensitive is to point out that they can be used to say different things in different contexts. When I utter 'I am hungry' I say something different than when you utter the same sentence. We express different propositions with different truth-conditions. Clearly such straightforward facts play some role in our treating 'I' as semantically context sensitive. But this method of inferring context sensitivity from variation in what is said threatens to greatly expand the set of context sensitive expressions because it turns out that an extremely wide range of expressions give rise to variation in what is said. What follows is a small selection of these types of expressions. All the expression types are discussed by Cappelen and Lepore (2005, 17-39), though the examples given for them differ.

Quantifier expressions are sometimes argued to have different domains of quantification depending on the context in which sentences containing them are uttered:

(1) No student paid attention to the lecture.

If a professor uttered (1) directly after teaching a class, she would ordinarily be taken to be saying that no student in the class she just taught paid attention to her lecture.⁵ However,

4 The semantic value of a sentence relative to context is determined in the same way as the semantic value of the sentence, with the addition that the meaning of all context sensitive expressions in the sentence has been supplied by the context of the utterance.

5 In specifying what was said, we have also let context fix the meaning of the definite description, 'the lecture', singling out the utterer's lecture. But one's view on the manner in which this meaning has been affected will

different utterances of (1) are capable of having a different domain for 'no students'. A different professor saying (1) after teaching his own class would generally be taken to assert something whose truth depends in no way on the first professor's class or the attentiveness of the students therein.

Uses of gradable adjectives (which admit of comparative and superlative forms) are often claimed to differ regarding a relevant comparison class for possession of some trait or degree of having a given property:

(2) That truck is huge.

A compact car driver on a highway filled with cars dwarfed by a large pickup would normally be taken to have said something true by uttering (2). But applications of 'huge' will likely be restricted to a class of cars larger than that pickup if (2) is uttered at a monster truck rally.

Possessive constructions are often taken to contribute different relations between the possessor and the possessed depending on the contexts in which sentences containing them are uttered:

(3) Alice's pencil needs to be sharpened.

In many contexts it would be natural to take an utterance of (3) to say something that is true just in case the pencil that is owned by Alice needs to be sharpened. But the relation between Alice and the pencil need not be ownership in every utterance of (3). If Alice borrowed a pencil from someone earlier, and this is common knowledge, then an utterance of (3) will say something that is true just in case the pencil borrowed by Alice needs to be sharpened. Indeed, almost any salient relation between two things can be the relation specified by a possessive construction involving them.

depend on one's view of the meaning of definite descriptions.

Moving more quickly, each of the following sentences (and others containing the relevant expression from each) has been argued to be context sensitive, where each expression has been discussed by Bach among others (1994):

- (4) The Labrador is ready [for what event?].
- (5) A cheetah would be fast enough [for what purpose?].
- (6) My burrito is too juicy [for what purpose?].

In each case, the bracketed question specifies a further piece of information that is typically supplied by context. Different contexts will generally result in different answers to these questions being contributed to what is said by the utterance of each sentence. To expand on just one of these, you would expect (4) to require different things of the relevant Labrador if it were uttered at a dog show than if it were uttered in a house prior to going out for a walk.

The examples above have become somewhat commonplace in philosophy of language. Anne Bezuidenhout provides a perhaps more extreme or unusual example of contextual variation in what is said:

We're at a county fair picking through a barrel of assorted apples. My son says 'Here's a red one,' and what he says is true if the apple is indeed red. But what counts as being red in this context? For apples, being red generally means having a red skin, which is different from what we normally mean by calling a watermelon, or a leaf, or a star, or hair, red. But even when it is an apple that is in question, other understandings of what it is to call it 'red' are possible, given suitable circumstances. For instance, suppose now that we're sorting through a barrel of apples to find those that have been afflicted with a horrible fungal disease. This fungus grows out from the core and stains the flesh of the apple red. My son slices each apple open and puts the good ones in a cooking pot. The

bad one he hands to me. Cutting open an apple he remarks: ‘Here’s a red one.’ What he says is true if the apple has red flesh, even if it also happens to be a Granny Smith apple.

(Bezuidenhout 2002, p. 107)

This example would of course generalize to other color terms. This would suggest that which part of an object must appear to be a certain color under which circumstances for a color term to be said to apply to it varies depending on the context.

It will be helpful to have a term for expressions like the above that can be used to say different things in different contexts, and “context sensitive” won’t do since we are not prejudging whether the above are semantically context sensitive. Instead, we will call such expressions *variable*. All expressions in the Basic Set are variable, and the above examples represent a small selection of all the expressions outside of the Basic Set that have been argued to be variable as well. For ease of reference, we will use “The Extended Set” to refer to the set of variable expressions not in the Basic set. We will discuss how different theorists analyze expressions in the Extended Set in the next section.

3. Accounts of Variation in What Is Said

This section will cover the three most prominent ways to account for the examples from the previous section, which I will call the Contextualist, Pragmatist, and Minimalist approaches respectively.⁶ It’s worth noting at the outset that a given theorist need not adopt the same approach in all cases. For instance, one can be a Contextualist about ‘ready’ but a Pragmatist about gradable adjectives.

The Contextualist attempts to explain variability in terms of semantic context sensitivity.

⁶ I draw the ‘Pragmatist’ label from Neale, who calls this approach “Linguistic Pragmatism” when applied generally (2004).

For contrast, I'll describe what the Contextualist approach looks like when it differs most sharply from both the Pragmatist and Minimalist ones. This approach is to posit unpronounced, context sensitive expressions in the logical form of the sentence containing the relevant expression and use these to account for how the sentence can be used to say different things in different contexts.⁷ For instance, they might claim that the correct logical form of (4) is:

(4LF1) [S [DP The Labrador] [VP is ready [PP x]]]

Here I've left out irrelevant syntactic detail. The 'x' is to be filled in with the activity or event that is being discussed in the context of utterance, say, a dog show. The semantic value of 'ready' will, roughly, be a function from activities and objects to truth values, namely the characteristic function of the set of ordered pairs in which the second member of the ordered pair is ready for the event that is the first member of the ordered pair.⁸ Thus, when we fill in the semantic value of the context sensitive items in the sentence (including the unpronounced context sensitive one) and apply the semantic composition rules, we can derive the truth-conditions of the sentence. In a context where the relevant activity is a dog show, the sentence will be true if and only if the Labrador is ready for a dog show.⁹ Because of this, for the Contextualist the semantic value of this sentence is a proposition and that proposition is identical with what the speaker said by uttering the sentence.

The details of the implementation of the Contextualist approach will change depending

7 The logical form of a sentence is a representation of its syntax that contains everything relevant to semantic interpretation of it. What the correct logical form of a sentence is will depend on the explanatory goals of semantics and syntax just as the correct semantic value of an expression does.

8 To allow for simple rules combining semantic values, 'ready' would actually be assigned a function from events to a function from objects to truth values. Further complications might be added depending on the scope of the linguistic phenomena we are explaining.

9 It can be argued that this is still not enough to give us truth-conditions because there are different ways in which you can be ready for a dog show—say, being ready to observe a dogshow versus being ready to participate in one—and we need to specify which to arrive at truth-conditions. If so, the Contextualist approach would build in further variables, or build in more information to the variable on offer.

on the expression involved.¹⁰ For quantifier sentences, they would posit an unpronounced context sensitive item that filled in the class of objects in the domain of the quantifier for that sentence (Stanley and Zsabo 2000). What would stay the same in each case is that the Contextualist posits a more complicated logical form for the sentence than it at first appears to have, resulting in the sentence having a proposition as its semantic value relative to context. This proposition will be what is said by uttering the sentence in that context.

Pragmatists and Minimalists, on the other hand, both posit a logical form for sentences involving the relevant expression that more closely matches how the sentence is spoken. For (4), this form would be something like:

(4LF2) [S [DP The Labrador] [VP is ready]]

The major difference with (4LF1) here, of course, is that we have dropped the context sensitive expression specifying what the Labrador is ready for. Where the Pragmatist and Minimalist come apart is in their assignment of a semantic value to 'ready'. The Pragmatist assigns 'ready' the same semantic value as the Contextualist, meaning that the semantic value of the whole sentence is in some sense incomplete for them. The semantic value of 'ready' is a function of both the object that is said to be ready and also the activity it is said to be ready for, but one of the two arguments of this function is missing. Because of this, the semantic value of the sentence as a whole is not a proposition for the Pragmatist. The sentence needs to have an activity specified for 'ready' in order to have determinate truth conditions, but nothing in the logical form does this specifying.¹¹ Kent Bach (1994) aptly calls the resulting proposition-with-a-missing-piece a *propositional radical*. This type of situation will happen in general for the Pragmatist. Any

10 And, as I mentioned above, I am only discussing here the version of Contextualism that differs most clearly from Pragmatism and Minimalism. Instead of adding unpronounced terms to the Logical Form, Contextualists sometimes posit that pronounced terms in the sentence that aren't clearly context sensitive in fact are.

11 For examples of the Pragmatist approach for various expressions, see Francois Recanati (2004), Stephen Neale (2004), and relevance theorists such as Robyn Carston (2002) and Dan Sperber and Deirdre Wilson (1986).

sentence which contains an Extended Set expression will be incomplete in the above manner and thus not semantically express a proposition. The Pragmatist will still say that what the speaker said by uttering the sentence is a proposition, though, because they do not identify what the speaker said with the semantic value of the sentence. What a speaker says by uttering a sentence is a matter of pragmatics for them.

The Minimalist, on the other hand, gives a different interpretation to the semantic value of 'ready'. For them the semantic value of 'ready' will be a function with only one argument, the characteristic function of the set of ready things. Because of this, the semantic value of sentence (4) will still be a proposition according to the Minimalist, since we are no longer missing an argument for the semantic value of 'ready'. And this proposition will not depend on context to fill in what activity the Labrador is said to be ready for. That said, the minimalist does not identify this semantic value with what is said by uttering the sentence.¹² The speaker of (4) said something about readiness to perform in a dog show, but the proposition semantically expressed by the sentence they uttered has nothing to do with dog shows in particular. We will call the propositions that minimalists think are semantically expressed by declarative sentences *minimal propositions*. These are minimal in that context only affects the sentences expressing them in so far as it affects expressions in them from the Basic Set. In all other respects, the proposition expressed is context-invariant.

You will likely want to know more about the nature of these minimal propositions, such as what the truth-conditions of (4) would be for a minimalist. Cappelen and Lepore want to say very little about this and in fact to argue that it is not the job of semantics to say more. Consider

¹² Things get somewhat complicated here by the fact that the main minimalists we are discussing, Cappelen and Lepore, are “speech act pluralists” who think that many things are said by every utterance, and one of these many things will in fact be the semantic value of the sentence relative to context (2005, 4). However, they do not think this latter proposition is the most salient of the things a speaker says by uttering a sentence. Usually this will be a more specific proposition, such as one about dog shows in the case of sentence (4).

'ready' (2005, 155-157). Cappelen and Lepore will claim that 'The Labrador is ready' expresses the proposition that the Labrador is ready, a proposition which is true if and only if the Labrador is ready. However, at other times Cappelen and Lepore are somewhat more willing to speculate about the what minimal propositions require of the world to be true (2005, 160-162, 185-186). Here is what they say about 'ready' when in this speculative mode:

Think about what people who are ready have in common. To make this vivid, imagine A's being ready to commit a bank robbery, B's being ready to eat dinner, and C's being ready to take an exam...It's not just a pun that we feel comfortable describing them all as being ready. They really are all *ready*. That's different from their all being *done with* the tasks, or *excited about* them, or *prepared for* them, or *good at* them, etc. (2005, 166-167).

This passage suggests that the minimal proposition that, say, Sam is ready is true just in case Sam is ready for something or other. The claim is different from the claim that there is elision of an initial existential quantifier at the phonological level which is present in logical form. Such an addition would violate Minimalism's general distaste for positing items in logical form based only on semantic considerations. Rather, the claim is that the semantic value of 'ready' applies to anything that is ready for anything, with no restriction on the domain of the quantifiers.

This account can be extended naturally to other posited cases of minimal propositions for sentences containing Extended Set expressions. The minimal proposition that something is strong enough is true just in case that thing is strong enough for something or other. The minimal proposition that something is heavy is true just in case it is heavy by some standard or other. The minimal proposition that something is blue is true just in case some part of it is blue in some way or other under some condition. It's worth mentioning that these propositions are quite strange, in that they are typically obviously true (and their negations obviously false); surely just about

everything is ready for something or other. And even if some things do not appear any shade of blue in any condition, presumably the semantic value of 'blue' will still be true of many more things than 'blue' would ordinarily be applied to in a given context.

When discussing the role that minimal propositions play in communication I will treat the account of minimal propositions sketched above as Cappelen and Lepore's view. This treatment may seem unfair to them, as they are quite clear in refraining to officially endorse it. That said, many of Cappelen and Lepore's arguments rely on claims about what facts can be explained by minimal propositions for which other views cannot account. This will become particularly clear when we look in section 4.2 at their argument that minimal propositions are needed to explain communication across contexts. Arguments such as this can only be assessed if we have some account of what minimal propositions require of the world to be true. Since the above suggestion is the fullest account about which Cappelen and Lepore even speculate, I will use it when assessing such arguments.

The most important thing to take away from this discussion of approaches to variation in what is said is the great similarity between the pragmatist and minimalist approaches. They associate the same logical forms with sentences. They also think that what a speaker says in uttering a sentence is a matter of pragmatics; it is not identified with the semantic value of the sentence uttered even relative to context. They differ only in that the minimalist thinks the semantic value of the relevant sentence relative to context is a minimal proposition, whereas the pragmatist thinks it is a propositional radical. Cappelen and Lepore generally present themselves as disagreeing more greatly with their opponents than this. We will see in the next section how this misperception leads their arguments to fail against the pragmatist, undermining the case for minimal propositions. We will return in the final section to the question of whether the

Contextualist can also escape their arguments.

4. Cappelen and Lepore's Arguments Against Pragmatism

Cappelen and Lepore advance three types of argument against Pragmatists. The first involves their deployment of three tests meant to provide evidence for when an expression is context sensitive. They claim that these tests are failed by the expressions in the Extended Set. I will discuss Cappelen and Lepore's account of each test and its application. Then I will address their second argument, that Pragmatism cannot account for the success of communication across extremely dissimilar contexts. This will be followed by an examination of the various ways in Cappelen and Lepore claim that Pragmatism is self-contradictory.

4.1.1. The Inter-Contextual Disquotational Indirect Reports Test

Cappelen and Lepore's first test of context sensitivity concerns whether the putatively context sensitive term blocks "Inter-Contextual Disquotational Indirect Reports", which Cappelen and Lepore explain as follows:

Take an utterance u of a sentence S by speaker A in context C . An Inter-Contextual Disquotational Indirect Report of u is an utterance u' in a context C' [where C does not equal C'] of 'A said that S .' (2005, 88)

So, the test involves indirect speech reports that are inter-contextual in that the report's context differs from the reported act's context. The reports are also disquotational in that the same sentence that was used in the speech act being reported is again used in the complement clause of the report of that speech act. They then characterize their test as follows:

Suppose you suspect, or at least want to ascertain whether, e is context sensitive. Take an

utterance *u* of a sentence *S* containing *e* in context *C*. Let *C'* be a context relevantly different from *C* (i.e., different according to the standards significant according to contextualists about *e*). If there's a true disquotational* indirect report of *u* in *C'*, then that's evidence *S* is context insensitive. (To be 'disquotational*' just means you can adjust the semantic values of components of *S* that are generally recognized as context sensitive, i.e., we just test for the controversial components.) (2005, 89).¹³

A natural interpretation of the reasoning behind the Inter-Contextual Disquotational Indirect Reports test is as follows. Since Pragmatists think that context can influence the meaning of an utterance of a given sentence in a number of ways, and since there is no reason why the context of the speech report must be the same as the context of the speech act being reported, we should expect there to be situations in which it is infelicitous to report a speech act using the exact same sentence as the original. This would be because the sentence might take on a different meaning in the context of the speech report than it had in the original speech situation. This change in meaning could result in the report being false because something other than what the original speaker communicated was attributed to the speaker by the report.

A number of commentators (including Szabo 2006; Travis 2006; Bezuidenhout 2006) claimed that the context sensitive expressions they posited passed this test, *contra* Cappelen and Lepore. It is easy to find examples that seem to show that the expressions in the Extended Set cannot always be felicitously reported disquotationally. Imagine that Jason yesterday told his parents that his sister was ready for soccer practice by uttering, "Lisa is ready". Today Jason's mom asks whether Lisa is ready for her chemistry final and is told, "Jason said that Lisa is ready", with the reporter knowing full well of the difference in context. This report seems false

¹³ Note that they use "contextualists" here to refer to both what I call Pragmatists and Contextualists.

every bit as clearly as a report, “David said that that is his favorite movie” while demonstrating the DVD case of a completely different movie than David had demonstrated the previous day when uttering, “That is my favorite movie.” This would suggest that at least one Extended Set expression passes Cappelen and Lepore's first test.

However, in a reply to critics Cappelen and Lepore claim that these commentators have misinterpreted their test (2006, 52-54). They claim that, contrary to how these commentators dealt with the test, passing the test is not guaranteed by a single instance of a sentence being falsely reported by a disquotational report. Instead, passing the test requires that there not be any disquotational indirect speech report that is true even though the report and the reported sentence were uttered without relevant shared context. Cappelen and Lepore claim the existence of any of such report indicates that there is a minimal content for the relevant sentence which is shared in every context. Because the sentence expresses this minimal content in every context, the reporter can truthfully attribute that content to the speaker of the reported sentence.

On closer inspection, this interpretation of the test makes more sense for Cappelen and Lepore's purposes than the interpretation used by their critics. The first interpretation of the test relied only on the claim that utterances of the relevant sentence could say different things in different contexts to generate the falsity of the relevant disquotational report. But Cappelen and Lepore agree that the sentences under discussion contain variable expressions, so their own view commits them to the claim that the sentences will pass the first interpretation of their test. The interpretation they clarify in their response to critics makes more sense of how the test stems from the difference between Minimalism and Pragmatism.

That said, a good case can be made that the Cappelen and Lepore cannot use this test to show that the reported sentences express minimal propositions. Recall that passing the test

requires that there not be any true disquotational indirect speech reports that do not share the relevant context. What that relevant context is will vary based on the manner in which context is claimed to affect the meaning of sentences involving the expression being discussed. For sentences containing 'I', what needs to stay fixed between contexts to guarantee true reports is the speaker. For ones containing 'that', the demonstrated object must remain fixed.

What about sentences containing 'ready', but without a specified prepositional phrase complement for it? Pragmatists take the content of utterances of these sentences to include contextually supplied information about what the person is ready for (and the manner in which they are ready for it and so forth). So to assess whether 'ready' fails Cappelen and Lepore's first test, we have to look at contexts in which the original speech act and the report of it differ regarding whatever facts determine what the person is said to be ready for. This would mean that in the speech report either the subject is being claimed to be ready for something different than she was originally claimed to be ready for (as in the soccer/exam example above), or that only one of the two contexts fails to supply anything for the person to be ready for. We've already seen that reports of the former type are inappropriate; a disquotational report of a sentence with 'ready' can be false if the reporter intends to use the context of the report to mislead people about what the reported speaker meant. We could expect things to be similar for other expressions in the Extended Set.

In light of this, let's now examine speech reports in which the reporters are ignorant of features of the context of the reported speech act that are potentially relevant to that act's meaning. To help imagine such a situation, it might be useful to think of the reporting context as consisting of spies discussing recordings from a bug. For each case I will give the reported sentence, the report, and the missing potentially relevant context. It is common knowledge

among reporter and addressee that they lack the potentially relevant information. Assume that they know all features of context that might be relevant to the meaning of the reported act aside from the information explicitly highlighted as absent.

For the first two sets of sentences, the reporters don't know the location designated by 'here' or the object picked out by 'this'.

(S1) Tom likes it here.

(R1) Putin said that Tom likes it here.

(S2) This is Michael's favorite cat.

(R2) Putin said that this is Michael's favorite cat.

My assessment of these sentences is in line with what Cappelen and Lepore would predict. Neither (R1) nor (R2) seem to be true reports of (S1) and (S2). In the case of (R1), 'here' appears to refer to wherever the report took place, a place that Putin presumably said nothing about. (R2) seems to lack a truth value since it contains an element, 'this', that needs to be assigned a value by context but which hasn't received one.

For the next four sets of sentences, the reporters don't know the relevant activity for 'ready', the relevant comparison class for 'tall', the relevant relation picked out by the possessive, or the relevant manner of greenness picked out by 'green'.¹⁴

(S3) Tom is ready.

(R3) Putin said that Tom is ready.

(S4) Michael is tall.

¹⁴ By "relevant manner of greenness" I'm referring to the putative variations in what can be said by sentences involving color words described by Bezuidenhout, which I outlined earlier.

(R4) Putin said that Michael is tall.

(S5) Tom is holding Michael's book.

(R5) Putin said that Tom is holding Michael's book.

(S6) The apple is green.

(R6) Putin said that the apple is green.

Again, I find myself agreeing with what one would expect Cappelen and Lepore to say about these examples. Despite the reporters being ignorant of parts of context that both Pragmatists and Minimalist think influence what is said by the reported speech act, each of the speech reports seems true. The Minimalist account of this situation is that the reporters are reporting the minimal proposition that Putin asserted in addition to whatever other propositions he asserted in the context. The sentences that come next cast doubt on this account.

For the final two sets of sentences, the reporters don't know how Putin intended the ambiguities to be resolved. In the first they do not know whether 'bat' denotes a pet animal or a club-like object. In the second, they do not know whether Putin claimed Michael used binoculars to see a spy or if he claimed that Michael saw a spy who had binoculars.

(S7) Tom brought his bat to show and tell.

(R7) Putin said that Tom brought his bat to show and tell.

(S8) Michael saw the spy with the binoculars.

(R8) Putin said that Michael saw the spy with the binoculars.

Reports (R7) and (R8) sound every bit as appropriate as (R3)-(R6).¹⁵ In fact, all six of the reports that do not involve ignorance of the value of an item from the Basic Set seem true. But Cappelen and Lepore's account of the truth of (R3)-(R6) cannot be used to explain the truth of (R7) and (R8). According to Cappelen and Lepore's own view, in the last two cases the reporters are ignorant of something relevant to determining the minimal proposition expressed by the reported sentences, just like in cases (R1) and (R2). From the above we can conclude that we cannot explain when disquotational indirect reports are true by determining whether the reported sentence semantically expresses a minimal proposition in the context of the report.

Now both the Minimalist and the Pragmatist are left without an explanation for the difference between (R1)-(R2) and (R3)-(R8). But this situation should not worry the Pragmatist now that we have eliminated the Minimalist's account of the difference. A central feature of Pragmatism is that much context sensitivity is not triggered by items in logical form, whereas the context sensitive expressions from the Basic Set will be present in the logical form. No part of Pragmatism requires that these two kinds of context sensitivity have all the same effects in different linguistic contexts. In particular, it is not in itself an objection to Pragmatism if only the latter kind of context sensitive expression requires true disquotational indirect reports of sentences containing them to retain the relevant context of the reported speech act. This fact was only a problem for Pragmatists under the assumption that the best (and perhaps only) explanation of it required appeal to minimal propositions. Since that explanation fails, the Minimalist has not given the Pragmatist a reason to worry.

15 For what it's worth, I showed all eight pairs of sentences to nine linguistics and philosophy graduate students at a semantics reading group. They were told about the missing potentially relevant information in each case but were not told my own view on the sentences or to what kind of test they might be relevant. Except for one linguist who abstained from answering any of the questions due to not understanding what 'true' meant in this context, there was complete agreement with my own view about sentences (R1)-(R6). This also held for (R7) and (R8) except that one student was drawn to the view that these reports are indeterminate though he didn't feel sure about this.

4.1.2. The Collective Descriptions Test

Cappelen and Lepore's second test for context sensitivity is really two tests; one for context sensitive verb phrases and another for context sensitive singular terms. For a context sensitive verb phrase, v , “on the basis of knowing that there are two contexts of utterance in which 'A v -s' and 'B v -s' are true respectively, we cannot automatically infer that there is a context in which ' v ' can be used to describe what A and B have both done” (Cappelen and Lepore 2005, 99). For a context sensitive singular term, N , we should not be able to infer from there being two contexts of utterance in which 'N is F' and 'N is G' are true respectively that there is a context in which we can truly utter 'N is F and G'. In each case we call the sentence which combines the earlier two a “collective description”.

As with the first context sensitivity test, there are two interpretations for how to use this test, one which doesn't rely on any of the differences between Pragmatism and Semantic Minimalism, and one that does. Consider the following sentences:

(7) John is ready.

(8) Joe is ready.

(9) John and Joe are ready.

According to the first interpretation of the test, a supposedly context sensitive verb phrase like, 'is ready', should be such that there can be two contexts in which utterances of (7) and (8) are true even though in another context an utterance of (9) is false. Clearly there are such contexts. Just have the speaker of (9) claim John and Joe to be ready for something that they've never been ready for (such as running for president), whereas the speaker(s) for (7) and (8) truly attributed readiness for something more prosaic to each (such as eating dinner). This interpretation of the test is bad for Minimalists not just because it clearly delivers a result they disagree with, but also

because their own theory predicts that we should see these results. Again, all that 'is ready' passing the test relied on was that utterances containing it can be used to say different things in different contexts, which the Minimalist accepts.

So we instead should look for an interpretation of the test that relies on the difference between Minimalism and Pragmatism, the existence of minimal propositions. On this reading of the test, minimal propositions are supposed to explain why one can truly utter (9) whenever one knows that both (7) and (8) can be truly uttered. We have already seen that this is not the case if the activity associated with 'ready' differs significantly in the context of (9) from that of the other two. Instead we will look at situations in which the contexts of each sentence differ but not in such a way that the speaker of the collective description is claiming something entirely unrelated to what was claimed in the utterances collected.

Suppose John is the closer for a World Series Major League Baseball team. A family hears an announcer utter 'John is ready' when he throws his last warm-up pitch before heading to the pitcher's mound. A member of the family around the same time utters 'Joe is ready', said of the youngest child to describe his being ready for bed. For whatever reason, one of them decides to add, 'John and Joe are ready'. The Minimalist accounts for the intuitive truth of (9) by positing that (7) and (8) both semantically express minimal propositions in which the same property (readiness) is ascribed to their subjects. This guarantees that there is a property that can be picked out by 'ready' which will apply to both Joe and John. But a Pragmatist can account for the facts about these sentences without using minimal propositions. There is a property that applies to both John and Joe, being ready for the salient task that each is about perform. If the speaker of (9) intends the sentence to mean that they are each ready for these respective tasks (or the single task of either pitching or going to bed), then it will be true.

More generally, it is easy to utter true collective description sentences containing Extended Set expressions because the speaker will generally be able use the relevant verb phrase in a way that it applies to both objects (similarly for singular terms). To apply this thought to another example, a Pragmatist should have no problem accepting that 'are tall' can be used to mean 'are tall for the salient groups each belong to', accounting for why it is easy to utter true collective descriptions of sentences applying 'is tall' to different subjects. Unless Cappelen and Lepore have an unstated reason why these expressions cannot be used in this way, the Pragmatist seems perfectly capable of explaining the data concerning collective descriptions.

4.1.3. Real Context Shifting Argument Tests

When listing variable expression in Section 2, we gave what Cappelen and Lepore call Context Shifting Arguments (CSAs) for them. A CSA involves examining different utterances of a sentence made in different contexts to see whether or not the utterances say different things depending on the context. What Cappelen and Lepore call “Real Context Shifting Arguments” (RCSAs) are just like that, except that one of the contexts is that of the person presenting the RCSA, who goes on to use the sentence in their own context.¹⁶ More specifically, RCSAs are supposed to motivate an instance of the Inter-Contextual Disquotation schema:

(ICD) There are (or *can* be) false utterances of $\lceil S \rceil$ even though S. (Cappelen and Lepore 2005, 105, footnote removed).¹⁷

An example RCSA could go as follows.

¹⁶ Cappelen and Lepore call them “Real Context Shifting Arguments” because, unlike CSAs, they take RCSAs to provide good evidence of semantic context sensitivity (2005, 104-108).

¹⁷ In practice, Cappelen and Lepore don't require the ICD to be exact instances of the schema. Instead they just require them to involve contrasting the truth or falsity of one's own use of a sentence with that of another person whose utterance is mentioned.

Here. There isn't any water here because a drought has gone on for years. But years ago, when I lived in the rain forest, I remember saying, 'There is water here', which was true. So there have been some true utterances of 'There is water here' even though there isn't any water here.

The upshot of this RCSA would be that 'here' is semantically context sensitive.

It's fairly clear why one would expect variable terms to admit of RCSAs. If sentences containing the term can be used to say different things in different contexts, then in particular the writer's own use of the sentence within the RCSA could be used to say something different from what was said by the other utterance that is mentioned in the RCSA. In many cases, we should be able to fill out the contexts so that what is said by one of them is false even though what is said by the other is true.

Recall that we noticed two different interpretations for the previous two context sensitivity tests. In each case, one of the interpretations failed to make use of any of the differences between Minimalism and Pragmatism and hence the test could not be used to differentiate the two views. The other interpretation involved explaining some putative linguistic data by appealing to minimal propositions. If you reread the paragraph before this one, it should be clear that the RCSA test is of the former type. Motivating the test requires nothing but the claim that the relevant utterances can be used to say different things in different contexts, a claim that, yet again, both Minimalists and the Pragmatists accept. However, unlike with the previous two tests, I can think of no further way of interpreting the test on which it actually makes use of minimal propositions. That is, there doesn't appear to be any interpretation of the test according to which examining its application to different types of sentences results in an argument against Pragmatism. Because Pragmatists and Minimalists both believe that the sentences in question

can be used to say different things in different contexts, they should both expect all of the relevant sentences to admit of RCSAs.

What Cappelen and Lepore actually claim is that good RCSAs can be made for expressions in the Basic Set but not for the Extended Set. Whether or not they are right about this, this fact would be just as odd for them as it would be for their opponents. Perhaps Cappelen and Lepore's final point of caution in their section on RCSAs can help us here:

The tests in this chapter obviously have to focus on what's communicated; we are after all communicating with the reader, and so the tests are tests that go via communicated content. Our view...is that there is a sharp distinction between communicated content and semantic content. Nonetheless, we use communicated content (the content we succeed in communicating to our readers) to 'get at' semantic content (semantic context sensitivity in particular). There is of course no other way to proceed. The purpose of the tests is to generate contexts in which semantic content is salient. Think of the tests like this: They are ways to get the audience to notice semantic features of sentences uttered. They create contexts in which our attention is drawn to features of the semantic content expressed by the utterances in question. (2005, 113).

This response will not help. The problem for Minimalists is that, for the same reason as with the Pragmatists, their view seems to predict that expressions in the Extended Set should pass the RCSA test. Their reply is that when thinking about the test we are focused on the semantic content of (the minimal propositions expressed by) the sentences involved. But thinking about their examples of RCSAs for expressions outside of the Basic Set I was not aware of focusing any more than usual on semantic content. Rather, I was as usual attempting to determine the content that the authors intended to communicate to me. Cappelen and Lepore seem committed

to the view that the nature of RCSAs themselves focuses interpreters on semantic content, but they do not give an explicit argument for why RCSAs should do so.

In any case, as with the Indirect Disquotational Report test, the Pragmatist should not find it too surprising that sentences containing expressions from the Basic Set have different linguistic properties than those that do not. The burden of proof is on Cappelen and Lepore to demonstrate that any given difference in linguistic properties between Basic Set expressions and others is best explained through positing minimal propositions. Without a clearer account of how the RCSA test focuses readers on minimal propositions, Cappelen and Lepore have not met that burden in this case.

To summarize section four so far, we have now examined all three of Cappelen and Lepore's context sensitivity tests. All three of these tests had a natural interpretation that relied only on the relevant expressions being variable. But because Minimalists and Pragmatists agree about which expressions are variable, the tests failed to differentiate between Minimalism and Radical Pragmatism on these interpretations. The first two tests but not the third could additionally be construed in such a way that minimal propositions play an explanatory role in them, allowing the Minimalist to potentially give a better account of the linguistic facts highlighted by the test than the Pragmatist. However, in each of these cases it was found that the Pragmatist is in fact no worse off in terms of accounting for the relevant linguistic facts. I conclude that Cappelen and Lepore's context sensitivity tests do not provide an argument against Pragmatists. The final two subsections of section four will examine Cappelen and Lepore's remaining two arguments against Pragmatism.

4.2. Can Pragmatists Account for Shared Content Across Contexts?

Cappelen and Lepore's argument from shared content goes as follows (2005, 123-125). According to Pragmatism, a great many features of context are capable of being relevant to determining what is said by a given utterance, including knowledge of what's been said earlier, knowledge about the person speaking, knowledge of observable features of the speech environment, and general world knowledge. We are often capable of understanding the contents expressed by people even when we lack access to many of these potentially relevant features of content. This fact would be “miraculous” if Pragmatists were correct (2005, 123).

As usual, we must get more precise on how this argument makes use of a difference between Pragmatism and Minimalism, since Minimalists agree that figuring out what is said can require knowledge of all the features that Pragmatists cite. The answer is that minimal propositions play the role of communicated propositions of last resort for the Minimalist. Cappelen and Lepore claim that the “proposition semantically expressed is the content that the speaker can expect the audience to grasp...even if they have mistaken or incomplete communication-relevant information” and that it is “the content which can be grasped and expressed by someone who isn't even a participant in the context of utterance” (2005, 184-185). We are capable of understanding what someone said even when we fail to know some relevant feature of the context because one of the things they said by uttering the sentence is a context insensitive minimal proposition which we know is expressed by the sentence they uttered.

If you recall how strange minimal propositions are from section three, this claim about their role may seem odd. The minimal proposition that John is ready is true just in case one of infinitely many possible different events or activities is such that John is ready for it. Obviously this is the case. How can a listener be taken to understand what someone communicated if the

only proposition they can determine was said was something obviously true and that the speaker was not even consciously trying to communicate? Cappelen and Lepore attempt to answer this question in the case of an utterance of 'A is red':

[Knowing the minimal proposition expressed by 'A is red'] is a starting point. Suppose, for the sake of argument, that the proposition *that A is red* is trivially true...The audience knows that the speaker is talking about A and its redness, not, for example, about oysters, France, or Relevance Theory. There's a lot of stuff to talk about in the universe... It is trivial that A is red on some surface or other under some condition or other. The audience can assume that the speaker knew that this was trivial and was not interested in conveying such trivialities with his utterance and can, therefore, infer that there's work to be done in order to figure out exactly what the speaker was trying to communicate. In general, audiences know what to look for in such situations; they know what kind of information would help narrow down more closely what the speaker wanted to communicate (2005, 185-186).

The above represents a plausible enough account of what interpreters must do to figure out what something means assuming they lack relevant contextual information.

But Cappelen and Lepore don't stop to consider whether Pragmatists can offer a similar account without appealing to minimal propositions, and it seems that they can. An interpreter of 'A is red', according to Pragmatists, will still have access to an interpreted logical form of that sentence (that is, a logical form with the lexical entry for each expression in it specified). Surely this is enough to infer that the utterance was about A and its redness because 'A' and 'red' will both be parts of the logical form and the interpreter will know the semantic value of these terms. The interpreter will also have a sense of how these semantic values constrain what can be said by

uttering 'A is red'. In particular, this will involve knowing that, although you can say many different things by uttering 'A is red', any proposition you express by doing so will require of the world at least that A is red on some surface or other under some condition or other. This fact itself is not very informative and unlikely to be what the speaker intended to communicate, but the interpreter can use it as a start in the process of determining what the speaker so intended.

Cappelen and Lepore appear to be treating Pragmatists as though they don't have any notion of the context insensitive meaning of a sentence. But Pragmatists do have such a notion, merely one that assigns sub-propositional meanings to sentences. The above paragraph illustrates how even these meanings are capable of doing the work that Cappelen and Lepore's minimal propositions can do to explain understanding across contexts. More generally, we have been assuming, roughly, that the minimal proposition expressed by a sentence is true just in case one of the many things that can be said by uttering that sentence is true. Pragmatists have a different notion of the work done by the meaning of sentences (interpreted logical forms); they think that these meanings delimit what can be said by uttering the sentence. But since Minimalists and Pragmatists think that a similar range of things can be said by uttering a given sentence, it is clear that their respective notions of sentence meaning can do the same work. Knowing what a sentence means, for either theorist, will mean knowing something about what the sentence is about (such as A and redness) and also knowing that the sentence is being used to say one of the many things that that sentence can be used to say, as determined by the meaning of the sentence. This isn't very much to go on, but the important point is that same resources exist for both the Pragmatist and the Minimalist. From the utterance of a sentence, Pragmatism predicts that speakers will be able to infer that the speaker of the sentence believes the minimal proposition

that Cappelen and Lepore would assign to that sentence.¹⁸

It is worth adding that even according to Cappelen and Lepore's own view, it is plausible that knowing a sub-propositional context insensitive sentence meaning can provide as much information about what was communicated by an utterance as knowing the minimal proposition expressed by an utterance. To see this, let us consider two more cases of disquotational indirect reports of sentences spoken by Putin overheard by spies:

(10) She will make a good prime minister.

(11) The KGB is ready.

For (10), we are assuming that the reporter does not know the referent of 'she', and in (11) we are assuming that the reporters do not know what the KGB is ready for. From this we can conclude that according to Cappelen and Lepore, the spies would know the minimal proposition expressed by Putin's utterance in (11) but not in (10) (because 'she' is a semantically context sensitive word and we don't know its value in this context).

However, this difference in whether we know the proposition expressed would not correspond to a difference in the amount of information provided by knowing that Putin uttered each sentence. Consider the information provided by sentence (10). Even if you don't know the referent of 'she', that Putin uttered (10) tells you that Putin thinks that some particular woman would be a good prime minister (of somewhere). That some particular women would be a good prime minister of somewhere is itself a perfectly fine proposition and one we are able to conclude that Putin endorses from his utterance. The proposition that one can infer Putin endorses from his utterance of (10) is extremely unspecific, but so is the minimal proposition that Cappelen and Lepore might speculatively assign to (11). As discussed earlier, the minimal

¹⁸ This requires the assumption that the sentence was being used literally, but the Minimalist also needs this assumption for their description of how the listener determines what was said.

proposition expressed by (11) would be the claim that the KGB is ready in some way or other for something or other, which, surely, they are. That Putin believes this is even more trivial than that he believes the earlier unspecific proposition that some particular woman would be a good prime minister of somewhere. Many sexists have doubted that any woman could be a good prime minister of anywhere and even some that accept the possibility in the abstract might not have in mind any particular woman they think it is true of. On the other hand, anyone who understands the full range of things that a given person or agency can be said to be ready for can tell that it is obvious that there are some things that the KGB is ready for. Thus, knowing that Putin uttered (10) is much more informative than knowing that Putin uttered (11) even if Cappelen and Lepore are right in claiming that we know the proposition that the latter expresses but not the one that the former expresses. This fact is troubling for Cappelen and Lepore's argument that we need minimal propositions to account for the successful communication across contexts. If knowing the range of things that can be said by an utterance of (10) without knowing what was actually expressed by it can be more informative than knowing the minimal proposition expressed by (11), it seems plausible that sub-propositional meanings can more generally provide interpreters with as much information as minimal propositions.

4.3. Is Pragmatism Inconsistent?

Cappelen and Lepore's final argument against Pragmatism consists in three separate inconsistency charges that they level against the view. The previous discussion provides most of the information needed to dispense with these charges quickly. That said, before going forward I should point out that these charges are meant to apply to what Cappelen and Lepore call "Radical Contextualism", the view that some combination of Contextualism and Pragmatism apply to all

expressions in open linguistic classes. They argue that anyone tempted to one of these views about any expression outside of the Basic Set should ultimately be a Radical Contextualist in order to be consistent. For the sake of argument, I won't question that assumption in this section and will assume the Pragmatist is committed to Pragmatism about all relevant expressions throughout it.

The first inconsistency charge is that Pragmatism implies that some utterances of the sentences used to characterize Pragmatism will be false (Cappelen and Lepore 2005, 128-130).¹⁹ This claim is not defended by any specific context shifting arguments involving the terms used in the central claims of Pragmatism; it is merely offered as plausible in light of how widespread contextual influence on what is said is according to the Pragmatist. The charge is supposed to be a problem because the Pragmatist would need to provide an argument that “there is something special about” the supposedly true proposition they express by uttering these sentences as opposed to the false propositions they could otherwise utter with them (2005, 130). Specifically, the Pragmatist would have to establish that the propositions that they express with the relevant sentences are more important to the concerns of philosophers of language and semanticists than the other (false) propositions that can be expressed with them.

As usual, Cappelen and Lepore's claims appear at first to apply to Minimalism just as much as Pragmatism, as Minimalists are committed to as much variation in what can be said by uttering a given sentence as Pragmatists are. If Cappelen and Lepore are right, then presumably the sentences used to characterize Minimalism can be used to express different propositions that they would consider false. So, for the claim to actually target Pragmatism alone, we should

¹⁹ Cappelen and Lepore phrase the charge differently, as the objection that some utterances of 'RC[Radical Contextualism] is false' will be true. Their way of putting the argument doesn't appear to make it any stronger than it is as phrased in this paragraph..

interpret Cappelen and Lepore as claiming that the minimal propositions expressed by the sentences used to characterize Minimalism are all true, and that this protects them from the worry that there would be nothing special about what they say by uttering these sentences.

The problem here is that the previous claim assumes that the minimal propositions associated with the sentences used to define Minimalism are more relevant to the concerns of philosophers of language and semanticists than the other propositions that these sentences can be used to communicate, and Cappelen and Lepore have done nothing to defend this claim. Further, this claim would be surprising given that the minimal proposition expressed by a given sentence is generally not what is at issue in debates that use that sentence. For example, debates involving 'John is ready' typically involve figuring out whether John has practiced enough for a piano recital, not whether there is any activity at all he is ready for. Cappelen and Lepore fail to motivate why debates involving the sentences they use to state their views should be any different (that is, why we shouldn't expect one of the propositions other than the minimal proposition expressed by an utterance to be what is of the most interest). In light of this, if there is a demand on the Pragmatist to show why they are interested in the correct propositions, then there is an equal demand on the Minimalist.

Luckily for both views there does not seem to be any such demand. If someone wants to argue (A) that the sentences used to characterize a view can be used to express different views and (B) that those other views are more theoretically interesting than the ones being discussed, then the burden of proof is on the arguer to demonstrate both (A) and (B) by producing the opposing view and demonstrating its merits.

Cappelen and Lepore's second inconsistency charge is that Pragmatism, contrary to its own claims, is committed to there being some way of specifying truth conditions for context

sensitive sentences. They give the example, “‘Smith weighs 80kg’ is true in a context C iff Smith counts in C as weighing 80 kg” (2005, 133). I won't give this claim much attention because nothing in how I've defined Pragmatism commits it to the view that truth conditions along those lines cannot be specified for context sensitive sentences. Cappelen and Lepore attribute the claim that such truth conditions are impossible to specify to several prominent theorists including John Searle and Charles Travis (2005, 131). But even if they are right in so attributing this claim and even if Searle and Travis count as Pragmatists, it does not follow that Pragmatism is itself committed to it.

Cappelen and Lepore might object that they argued earlier in their book for the claim that what they call “Moderate Contextualism” inevitably collapses into “Radical Contextualism”, and so even theorists who don't take themselves to be committed to the claims of Travis are in fact so committed. However, their earlier argument about the collapse of Moderate into Radical Contextualism cannot support this claim. That argument simply attempted to show that if one finds CSAs for some expressions outside of the Basic Set plausible then one should find CSAs for nearly every expression plausible (2005, 39-40). But Cappelen and Lepore only dispute CSAs insofar as they are intended to demonstrate semantic context sensitivity. They accept CSAs as sound arguments for claims about contextual variability in what is said. So on this construal Cappelen and Lepore's arguments about the collapse of Moderate into Radical Contextualism would only show that nearly every sentence can be used to say many different things. No Travis-style claim about the impossibility of specifying truth conditions follows from that.

Cappelen and Lepore's final inconsistency charge is that Pragmatists inevitably end up using expressions in their accounts in ways that conflict with their theoretical claims (2005, 136-140). As an example, they cite a CSA given by Bezuidenhout about 'blue' that attempts to show

that its application to a given quantity of ink depends on the context. In particular she argues that context affects whether an utterance of 'The ink is blue', will be true just in case it looks blue when written on paper or just in case it looks blue in its container. Cappelen and Lepore note that in giving her CSA Bezuidenhout uses the word 'blue' multiple times in such a way that it is intended to apply to ink from both of the different contexts. They essentially argue that this commits her to the view that 'blue' does in fact have a context insensitive meaning, one that can be used to describe the many different things that are said to be “blue” in different contexts.

As far as I can tell, the above argument represents a slightly altered version of Cappelen and Lepore's second context sensitivity test argument, the collective description test. There Cappelen and Lepore argued that minimal propositions explained why it is in general so easy for there to be a true utterance of 'A and B are blue' if there are true utterances of 'A is blue' and 'B is blue'. Here they are claiming that minimal propositions account for the possibility of doing what amounts to the same thing over the course of a paragraph, applying 'blue' to both A and B. If my arguments in section 4.1.2 worked in the case of collective descriptions, they should work again here.

5. Application to Contextualism

As mentioned earlier, Cappelen and Lepore have a broader target than Pragmatism, a target they refer to as 'Radical Contextualism'. Pragmatism, when applied fully generally, is a form of “Radical Contextualism” as they use the term, but the latter view also includes theories according to which all context sensitivity in what is said can be traced to context sensitive items in the logical form of the sentence uttered. Thus, such views have a richer conception of the logical form of such sentences than both Minimalists and Pragmatists. Stanley and Szabo (2000)

offer an example of such a view, though in a moderate rather than radical form.

Since my defense of Pragmatism frequently relied on Pragmatism's great similarity to Minimalism—including each view's eschewal of hidden context sensitive variables in logical form—it would be natural to suppose that my arguments offer no help to the more distinct Contextualism. In this section, I make the case that my arguments serve both views. I'll begin where the case for this is clearest, with the arguments from shared content across context and inconsistency arguments.

Recall that I rebutted Cappelen and Lepore's argument that Pragmatists cannot account for understanding meaning without knowing relevant context by pointing out that Pragmatism supplied all the tools necessary to infer exactly as much meaning from an utterance as Minimalism does. In particular, the Pragmatist notion of sentence meaning allows a hearer to infer that the speaker believed the minimal proposition associated with the uttered sentence assuming the speaker spoke literally and truthfully. If that's so, then a Contextualist should similarly have no problem with this argument, because their view provides utterance interpreters with all the tools that Pragmatism does in addition to richer logical forms for the uttered sentence. These forms should if anything allow the listener to infer even more from the utterance of that sentence (again assuming literalness and truthfulness). Thus, Contextualists can account for understanding what was said across contexts as well as Pragmatism and Minimalism can.

The inconsistency charges against Contextualism can be dealt with even more easily, as each of my responses from section 4.3 apply directly to that view as well. The *tu quoque* and burden of proof responses to the first inconsistency charge need no modification to be applied to Contextualism. And the second inconsistency charge relies on the assumption that Radical Contextualism disallows context sensitive sentences from having specifiable truth conditions,

which is no part of the definition of the view. Similarly, my argument that the third charge fails if their argument for the second context sensitivity test fails applies here as well.

In light of that last point, we should move on to the context sensitivity tests, where you may suspect my rebuttals will stop applying to Contextualism. In section 4.1, I argued that because the only significant theoretical difference between Pragmatism and Minimalism was that the latter view is committed to minimal propositions, any context sensitivity test meant to support Minimalism over Pragmatism must be motivated by this commitment. And I found that no such motivation was available for any of the tests. But since Contextualism has additional theoretical differences from Minimalism, it's not clear that such an argument will help the view.

To get around this difficulty let's think more generally about what would make a linguistic test a test of context sensitivity. Whether they are pronounced or articulated only in logical form, context sensitive expressions have two notable features related to their impact on the meanings of sentences containing them. First, for the sentences containing them to express compositionally determinate propositions, the expressions must be assigned meanings by context. Second, the sentences containing them should be capable of being used to say different things in different contexts of utterances, depending on what value the context determines for the expression. These two features of context sensitive expressions suggest two ways in which one can test for the presence of context sensitive items in a sentence. You can assess whether the sentence fails to express a minimal proposition in the absence of contextual supplementation, or you can assess whether the sentence can be used to say different things in different context. But that the relevant sentences can be used to say different things in different contexts is common ground among all three views under discussion. So this feature of context sensitive expressions cannot be used in a test meant to support Minimalism over either Pragmatism or Contextualism.

If the above is correct, then even though Contextualism differs strongly from both Minimalism and Pragmatism in its view of how context sensitivity is implemented in language, it would still be the case that any dialectically effective Minimalist argument from context sensitivity tests would have to rely on Minimalism's distinctive claims about minimal propositions. And if that's so, then each of my arguments against there being any interpretation of the tests that succeeds in motivating the existence of minimal propositions will defend Contextualism from Cappelen and Lepore's arguments as well.

This isn't to say that the tests couldn't possibly provide any information relevant to deciding among the views. Recall that in discussing the first test we noted that context sensitive items from the Basic Set seem to behave differently in indirect reports than other expressions that appear to give rise to variation in what is said. This result might be somewhat less amenable to Contextualism than to other two views under discussion. The other two views already assume there is a difference between the linguistically determined context sensitivity due to items from the Basic Set and what they view as the non-linguistically determined variation in what is said that other expressions give rise to. In light of this distinction, neither view should be surprised by a result showing that these two types of context sensitivity give rise to differences in other areas. Since Contextualism views all context sensitivity as being of the same type as the context sensitivity of the items in the Basic Set, they might have a somewhat harder time accounting for the results of the first test. That said, even this outcome would be far from the knockdown argument against Contextualism that the test might be seen as presenting.

6. Conclusion

Let's take stock of what this paper has and has not shown. It certainly has not shown that

Pragmatism is the correct view. A much more comprehensive examination of it and its competitors would be needed for that. It also hasn't shown that Minimalism is false, though it undermines it by suggesting that Minimalism requires commitment to more theoretical assumptions than we need to make to explain all the relevant data. Instead, if my arguments have worked then this paper has shown first that Minimalism and Pragmatism are considerably more similar than is often thought and second that the few differences between them do not afford Minimalism the advantages that Cappelen and Lepore claim for it.

I argued for the above claim in two ways depending on the nature of the Cappelen and Lepore argument I was rebutting. In some cases, I argued that Pragmatists can account for facts that minimal propositions are claimed to explain. In other cases, I argued that Cappelen and Lepore's arguments were misdirected, either because they didn't apply to Pragmatism at all or because they applied to Minimalism as well. If my arguments have succeeded, they suggest that whether a sentence semantically expresses a proposition has less theoretical importance than has generally been assumed.

CHAPTER 2

PROBLEMS FOR METASEMANTIC EXPRESSIVISM

1. Introduction

Many of the most persistent metanormative issues arise from the fact that normative judgments appear to have features in common with both desire-like and belief-like states of mind. Normative judgments appear desire-like in that they seem capable of motivating actions in the characteristic fashion of desires. They seem belief-like in that they can take part in valid arguments and are expressed by sentences that share all the linguistic properties of straightforwardly descriptive sentences.

Traditional metanormative positions tend to take one of these features of normative judgments and claims as more central than the other.²⁰ Traditional expressivists, who claim that normative statements express desire-like states of mind, focus on the first feature. Cognitivists, who claim that normative statements express belief-like states of mind, focus on the second. Unsurprisingly, each view tends to have a hard time accounting for the feature of normative judgments and discourse focused on by the other view.

Recently, philosophers have attempted to get around these problems by proposing views that aim to take each feature of normative judgments fully seriously. This paper focuses on the attempt to do so by treating expressivism as a theory of the metasemantics of normative claims rather than the semantics. That is, these views attempt to understand expressivism as explaining

²⁰ These traditional positions were also traditionally called “metaethical” positions rather than “metanormative” ones. On my use, the normative encompasses a broader class of judgments and discourse than the ethical, including for example all things considered judgments of what one should do or think. These sorts of judgments quite generally give rise to the sorts of puzzles discussed in this paper.

The category of the normative is notoriously difficult to define, and I won't say more than that normative judgments are closely connected to judgments of what one ought to do, think, or feel.

why normative statements have the meanings they do, rather than giving a non-standard account of what those meanings are. In this paper, I argue that metasemantic expressivism cannot actually provide an account of why normative claims have their meanings. In doing so, I will devote the largest space to Michael Ridge's version of metasemantic expressivism as laid out in his *Impassioned Belief*, as it is the most well developed such view.

The paper will proceed as follows. Section two explains in more detail the two above features of normative judgment and the problems they cause for traditional metanormative theories. Section three explains the difference between semantics and metasemantics on the way to presenting Ridge's metasemantic expressivism. Section four presents problems for Ridge's view and those of a few other recent metasemantic expressivists. There I argue that such views are incapable of giving a metasemantics for normative claims. Finally, in section five I turn to the question of whether metasemantic expressivism would even have significant theoretical advantages if it could be made to work, arguing that it would not. This argument involves, among other things, giving an account of how much deference philosophers studying moral discourse should pay to semantic theories.

2. Internalism and the Frege-Geach Problem

To start with, we need a clearer account of the difference between desire-like and belief-like mental states. One popular account is in terms of direction of fit.²¹ Anscombe's famous example of the difference between a man's shopping list and a list with the same items compiled by a detective following him illustrates the distinction well:

²¹ Kim Frost (2014) among others has questioned the cogency of the notion of direction of fit. However, since my project is critical of metasemantic expressivism and its targets rely on some notion of direction of fit, this doesn't threaten my claims.

if the list and the things that the man actually buys do not agree...then the mistake is not in the list but in the man's performance...whereas if the detective's record and what the man actually buys do not agree, then the mistake is in the record (1963, 56).

Another way of putting the claim is that the shopping list has a world-to-mind direction of fit and the detective's record has a mind-to-world direction of fit. A similar difference occurs between desire-like and belief-like mental states. The former have world-to-mind direction of fit because they function to change the world.²² Thus, if there is a discrepancy between desire and the world, then the function of desires will be served by changing the world to match your desires. Beliefs, on the other hand, have mind-to-world direction of fit because their function is to represent the world.²³ If your beliefs fail to match the world, then they will only come to fulfill their function if you change your beliefs to match the world. I will call states with a world-to-mind direction “non-cognitive”, and ones with the other direction “cognitive”.

We can now ask which direction of fit normative judgments have. The plausibility of internalism about normative judgments supports the claim that they are non-cognitive. On one of its many construals, internalism states that normative judgments are necessarily capable of motivating action.²⁴ This thesis may seem undermined by the fact that people frequently don't act on their normative judgments. For instance, many judge that they should give more to charity than they do. But as we've defined it, internalism requires only that normative judgments always have the *capacity* to motivate action. The degree of motivation normative judgments provide

22 Or at least to bring it about that the world is a certain way. If the world is already that way, they function to maintain it.

23 At least, the representational beliefs have this function. It is not obvious that our pre-theoretic notion of beliefs perfectly coincides with the technical notion described here. Normative judgments are generally described as beliefs (e.g., the belief that you should not lie), so if they turn out to have a world-to-mind direction of fit then it may turn out that some beliefs aren't representational.

24 See, e.g., Ridge (2014) chapter two.

may be ridden overridden by other motivations such as greed, but we still expect there to be something that those other motivations overcame to lead to your action.

Cognitivists, who think that normative judgments are cognitive, have difficulty explaining the appeal of internalism. According to them, the function of normative judgments is to represent the world, not to change it. We would thus expect people that make the normative judgment that they should give to charity to only be motivated to give to charity if they also have an additional non-cognitive state, such as the desire to give to charity. Perhaps cognitivists could postulate that people just happen to have desires that accord with what they think they should do. But then it is hard to see why people's motivations would tend to change in accordance with changes to their normative judgments.²⁵ If a person lacked a desire to give to charity prior to deciding that they should give to charity, then why would we expect them to suddenly get such a desire once they decide they should? An appealing answer is that the person's judgment that they should give to charity is itself a state that motivates action.²⁶

Non-cognitivists think just that, that normative judgments are non-cognitive. Expressivists add to this an explanation of the meaning of normative utterances in terms of the states of mind those utterances express.²⁷ Such a view can easily explain the appeal of internalism about normative judgments. If normative judgments just are states whose function is to change the world, then it is no surprise that they are capable of motivating action. But

25 See Smith (1994) for a presentation of this argument.

26 The preceding should not have convinced any cognitivists of either internalism or the claim that cognitivists cannot account for it, and it wasn't meant to. But hopefully it did give some idea of why alternatives to cognitivism have been sought by so many metanormative theorists.

27 As we'll see in the next section, "explanation of the meaning" is purposely ambiguous between semantic and metasemantic versions of expressivism. Also note that the relevant notion of "expressing" a mental state here is one according to which an assertion can express a mental state that its speaker lacks. I'll say more about this later.

expressivists face a serious issue of their own in the Frege-Geach problem, which involves accounting for the meanings of normative expressions embedded in unasserted contexts.²⁸

The Frege-Geach problem is easiest to understand as applied to the simplest forms of expressivism, like A.J. Ayer's (1952) emotivism. According to emotivism, normative utterances function only to express the emotions of the speaker. For instance, uttering 'Killing is wrong' conveys more or less the same thing as 'Boo killing!'. And like the latter utterance, the former is not even truth evaluable. Whatever plausibility such a view has for simple predications of normative expressions, it quickly loses it when we look to more complicated types of sentence containing normative terms.

Consider: 'Killing is not wrong', 'If killing is wrong, I won't kill', 'I wonder whether killing is wrong', and 'John hopes that killing is not wrong'. It is not at all clear what such sentences could even mean according to emotivism. For one thing, none of the sentences express a negative emotion toward killing. For another, as we will see in the next section, the emotivist cannot appeal to the standard compositional meanings provided by semantics because they reject a truth-conditional account of the meaning of the normative sentences. The cognitivist, however, can easily account for such meanings using the same compositional accounts as used for other declarative sentences, since they need not think there is any important difference between normative declaratives and others.

The emotivist might try get around the first issue by claiming that 'wrong' has a different meaning in the unasserted contexts above from its meaning in atomic sentences. But if they go

28 Geach (1965) gave the canonical account of this issue, attributing his insight to Frege's distinction between force and content. Frege did not actually write on expressivism, which didn't yet exist. For further discussion, see Dreier (1996), Unwin (1999), and Schroeder (2008).

this route, then they will not be able to explain the validity of arguments involving 'wrong'. For instance:

(1) Killing is wrong.

(2) If killing is wrong, so is letting die.

(3) Letting die is wrong.

(3) follows from (1) and (2), but if 'wrong' means something different in the unasserted context of (2) from the asserted one of (1), then the argument can't be valid. It would simply trade on an equivocation over different senses of 'wrong'. The ambiguity proposal can only explain the meaning of complex normative sentences at the cost of losing any hope of explaining the validity of normative arguments.

As before, what I've said shouldn't convince any expressivists to abandon their view. There are far more sophisticated versions of expressivism with far more plausible accounts of the Frege-Geach problem.²⁹ But it is extremely controversial whether any of the expressivist proposals have succeeded in solving the problem, or even whether any have accounted for the meaning of something as simple as negated normative sentences.³⁰ What I do hope to have done, in explaining the issues of both internalism and the Frege-Geach problem, is to have given some sense that there are two major issues about normative judgments and assertions that need explaining and that views that have been good at explaining one of these issues have fared poorly at explaining the other.

This situation has led some to seek alternatives to the main traditional views. One such alternative is to propose a hybrid view, according to which normative judgments involve both

²⁹ See, for instance, Gibbard (1990; 2003) for such accounts.

³⁰ Unwin (1999) and Schroeder (2008) cover this issue in detail.

cognitive and non-cognitive states.³¹ For instance, the judgment that killing is wrong might comprise both a non-cognitive disapproval of killing and the cognitive belief that killing fails to maximize utility. The hope is that the non-cognitive portion of the normative judgment can explain the motivational force of that judgment, whereas the cognitive component can account for the Frege-Geach problem. Most naturally, embeddings of normative sentences can be understood as involving thoughts corresponding to embeddings of the cognitive portion of the judgment. As we will see, Ridge presents one such view, but he also combines it with a metasemantic account of the commitments of expressivism. It is to that account that we now turn.

3. Ridge's Metasemantic Expressivism

In order to explain how Ridge's metasemantic view works, we must first clarify the difference between semantics and metasemantics. Ridge glosses semantics as aiming “to provide a systematic characterization of the literal meanings...of words and sentences” (2014, 102). On this account, the semantic value of a sentence paired with a context is its literal meaning. This is on the right track, but it would be surprising if the theoretical notion of semantic value matched up exactly with our everyday understanding of literal meaning.³² Compare the case of 'weight' as discussed before Newtonian physics. There are clearly a number of connections between the ordinary understanding of weight then and the technical notion from Physics, but they are far from identical. The ordinary understanding, for instance, would not have accepted the claim that objects would have different weights depending on what bodies they were near.

31 See Schroeder (2009) for a survey of such views. Ridge (2014) calls them “Ecumenical” views.

32 'Literal meaning' is sometimes also used in philosophy of language to refer to the technical notion of semantic value. This is fine, but then it doesn't clarify what semantic values are to say they are literal meanings.

In light of this issue, it is better to understand the notion of semantic value in terms of the explanatory roles it plays within semantic theories.³³ Assigning semantic values to expressions is part of giving an overall semantic theory that accounts for a variety of linguistic phenomena. Most prominently, semantic theories give an account of the entailment relations between sentences, and entailment relations explain a number of linguistic facts. For one thing, by demonstrating which sentences contradict each other, they add to an explanation of why people don't generally utter the pairs of sentences the theory deems contradictory.³⁴ In a similar fashion, accounting for entailment relations between sentences can help explain why people generally don't utter a sentence that straightforwardly follows from their previous sentence, at least not without drawing attention to this fact. Additional explananda of semantic theory include ambiguity, synonymy, and the role of semantic values in communication.³⁵ The view sketched here allows us to understand semantic values as whatever objects play the semantic value role in theories that explain these features of language.³⁶

Before moving on, I should note that Ridge generally discusses the meanings of normative claims rather than of normative sentences. This is because, on Ridge's view, there are no normative sentences, only normative uses of sentences. He gives a context sensitive account of the meanings of seemingly normative terms such as 'good', 'reason', and 'ought' according to which whether sentences involving these make normative claims depends on the context of utterance. However, on my understanding of semantic values, claims don't have semantic values,

33 I take this understanding of semantics and metasemantics from Yalcin (2014). Read his article for a fuller account of the view sketched here.

34 The theory could do this in tandem with an account of the relationship between semantic values and communication that would show that asserting contradictions could not serve the ordinary goals of communication. Consider in particular Stalnaker's account of assertion (1984).

35 Semantic theories attempt to explain these notions compositionally by showing how the semantic values of complicated sentences are determined by the semantic values of their parts. Compositionality allows semantic theories to explain how finite beings like us can understand indefinitely many sentences.

36 One way of understanding how this gives rise to a meaning for 'semantic value' is found in Lewis (1970).

only linguistic items do. So I will continue to speak of the semantic value of normative sentences and expressions, and Ridge can understand this as talk of the semantic value of these sentences relative to a context in which they have a normative meaning.

With the above view of semantics in mind, we can understand metasemantics as explaining the facts in virtue of which semantic facts hold. In particular, a complete metasemantic theory would have to explain why expressions have the semantic values they do. There are different understandings of what this kind of explanation amounts to. On one understanding, it is a request for a causal account of how an expression ended up with the semantic value it did. On another, the explanation is of a more metaphysical sort, accounting for the facts that constitute or ground an expression's having the semantic value that it does.³⁷ The term 'metasemantics' has been applied to each of these, but I intend the latter understanding.

We are now in a position to understand the difference between semantic and metasemantic versions of expressivism. Semantic versions of expressivism are accounts of the semantic values of normative terms. They are expressivist in that the semantic values they posit are the mental states expressed by normative utterances, where these mental states are non-cognitive ones. The goal of such accounts is to provide a systemic way of deriving the mental state expressed by each normative sentence from the component parts of that sentence. Textbook semantics, on the other hand, gives a truth-conditional account of the semantic values of normative sentences. This is for the simple reason that normative declarative sentences appear to have all the same linguistic features as non-normative ones, and the textbook semantics for declarative sentences in general is truth-conditional. Semantic expressivism, thus, offers a revisionary view of the meanings of normative expressions.

³⁷ See Fine (2012) for more on grounding. Some theorists take the causal account to be part of the ground, in which case it belongs here as well.

Metasemantic expressivism is motivated in large part by the problems that the above revisions engender.³⁸ We've already seen that expressivism has a hard time providing compositional meanings for embedded normative clauses. Textbook semantics has no problems with such cases. Additionally, some philosophers have argued that philosophy has a poor track record when it comes to challenging scientific results, and as such we should be wary of doing so. According to this argument, since Semantics is an independent scientific discipline, we should greatly defer to what semantic theories tell us. Alejandro Carballo voices the concern:

Naturalists have long been skeptical of any philosophical view that conflicts with the deliverances of the natural sciences. Philosophy, on this view, should never proceed on the assumption that it has the epistemological high-ground with respect to the sciences. If constraints internal to semantic theory tells us that 'Torture is wrong' has a certain set of worlds as its semantic value, then so be it (2014, 123-124).

Ridge makes a similar claim (2014, 104). I'll say more about this argument in section five.

Metasemantic expressivism avoids these problems by retaining a standard first-order semantics for normative expressions. It requires no revisions and thus properly defers to the work of Semanticists. Additionally, since the standard semantics has accounts of the meanings of complex normative sentences, these should no longer pose a problem for this kind of expressivism.³⁹ Metasemantic expressivism manages to remain distinctively expressivist despite these accommodations by giving a different explanation from the cognitivist of why normative sentences have these semantic values. Cognitivist accounts will explain the meanings in terms of

38 In addition to Ridge (2014), metasemantic expressivism is considered in Carballo (2014), Chrisman (2012), Chrisman (2013), Chrisman (2016), Charlow (2013), Sepielli (2012), Silk (2013).

39 Note though that Ridge denies that going metasemantic is sufficient to solve the Frege-Geach problem (2014, 137-138). He thinks he needs to "earn the right" to make use of the machinery of first-order semantics by showing how his hybrid expressivism can account for complex meanings and valid arguments involving normative sentences. As I'll argue in the next section, I agree here.

the representational function of discourse involving those sentences. The expressivist account will appeal instead to the non-cognitive states of mind of those uttering the sentences in offering their account of the sentences' semantic value.

More precisely, metasemantic expressivism can be understood as the conjunction of three theses:

- (1) Normative judgments are non-cognitive.
- (2) Normative sentences have standard, truth-conditional semantic values.
- (3) Normative sentences have these semantic values in virtue of expressing non-cognitive judgments.⁴⁰

Thus far, I've explained only what it would be for something to count as a metasemantic expressivist view. But what more specifically would such a view claim? Here I'll focus on the most extensively worked out view, argued for by Michael Ridge in *Impassioned Belief*.

Ridge offers a metasemantics meant to explain the meanings of all sentences in terms of the mental states they express. Wayne Davis presents one way of doing this. On his view, like Grice's, sentence meaning is explained in terms of speaker meaning.⁴¹ What a speaker means by uttering a sentence is determined by their intentions to indicate their state of mind using the utterance. What a sentence means is a matter of what it is conventionally used by speakers to mean. Ridge slightly favors an alternative to this view, but the details of neither view will matter to my arguments. What does matter is that this account of metasemantics explains the meanings of sentences in terms of the mental states of sentence utterers.

40 As we will see soon, Ridge offers a hybrid form of expressivism wherein normative sentences express a mental state consisting of both a cognitive and a non-cognitive part. So, you should understand "non-cognitive judgments" in this thesis as including this kind of hybrid judgment.

41 See Grice (1957) and Davis (2003).

Ridge combines this metasemantic view with the theory of propositions advocated by Scott Soames.⁴² On Soames' view, propositions are cognitive event types. The proposition that grass is green is the type corresponding to mental predications of greenness to grass. More generally, on Ridge's construal, the proposition expressed by a sentence is simply the type of the mental state that sentence expresses. So, the proposition that pleasure is good is the mental state type of judging pleasure to be good.⁴³ If correct, Davis' account of what it is for a sentence to express a mental state and Soames' account of propositions together account for in virtue of what all declarative sentences having the meanings (propositions) that they do. Each sentence has a given proposition as its meaning in virtue of the fact that that sentence expresses the mental state type that is that proposition.

But what is expressivist about this metasemantics? Ridge's expressivism comes through *via* his claim that normative judgments are hybrid mental states that involve both a desire-like state of mind and a representational state of mind. Specifically, the judgment that pleasure is good will involve two, related mental states: a normative perspective and the belief that the standards determined by that perspective rank pleasure high.⁴⁴ The details of Ridge's notion of normative perspectives are not important here. All you need to know is that these states are non-cognitive and that they determine a set of standards by which one ranks various things. Though the relevant standards are determined by one's normative perspective, the belief that the standards rank pleasure high is an ordinary representational belief.

42 Soames (2010).

43 More carefully, the mental state of predicating goodness of pleasure. One predicates goodness of pleasure even when one merely considers whether pleasure is good while withholding judgment. Since it won't make a difference to my arguments, I'll be loose with this distinction here. Note also that I am using 'mental state' such that it can apply to events.

44 This is a simplification, but the parts I leave out should not matter. See chapter four of Ridge (2014) for more details.

An example will make this clearer. Suppose you are a utilitarian. You will thus have a utilitarian normative perspective, which for our purposes can be thought of as like a plan to bring about high utility outcomes or a state of approving of those outcomes. This perspective determines a set of standards for ranking various properties or states of affairs, namely one according to which states of affairs with the most pleasure are ranked highest. The judgment that lying is wrong for such a person would consist in having a utilitarian normative perspective and also believing that utilitarian standards rank lying lowly. This belief represents the world as being such that lying does not bring about the most pleasure.

There are many other ways to judge that lying is wrong, though. For a Kantian, such a judgment would consist in a Kantian normative perspective along with the belief that Kantian standards rank lying lowly (a plausible belief). Most people, of course, don't have normative perspectives as fully worked out as utilitarianism or Kantianism. But such people still have some relevant motivational state that in some cases determines a ranking of outcomes. So they too can make the judgment that lying is wrong, only with a different pair of mental states constituting that judgment. In total, there will be as many hybrid states constituting the judgment that lying is wrong as there are normative perspectives, which is to say, indefinitely many.

Tying this back to the earlier discussion, the proposition expressed by 'lying is wrong' is the event type of judging that lying is wrong. As we've just seen, there are indefinitely many ways to make this judgment, so no particular normative perspective or related belief will be part of the proposition. Instead, the proposition is simply the type of mental state one is in when one has some normative perspective or other and believes that the standards determined by that perspective rank lying low. And the reason that is the proposition expressed—the metasemantic account of why the sentence has the meaning that it does—is that the sentence is conventionally

used to indicate that mental state. Ridge thus offers a metasemantics that explains how non-cognitive states of mind in part determine the meanings of normative sentences.

4. Problems for Metasemantic Expressivism

The biggest issue facing metasemantic expressivism is that expressivists are not in fact capable of explaining in virtue of what normative expressions have their semantic values. I'll introduce and defend this claim as applied to Ridge in the first subsection, after explaining how it differs from another argument that has been raised against Ridge. In the next one, I expand this objection to cover versions of metasemantic expressivism that attempt to explain in virtue of what sentences have a given set of possible worlds assigned to them by semantic theories.

4.1. Ridge's View

In assessing whether metasemantic expressivism has the benefits promised, the biggest issue for Ridge arises from the fact that he doesn't explain what grounds the semantic values that actually get assigned by semanticists. Instead, he explains what grounds a given Soames-style proposition being assigned to the sentence. In a critical review, Andrew Alwood picks up on this point and leverages it to claim that Ridge's view is not actually metasemantic (2016). Recall that semantic expressivism treats mental states as the meanings of normative sentences. Ridge's view is supposed to be metasemantic in that he instead retains propositions as the meanings of such sentences while offering an account of why each sentence has the meaning it does. This may seem like a clear contrast; but according to Alwood, Ridge's endorsement of Soames' view of propositions undermines it.

The view of propositions Ridge endorses characterizes them as mental event types. Due to this, to say that the meaning of 'killing is wrong' is a proposition is to say that its meaning is the event type of judging that killing is wrong—a judgment that is partially constituted by a non-cognitive mental state. But that's precisely what semantic expressivism claims: that the meanings of normative sentences are non-cognitive mental states.⁴⁵ In short, Alwood's worry is that once you understand Ridge's view of propositions, his own view appears to simply be a way of relabeling standard semantic expressivism in terms of propositions. For Ridge—just as much as traditional expressivists—normative sentences are assigned different first-order meanings than non-normative ones, since the meanings assigned to only the former are partially constituted by non-cognitive mental event types.

It is true that Ridge also offers a metasemantics that promises to explain why each mental state is associated with each sentence. But metasemantic expressivism was not advertised as the conjunction of expressivism and any metasemantic view. Presumably most expressivists think that there is some story or other about what constitutes or grounds normative sentences having the meanings they do. Gibbard even offers one in *Wise Choices, Apt Feelings*, adverting to a Millikan-style account of metasemantics.⁴⁶ And, as with Ridge, the reason normative sentences have the meanings they do comes partly from the non-cognitive states of speakers. But this does not make the view Gibbard presents there a metasemantic form of expressivism because, again as with Ridge, the meanings of normative sentences themselves differ in kind from those of descriptive ones.

45 Ridge is a hybrid expressivist, so for him the meaning is also partly constituted by a cognitive state. But this difference has nothing to do with his view being metasemantic; it simply stems from his providing a hybrid view.

46 See Gibbard (1990) and Millikan (1984).

Alwood's objection points toward the biggest problem for Ridge, but it doesn't correctly identify it. Metasemantic expressivists are committed to the claim that normative sentences have (standard) semantic values in virtue of expressing non-cognitive mental states. Their distinctive claim is that assigning a truth-conditional semantic value to a sentence does not commit you to the claim that that sentence represents the world as being a certain way. Instead, they believe a sentence can have its truth-conditions in virtue of expressing a mental state that functions to motivate action. It is no objection to such a view to note that they see semantic values as being intimately tied to non-cognitive mental states.

However, Ridge's view is still threatened by the fact that semanticists don't actually assign Soames-style propositions as the meanings of sentences. Unless Ridge can explain how the actual meanings assigned by semantic theories are determined, he hasn't offered a metasemantics for the standard theories. To that end, let's take a closer look at semantic values as they are found within the discipline of semantics.

For simplicity, we will try to account for the semantic value of the normative predicate, 'good'. The textbook account of the semantic value of 'good' is that it is an extension: the set of good things.⁴⁷ So our question is: what should Ridge's metasemantics of predicate extensions be? Ridge doesn't directly answer this question, but he does give a theory of judgments of an object being in the extension of a predicate. Such judgments, according to Ridge, are normative, and

⁴⁷ See Heim and Kratzer (1998). This is a significant simplification in several ways. For one, the actual semantic value is the characteristic function of this set—the function that maps individuals to the number one if and only if those individuals are good—rather than the set itself. Giving predicates functional meanings along with providing rules for when those functions are to be applied to the meanings of other expressions in the sentence helps the theory account for compositionality.

Another simplification is that the value might also end up being a function of worlds, events, or times depending on how these help explain the meaning of the predicate in more complicated contexts. We will look at how adding in worlds affects things shortly. And finally, I ignore issues of context sensitivity. None of these simplifications should threaten the point made in this section.

thus consist in hybrid states of a related pair of non-cognitive and cognitive attitudes.⁴⁸ The specific judgment assigned is a normative perspective along with “The (representational) belief that [an object] *O* falls into the extension of [the predicate] 'F' on any admissible specification of its referent” (2014, 212). Being an “admissible specification” depends in part on according with the standards determined by the normative perspective. For a utilitarian perspective, only utility maximizing acts will fall under the extension of 'good', for instance. Given that something being in the extension of 'good' entails that it is good and *vice versa* (assuming no context sensitivity), it makes sense that a normative theory of judgments of goodness would lead to a normative theory of judgments of the extension of 'good'.

So, in virtue of what does 'good' have its extension? You might think that for a particular use of 'good', its extension is determined by the utterer's normative perspective, as this partially determines what we judge to be good. But that answer would reduce expressivism to subjectivism, the view that 'x is good' means the same thing as 'I approve of x'. Since expressivism was developed largely in response to the problems with simple subjectivism, this will not do. Suppose someone utters 'Torture is good'. On the above view, their use of 'good' has an extension determined by their normative perspective. Torture will thus fall under the extension of their use of 'good'—making their utterance true—just in case their normative perspective ranks torture high. The view would force you to say that someone with an abhorrent normative perspective speaks truly when they say 'Torture is good'. Since Ridge argues at length against these consequences of subjectivism, he would not accept them as consequences of his on view. We need another account of the metasemantics of normative predicate extension.

48 This includes judgments involving the extensions of non-normative predicates like 'blue'. But, though such judgments technically express both a normative perspective and a representational belief, for non-normative predicates the normative perspective has no effect on what the representational belief is.

Because Ridge adopts Soames' view of propositions, it is natural to think that he can simply make use of that view's account of how predicate extensions are determined, but this won't work either. In Soames' view, the extensions of predicates are determined by the event types of predication that those predicates express (2010, chapters five and six). For instance, 'blue' expresses the event of predicating blueness, and due to this the extension of 'blue' will be the set of things that have the predicated property—the set of blue of things. Soames can say the same thing about normative predicates like 'good' if he is not an expressivist. But expressivists reject the claim that you can explain the meaning of normative predicates in terms of their standing for normative properties. Instead, if they allow talk of normative properties at all, they will say that they have to explain how such talk can be made sense of within their framework. So metasemantic expressivists cannot make use of the extensions of normative properties to explain the extensions of normative predicates, as would be needed for the above explanation to work.

The failure of the two explanations above was not a coincidence. Rather, it points to the impossibility of Ridge or other metasemantic expressivists giving an account of in virtue of what normative predicates have their extensions. The reason is that what extension 'good' has is not even a semantic matter on Ridge's view. It would only be a semantic issue if there were some linguistic convention that partially determined an extension for 'good'.⁴⁹ For Ridge, the linguistic conventions governing 'good' determine only that one have a normative perspective and that one predicate 'good' of something only if one's normative perspective ranks it high. That is, each normative perspective corresponds to a (potentially) different extension of 'good'. Since linguistic conventions do not determine a correct normative perspective, the correct extension is similarly undetermined. There cannot be a metasemantic explanation for something that isn't even a fact of

49 Or at least a function from contexts to extensions.

semantics. If this is right, then not only does Ridge not provide a metasemantics for extensions, he cannot possibly provide one. His expressivism commits him to the view that extensions of normative predicates are not determined by any metasemantic fact.

Another way to see this point is to compare metasemantic expressivism to another recently proposed view about the metasemantics of normative expressions: the claim that normative expressions are *referentially normative*.⁵⁰ To claim that an expression is referentially normative is to claim that its reference is determined by its normative role. Here, “normative role” refers to certain non-cognitive mental states that are associated with the use of the expression. For instance, one might claim that the normative role of 'ought' is such that a person who judges that they ought to bring about some scenario, P, will generally intend to bring it about that P. The claim that 'ought' is referentially normative would then entail that any term that shares the normative role of 'ought' would have to have the same reference as it, even if the terms were consistently applied to completely different objects.

Both the view just sketched and metasemantic expressivism endorse the claim that the non-cognitive states associated with normative expressions play a central role in determining the semantic values of those expressions. But the views are actually quite different. Claiming that normative expressions are referentially normative commits one to a very strong form of metanormative realism. Not only are there objective normative properties, but these properties are reference magnets. Any expression that a community uses with a normative role will have its reference drawn to these normative properties. Suppose that both I and someone in a distant galaxy guide our actions using a concept with the normative role of 'ought', and suppose that our actions are so-guided toward completely different things (perhaps the alien wishes to minimize

⁵⁰ See Wedgwood (2001), Eklund (2017), and Williams (2018) for discussions of this idea. The term is Eklund's.

the well-being of the universe). According to the view under consideration, our concepts will pick out the same property. And at least one of us must be greatly mistaken about what states of affairs have that property. It is hard to get more realist than that. Metasemantic expressivism, on the other hand, is an attempt to make a form of metanormative anti-realism consistent with semantics.

The issue is that it is hard to see how metasemantic expressivism could possibly provide explanations of context-insensitive extensions of normative predicates without just turning into a version of the view just discussed. If the non-cognitive states expressed by normative utterances determine an extension for 'good', then the terms others use to express the same non-cognitive states should have the same extension. Perhaps there is a way that such a view could be developed that would not simply be a version of the staunchly realist view described in the preceding paragraph, but I don't currently see how. It thus seems to be a core feature of expressivist views that they cannot explain the extensions of normative predicates in the manner needed to accord with orthodox semantics.

4.2. Metasemantic Expressivism and Possible Worlds

The metasemantic expressivist views discussed by Chrisman and Carballo may be in better shape. They both share Ridge's view that adopting an overall metasemantics in which the thought expressed by a sentence determines that sentence's meaning allows one to be an expressivist without abandoning standard semantics.⁵¹ But unlike Ridge, they frame their discussion directly in terms of explaining the truth conditions of normative sentences, rather than explaining how they get assigned Soames-style propositions.

⁵¹ See Chrisman (2016, 171-175) and Carballo (2014, 121). Chrisman, like Ridge, labels this form of metasemantics "ideationalism". Carballo uses "mentalism" for the same idea.

Of the two, Carballo goes into greater detail in explaining the relationship between truth conditions and non-cognitive mental states. He describes the following metasemantic picture:

Part of what it is for the meaning of ‘Torture is wrong’ to be given by the set of worlds in which Torture is wrong is for that particular abstract object to adequately characterize certain relevant features (what I will call ‘linguistically relevant features’) of the mental state expressed...by an utterance of ‘Torture is wrong’ (2014, 138).

This isn't yet expressivist; it is merely consistent with expressivism. It will be expressivist if the mental states expressed by normative sentences are non-cognitive.

The restriction to linguistically relevant features of the expressed mental state is important. Expressivists think there are important differences between the mental states expressed by descriptive sentences and normative ones. The former are cognitive and the latter are not. But Carballo thinks that sets of possible worlds adequately characterize the linguistically relevant features of both cognitive and non-cognitive mental states. So for Carballo's proposal to work then, being cognitive must not be among the linguistically relevant features of the mental state expressed that make assigning a set of worlds to a sentence appropriate.

Carballo adds that the relevant features “will be those that play a role in explaining the communicative effect of an utterance of that sentence” (2014, 138). But this is still somewhat unclear. Metasemantic expressivists will likely think that normative sentences differ from descriptive ones in their communicative effects in that utterances of them have a greater tendency to motivate action. Since Carballo claims that sets of possible worlds adequately characterize both types of mental state, these communicative effects cannot be among the ones Carballo is talking about.

Due to these difficulties, it will be helpful to step back and ask what facts about belief contents, sentence meanings, and sets of worlds makes the last of these appropriate for modeling for the first two. One is that for every distinct belief content or sentence meaning, there is a distinct set of worlds.⁵² That snow is white and that grass is green are different contents and they are assigned different sets of possible worlds. Similarly, 'Snow is white' has a different meaning than 'Grass is green' and each sentence is assigned the same set of worlds as the belief expressed by it.

Sets of worlds are also good at representing entailment and consistency relations. That grass is green entails that it is green or white. This is mirrored by the set of worlds in which grass is green being a subset of the set of worlds in which it is green or white. Inconsistency between contents or sentences is modeled by the intersection of the relevant sets of worlds being empty. Contents or sentences being contradictions is modeled by their being assigned empty sets of worlds.

Finally, possible worlds can be used to model the effect of assertion on the mental states of conversational participants.⁵³ Asserting that Bob moved yesterday tends to bring about in one's interlocutors the belief that Bob moved yesterday. In modeling this fact it is helpful to have the same kind of object, a proposition, play the role of both the content of assertions—usually regarded as the semantic value of the sentence asserted—and of beliefs. If a single object plays both roles, we can model the assertion that P as an attempt to make one's interlocutors believe that P, where P is both the content of the belief and the semantic value of the sentence that

52 This is a famously controversial claim, especially when it comes to necessary propositions. The set of all possible worlds gets assigned to both the belief that $2+2=4$ and the belief that Mark Twain is Samuel Clemens even though these beliefs seem to have distinct contents. Since this problem with modeling contents with possible worlds is a general one with seemingly little to do with the difference between normative and descriptive sentences, I won't discuss it further here.

53 This is the feature that Carballo focuses on his discussion (2014, 139-140).

expresses it. We saw above that sets of worlds are independently appropriate models of several features of both belief contents and sentence meanings. Sets of possible worlds are thus good candidates for modeling some aspects of conversation as well.

Metasemantic expressivists will argue that normative sentences and beliefs as they understand them have all the features just mentioned that make sets of worlds appropriate models of content. Each distinct normative content can be assigned a distinct set of possible worlds. Normative beliefs and sentences have the same sorts of entailment and consistency relations as descriptive ones, and so these relations can also be modeled with set theoretic relations between sets of worlds. Finally, assertions of normative sentences tend to bring about changes in the belief in a similar manner to descriptive ones, making sets of possible worlds an appropriate model of normative assertion.

But it is important to note that this approach in an important sense assumes a solution to the Frege-Geach problem rather than presenting one. The Frege-Geach problem does not merely concern constructing an appropriate model of compositional meanings and entailment for normative sentences. It also requires establishing that normative judgments, as the expressivist understands them, have the relevant features to be modeled. Even if the expressivist has made it clear what mental state corresponds to the judgment that killing is wrong, it may be not at all clear what mental states correspond to the more complicated judgment that if killing is wrong so is letting die. And even if the expressivist succeeds in explaining what these mental states are, there remains the question of whether the standard inferences involving them are valid or rational. It only makes sense to assign these judgments possible worlds as contents if these questions can be addressed. Since there is no reason to believe that the metasemantic expressivist

is in a better position to answer them than the semantic expressivist, this undermines the claim that metasemantic expressivism solves the Frege-Geach problem.

The approach also doesn't help matters when it comes to Ridge's problem with the metasemantics of normative predicate extensions. For an expressivist, linguistic conventions will determine only that one apply 'wrong' to something in a world just in case one has a negative non-cognitive mental state concerning that thing.⁵⁴ They will not determine which non-cognitive state is semantically correct, because none of them is more semantically correct than the others. Accordingly, the claim that an act of torture falls under 'wrong' in a world just in case it fails to maximize utility in that world will not be privileged by the semantics over the claim that that act of torture falls under 'wrong' in all worlds. Expressivist metasemantics thus faces equally large problems in the possible worlds framework.

5. Deference to Semantics

If the foregoing is right, then metasemantic expressivism is not a tenable view. But there remains the question of how much of the work of metasemantic expressivists can be salvaged. In section three, we discussed two benefits of metasemantic expressivism over semantic expressivism. One was that it shows a proper deference from philosophers to semantic theory. The other is that it might allow expressivists to take advantage of already existing semantic machinery to explain the compositionality of normative sentences and thereby solve the Frege-Geach problem.

The second of these is not an issue for Ridge, because he explicitly doesn't think that metasemantic expressivism solves the Frege-Geach problem in this manner. Rather, he thinks that he must first “earn the right” to use first-order semantic accounts by explaining how his view

⁵⁴ An example of a “negative non-cognitive mental state” would be a normative perspective that ranks the thing lowly.

can make sense of unasserted normative expressions and valid arguments involving them (2014, 137). To do so, Ridge makes use of the hybrid mental states that he claims normative sentences express. The idea is to offload semantic complexity onto the representational component of the state. For instance, the judgment that if killing is wrong so is letting die will consist of a normative perspective and the belief that if the standards determined by that perspective rank killing lowly then they also rank letting die lowly.

Note that the non-cognitive component, the normative perspective, is identical to what it would be in the case of the simpler judgment that killing is wrong. All complications stemming from embeddings are handled by the representational portion of the normative judgment. This allows Ridge to give the standard accounts of the compositional meanings of sentences expressing these judgments and the validity of arguments containing them. There isn't space here to assess Ridge's solution to the Frege-Geach problem in full. But what I hope has become clear is that it in no way relies on his view being metasemantic. He could give the exact same solution while offering his proposal as a revisionary account of the meanings of normative sentences as being given by the (partially) non-cognitive states they express. So at least regarding the Frege-Geach problem, abandoning the claim that his proposal is a metasemantic one does Ridge no harm.

All that remains, then, is to see whether turning his proposal into a semantic one would amount to improper philosophical meddling in a scientific discipline. One reason to suspect that it would not is that in the very first chapter of his book Ridge himself revises a standard semantic analysis. He argues that we should replace Angelica Kratzer's account of deontic modals—expressions like 'should' or 'must' as used to make ethical claims—with his own account.⁵⁵

⁵⁵ See Kratzer (2012).

Ridge's account is similar to Kratzer's in that both appeal to two sources of context sensitivity for sentences containing deontic modals. For Kratzer, context affects both background knowledge and how high ranked various outcomes are. Ridge too has context affecting background knowledge, but he gives a different account of how it affects rankings of outcomes by way of his notion of standards. This allows him to make more fine-grained distinctions than Kratzer can, and he claims that such distinctions are needed by semantic theory. Whatever degree of deference is owed to semantics, Ridge must think that his revisions of the standard semantics for deontic modals are consistent with it.

Another reason to suspect that a significant degree of semantic meddling is acceptable is that linguists themselves appear to think so. Semantics is still a very young discipline, one which continues to harbor widespread debates about which foundational assumptions lead to the best analysis of various linguistic facts. In particular, semantics includes people working in the frameworks of Davidsonian truth-theoretic semantics, Montagovian model-theoretic semantics, as well as more recent dynamic approaches to semantics.⁵⁶ Important assumptions about how to explain semantic facts remain up for grabs, and it is not clear that an expressed mental state semantics—as would be offered by a semantic expressivist—cannot be one of the options.

Moreover, if you look to see what work is required for those presenting a new framework or new intra-framework approach to a problem, the answer looks promising for the expressivist. To be taken seriously, new frameworks and approaches simply need to not lose significant amounts of the ground gained by past ones. For example, if you are proposing a dynamic semantics, you cannot simply abandon having any account of the compositional meaning of 'and' and other basic sentential connectives. And unsurprisingly, one of the first tasks of advocates of

⁵⁶ See respectively Larson and Segal (1995), Heim and Kratzer (1998), and Groenendijk and Stokhof (1991).

dynamic semantics was explaining how it accounts for such meanings. Relatedly, if Ridge is right that his semantics for deontic modals is simply a more fine-grained version of Kratzer's, then we can expect that it will retain the advances made by Kratzer's approach.

If the above arguments work, then it appears that semantics is only owed the deference of making sure you can account for the same phenomena as the theories you are replacing. But to do so is just to solve the Frege-Geach problem. Any successful solution to this problem requires an explanation of the meanings of normative expressions in all the varied linguistic contexts in which they appear.⁵⁷ And such solutions will also have to be compositional and retain some account of entailment relations between sentences involving normative expressions. If a new proposal for these expressions can do all this, then it has fully discharged its obligations of deference to semantics. Thus, if Ridge has actually succeeded in solving the Frege-Geach problem without appeal to metasemantics, then he has no further reason to worry about excessively meddling with first-order semantics.

6. Conclusion

There are good reasons to seek metanormative theories that avoid the problems besetting the classic ones, but metasemantic expressivism does not appear to be a promising approach. It cannot succeed at its most basic task of explaining the semantic values of normative expressions. And even if metasemantic expressivism could do this, it would not have the theoretical benefits claimed for it. For one thing, such views don't help with the Frege-Geach problem since the assignment of standard semantic values to normative expressions presupposes that such a solution has already been found. For another, the worry that semantic expressivism does not pay

⁵⁷ Or, at least, most of the contexts they appear in that have already been given an account by rival semantic theories.

enough deference to semantics does not present major difficulties for these views beyond the Frege-Geach problem. If this is right, then metasemantic expressivism is not a real option for metanormative theory and it is no great loss that it isn't.

CHAPTER 3
REFERENTIAL NORMATIVITY AND THE METAPHYSICAL COMMITMENTS OF
NORMATIVE REALISM

1. Introduction

Imagine a community that speaks a language much like English except for some features of the words with the spelling and pronunciation of English's normative expressions ('good', 'ought', 'wrong', etc.) and the concepts they express. These expressions retain the same connection to action that their use bears in English. So when a speaker there predicates 'good' of a future situation that tends to count in favor of their planning to promote it. And the speakers admonish the actions of others, hoping to discourage them, by calling the acts 'wrong'. The difference between us and them is that they apply their seemingly normative expressions to very different things than we do.⁵⁸

For most predicate expressions, we would take this difference to be a reason to suppose that their expressions have different extensions from our own. But much recent discussion among metanormative theorists has argued that there are reasons to think our expressions and theirs must have the same extension. For some this is a matter of intuitions about the case itself, whereas for others it follows from certain normative-realist intuitions. If such a view is right, the normative roles of these expressions would have to fix a particular content in order to ensure that our normative expressions and the alternative ones share an extension. Following Matti Eklund, I'll call expressions and concepts whose normative role fixes their reference *referentially normative* (2017, 10).

58 Of course, there is a great deal of variety in our own application of normative expressions. People certainly don't always apply the word 'good' to the same objects. That said, there is also a great deal of overlap in these applications, so we can imagine this other group also has a great deal of overlap in their applications, but the two overlaps are significantly disjoint. In the body of the paper I'll spell out a more specific hypothetical to make this concrete.

The notion of referential normativity has the potential to reconfigure metanormative debates. Eklund's arguments suggest that referential normativity is what more radical normative realists find missing in standard forms of naturalistic normative realism. Concurrently, Robbie Williams has argued that referential normativity requires neither revisionary metaphysics nor revisionary philosophy of language. He claims instead that it falls out naturally from an independently motivated metasemantic view and that this view allows for natural properties to be the referents of normative concepts. Putting these two together presents us with the hope that a metaphysically conservative theory could fulfill the desires of even those normative realists that have balked at hitherto advanced versions of normative naturalism. If this hope is met, we could in a sense replace the metaphysical concerns of metanormative debates with new ones from the philosophy of language. This paper aims to dampen the cause for enthusiasm at which I just now gestured. I will be arguing that the moral realist intuitions described above do require normative properties to be in some sense special, and that in the absence of this referential normativity fails to satisfy those intuitions.

On the way to arguing for this, I will first explain the notion of referential normativity and the two types of arguments that support the view that normative concepts have this feature. Then in section three I cover Williams' proposed theory of content that he claims implies that normative expressions are referentially normative, noting some worries for his theory even if it avoids my later objections. In the following section I cover reasons to think that a metaphysically conservative construal of referential normativity along the lines Williams provides fails to fulfill the normative realist motivation for referential normativity. More tentatively, I then speculate on some potential ramifications my arguments have for traditional arguments surrounding moral-twin-earth-like cases.

2. Why Referential Normativity

The notion of referential normativity relies on that of a concept's having a normative role. A number of different accounts have been offered of what normative roles amount to, but they all in some way concern the connection between applications of the concept in question and action or assessment of action.⁵⁹ Some notions of normative role focus on the connection between applying a concept and having some reactive attitude toward what it is applied to. Williams, for instance, illustrates the idea of normative role by positing that applying the concept of wrongness to an action causes one to blame the agent for performing that action (2018, 42).⁶⁰ Other notions of normative role focus on connections between applying the concept and deliberation.⁶¹

This paper will not assume any account of normative role is correct, but for concreteness it will help to have one in terms of which to frame the debate. For that purpose, I will use a deliberation account that is essentially a simplified version of the one offered by Gibbard (2003), and the central normative concept we will apply it to is the all-things-considered ought concept. I will say that judgments of the form 'I ought to P' tend to cause their maker to plan to P. Conversely, judgments of the form 'It is not the case that I ought to P' cause their maker to refrain from planning to P. Of course, we also make third-person ought judgments, such as the judgment that Sally ought to take the route 30 bus. We can understand these judgments as causing hypothetical plans to act in the suggested manner were you to be in the same circumstance as

59 I will only be assuming that there is some consistent and distinctive causal connection between the application of the concept and the states caused by doing so according to its normative role. The paper is not taking a stand on the internalism debate about normative judgments. See Williams for some discussion of this point (2018, 42-43).

60 To some, this might seem like an implausible connection between normative judgments and blame. Perhaps rather than those judgments causing one to blame, the judgments and the blame have a common cause in prior sentiments of the judge. Such a view would be consistent with referential normativity as long as normative concepts differ from non-normative ones in how they are causally linked to deliberation or assessment of action, such that we could identify distinctive normative roles for these concepts. In this paper, I will often follow Williams in speaking of our normative judgments causing certain mental states, but the reader should feel free to substitute this for the broader idea that the two are distinctively causally linked.

61 Examples of this kind of account can be found in Wedgwood (2001), Gibbard (2003), Scanlon (2007), and Ridge (2014).

Sally.⁶² I will refer to concepts that have this normative role as “ought-concepts”. If you find this account of normative role unsatisfactory, feel free to substitute your own. It should not affect any of the arguments.

So why have some theorists argued that the notion of normative role plays such an important role in fixing content? The main arguments offered relate to versions of the “moral twin earth” thought experiment.⁶³ Suppose we have some community, K, who possess an ought-concept, *ought_K*, that they apply to actions that conform to the categorical imperative. So, in response to judging that some action conforms to the categorical imperative they tend to judge that they ought_K to perform that action, and in response to judging that they ought_K to perform an action they tend to form a plan to perform it.⁶⁴ And suppose there is another community, U, with an ought-concept, *ought_U*, that they apply to actions that maximize utility. So, in response to judging that some action maximizes utility they tend to judge that they ought_U to perform that action, and in response to judging that they ought_U to perform an action they tend to form a plan to perform it. There are two ways in which this sort of example has been taken to argue in favor of the referential normativity of ought-concepts.

The first is by way of intuitions about the case. Say a member of U judges that one ought_U

62 Other normative concepts would have some distinct relation to forming plans as part of their normative role. One might propose that judging a future situation to be good tends to cause one to be more inclined to plan to bring it about *ceteris paribus*.

63 See Horgan and Timmons (1992) for the original discussion under that name and Hare (1952) for a precursor of this sort of example. The example I go on to offer differs in that it concerns concepts rather than linguistic expressions and that it concerns normative concepts in general rather than specifically ethical ones. Dowell (2016) gives a more recent discussion that offers a case against the first version of the argument for referential normativity that I present below.

Also, note that the original moral twin earth argument was not intended to show that normative expressions are referentially normative. Instead it was intended to establish a form of expressivism about morality. But it should be clear how if one holds fixed a commitment to normative claims and concepts being descriptive, it leads instead to referential normativity.

64 Plausibly, a number of different actions available to any individual will conform to the categorical imperative, and so this constraint will not single out a unique course of action. To make the example more realistic, feel free to add additional constraints that might pick out a determinate course (or set of courses) of action together with this one.

to lie to save the lives of others and a member of K judges that one ought_k not to lie to save the lives of others. It seems intuitive to some that the ought-concepts in these judgments share a content and so there is some shared content, P, such that members of U judge that P and members of K judge that not-P. Alternatively, you might argue that what is directly intuitive is that the K person disagrees with the U person when each makes their judgments. But if their ought-concepts differed in content, then it would appear that they do not really disagree with each other and are simply making judgments about different features of the act of lying to save the lives of others. So the claim that their judgments genuinely disagree also supports the idea that their ought-concepts refer to the same thing. Of course, this would just be a single case of two ought-concepts picking out the same thing despite being consistently applied to different actions. Extending it to a full case for the referential normativity of ought-concepts relies on the claim that the other details of the case do not matter to the intuitions on which the argument relies.

Recently, Eklund has proposed a distinct way in which the above sort of scenario motivates the referential normativity of ought-concepts. He argues that there are certain moral realist intuitions that would fail to be satisfied if normative expressions were not referentially normative. He introduces the intuitions as follows:

If you are a realist of the kind I want to consider, you want to say that Bad Guy – someone who does bad things and is motivated by bad desires – is somehow objectively out of sync with how to conduct one’s life: in no way are our differences with Bad Guy merely a disagreement in taste, or a matter of having different desires. Introducing a label, you are an ardent realist if you are a realist and this is part of your motivation. You think any acceptable normative realism should deliver the result that your difference with Bad Guy comes out as something more than a mere disagreement in taste. You think any

acceptable normative realism should deliver the result that there is something important that one of you (presumably you) objectively gets right and the other (presumably Bad Guy) objectively gets wrong. There are some facts that Bad Guy fails to properly appreciate (2017, 1).

In similarly metaphorical terms, we might say that the ardent realist believes that the world itself favors some views on normative issues, presumably not those of Bad Guy. Both this characterization and those in the quote fail to give a precise statement of what it is to be an ardent realist. That said, I share to some extent the intuitions gestured at by them and also the (soon to be discussed) intuitions about how these intuitions relate to moral-twin-earth-style examples. These intuitions also appear to be widely shared enough that they are worth discussing as something more than a personal idiosyncrasy. So despite lacking as straightforward a definition of ardent realism as would be desired, for this paper I will take it that shared intuitions about the relevant cases provide a solid enough handle for it to be productively discussed.

Now suppose that ought-concepts are not referentially normative, and in particular that our ought-concept has a different extension from that of K.⁶⁵ They correctly judge that one ought_k not to lie to save the lives of others, and we correctly judge that one ought to lie to save the lives of others. If this were the case, it seems that ardent realist intuitions could not be satisfied. In what sense would the world favor our way of doing things over theirs? Sure, we are objectively correct about what one ought to do. But equally so, they are objectively correct about what one ought_k to do. In fact, the citizens of K need not have any false beliefs supposing the case works this way. Any seeming dispute between us and them would actually be a case of us and them talking past each other.

⁶⁵ See Boyd (1988), Railton (1986), and Jackson (1998) for examples of theories that would allow for this possibility.

Eklund argues that the only way that ardent realism could be correct in such a case would be if there were an answer to some further question beyond the question of what one ought to do (2017, 22). This question would be along the lines of which ought-concept is the correct one to use to guide one's actions. But problems arise immediately when we reflect on this question. Presumably, in this question by “correct” we mean normatively correct.⁶⁶ But using our own concept of normative correctness, surely our own normative concepts are the correct ones. And the people of K use the normatively correct_k concepts, unlike us. Similarly, surely we use the normative concepts that ought to be used, but they use the normative concepts that ought_k to be used. So on all the natural ways in which we might try to gesture at the further question the ardent realist would need answered, the question ends up trivial. If you phrase the question in terms of our normative concepts the answer will favor our normative concepts, and if you phrase it in terms of theirs it will favor their normative concepts. Such trivial further questions fail to make the scenario one that would be acceptable to ardent realists, because the answer to such questions in no way suggests that either us or members of K are objectively getting things wrong. This suggests that any further question that would be acceptable to ardent realists would in fact be ineffable and unthinkable.⁶⁷

In sum, ardent realists are in trouble if normative concepts are not referentially normative. They can claim that there would be some further question about what normative concepts to use in such a case, but this further question would have to be one we are not even capable of thinking. But if they deny that there is such a further question, then it seems that ardent realism

66 That is, we are not simply concerned with correctness relative to some specific set of potentially arbitrary rules, but rather with the very same kind of correctness we are concerned with when asking what we ought to do.

67 See Eklund for an expanded argument for this conclusion (2017, 22ff). In the first place he is arguing here that the question would be ineffable, but since I have phrased this discussion in terms of normative concepts (conceived as vehicles for thought) rather than normative expressions as Eklund does, the unthinkable conclusion would follow more immediately.

must be false, since then there would be no sense in which we could say that the world somehow favors our way of doing things over those of K, or *vice versa*. In light of this, those with ardent realist intuitions have a reason to support the referential normativity of normative concepts. Importantly, these reasons apply independently of those offered earlier, concerning intuitions about content and about disagreement. We have no reason to believe that arguments against the force of those intuitions, such as those offered by Dowell (2016), also undermine Eklund's motivation for referential normativity.

It is natural at this point to wonder what would have to be the case in order for normative concepts to be referentially normative. In particular, you might wonder what accounts of what grounds the content of concepts would allow for this possibility, and what normative properties would have to be like to have this feature according to those accounts. On this question, Eklund is silent, though in places he tentatively suggests that there would have to be something special about normative properties for them to be referentially normative (2017, 133). In the next section, we will look at an account that goes into detail on these issues.

3. Robbie Williams' Account of Referential Normativity

In his “Normative Reference Magnets”, Robbie Williams argues that normative concepts are referentially normative, or “stable” as he puts it, though he is motivated by the first argument described above rather than Eklund's (2018, 41). Two key differences from Eklund that will matter in this paper are that he gives an explicit metaphysics of representation—an account of in virtue of what representational states have their contents—from which it follows that normative terms are referentially normative, and he argues that there does not have to be anything

metaphysically special about the properties ascribed by referentially normative predicates.⁶⁸

Williams' metaphysics of representation is a form of interpretationism, which states that “The correct interpretation of an agent *x* is that one that best accounts for *x*'s dispositions to act in the light of the course of experience *x* undergoes” (2018, 46). For his account, best accounting for these dispositions means best rationalizing them, that is, making their bearer maximally rational. In sum, an agent's concepts have the contents that would make their bearer most rational, according to Williams' account.

To spell this out more concretely, consider the vast number of possible assignments of referents to a person's concepts. Almost all of the assignments in logical space will be way off the mark, perhaps having the concept of penguin pick out the set of supernovae or the concept of president pick out the singleton containing only a particular tooth on a particular stegosaurus.⁶⁹ And these way off the mark assignments will generally be such that if they actually characterized what someone's concepts refer to, that person would be extremely irrational. The bases on which I form beliefs about penguins as well as the things I in turn infer from my penguin beliefs would be highly irrational on the supposition that my penguin concept picks out supernovae. From these considerations we can see that there is at least a partial ordering on reference assignments based on how rational the concept possessor is given that reference assignment. Williams' view is that the correct reference assignment is determined by what reference assignment is at the top of that ordering.⁷⁰

68 As Williams describes it, a metaphysics of representation also specifies what sorts of representational states there are and what kinds of things are subject to interpretation, but it is the claim about what makes something have a certain content that will matter in this paper (2018, 45-46). In what follows, I will generally also use “metasemantics” to refer to the relevant aspect of a metaphysics of representation, even though Williams is in the first place concerned with mental rather than linguistic representation.

69 Note that Williams' account is actually designed to return properties as the referents of predicate concepts, not the extensions of those properties. For simplicity, I'll talk mostly about concept extensions in this paper, but this detail of his actual account could matter to some of the arguments in section four, as will be noted.

70 Presumably, indeterminacy or ties at the top of the ordering will result in indeterminacy in reference.

There are several dimensions of rationality that will be relevant here. First there is the familiar distinction between epistemic rationality—governing how well supported a person's beliefs are—and practical rationality—governing the extent to which a person acts for good reasons. Second, there is the distinction between structural rationality, which governs features of the rationality of mental states without any reference to specific contents of those mental states, and substantive rationality, which does make such reference. To illustrate, one constraint of structural epistemic rationality might be to always believe P whenever one believes $P \& Q$. And a potential constraint of structural practical rationality is that if you want it to be the case that P and doing Q is the only way to bring it about that P , then you should do Q unless there are defeaters. For each of these examples, we can assess whether a person abides by this rationality constraint without assessing whether any particular atomic mental state is rational independently of the other mental states the person possesses.

Constraints of substantive rationality, on the other hand, do single out particular such mental states as being more or less rational on their own. For instance, it plausibly follows from substantive epistemic rationality that it is more rational to believe that all emeralds are green than that all emeralds are grue even if each belief would be perfectly structurally rational given the totality of one's beliefs.⁷¹ And a plausible constraint of substantive practical rationality is that you not be future Tuesday indifferent.⁷² Williams argues that interpretationism needs to consider both structural and substantive rationality in order to avoid certain skeptical paradoxes that would

71 To be grue is to be either green and first observed before time t (say, January 1st, 2019), or blue and not first observed after time t (Goodman 1983)

72 Future Tuesday indifference involves not attaching any weight to what one experiences on future Tuesdays, not even if it means accepting horrific pain on a future Tuesday in exchange for a tiny benefit on some other day Parfit (1984). Note that when that Tuesday comes around it will be a present Tuesday, not a future one, and so you will regard the pain while you are experiencing it as just as bad as a pain on any other day. See Street (2009) for an argument that avoiding this is not actually a constraint on practical rationality.

have the result that reference is radically indeterminate.⁷³

Finally, Williams' preferred way of characterizing these species of rationality is in terms of being reason-responsive. So, the degree to which a person is acting rationally is determined by the extent to which they are responding to the reasons they have, whether these are structural reasons to believe something or substantive reasons to act a certain way. Using this account of rationality, we could if we wished give a metasemantics with the same results as Williams' without making any appeal to facts about rationality. That is, instead of accounting for reference in terms of rationality and then accounting for rationality in terms of reason-responsiveness, we could skip the “middleman” and just account for reference in terms of reason-responsiveness. Since the application of Williams' account as is requires somewhat controversial commitment to there being substantive constraints on practical and epistemic rationality, the altered account might be more appealing to those who reject those commitments. That will be so at least insofar as substantive constraints on reason-responsiveness are more appealing than ones on being rational.

On the view we have been characterizing, if a concept's referring to some property would make its bearer most reason-responsive along a particular dimension of rationality, then *ceteris paribus* that fact makes the concept refer to that property. That is, this fact will incline the concept to have that reference; but if there are countervailing reasons-facts concerning its use, the concept might still have a different reference. For ease, we can refer to this situation as one where there is some amount of “metasemantic pressure” for the concept to refer to what would make its user most reason-responsive along some particular dimension. Or comparatively between two choices of referent, we could say there is metasemantic pressure for the concept to

⁷³ See Williams (2018, 47-49) and (2016) for a paper entirely on this issue.

refer to whichever of the two makes its user more reason-responsive along some dimension.

With this background in hand, we can now explain Williams' argument for the referential normativity of normative concepts. First, suppose for concreteness that as far as substantive practical rationality is concerned what we have most reason to do is to maximize utility. Now suppose that someone uses a concept, *ought_C*, with the normative role of our concept of the all-things-considered ought. And recall that for our purposes this means that judging that one *ought_C* to do something causes one to plan to do that thing. Plausibly, if we have reason to maximize utility we also have reason to plan to maximize utility.⁷⁴ So, a user of *ought_C* will be more reason-responsive if the concept refers to actions that maximize utility than if it refers to other actions *ceteris paribus*.⁷⁵ Suppose this were not the case and that *ought_C* did not refer to such actions. Then the person using it would fail to be maximally reason-responsive by taking something other than the fact that something tends to maximize utility as a reason for bringing it about.⁷⁶ Since on the view we are supposing what a concept refers to is determined by what would make its user most reason-responsive, this means that there is metasemantic pressure that inclines *ought_C* to refer to the maximization of utility, or in other words, to what one ought to do.

The above goes for anyone that uses a concept with the same normative role as that of ought. The pressure inclining that concept to refer to the same thing, what maximizes utility, will apply regardless of what typically causes the concept to token, whether that has to do with

74 I will later come back to what would result from denying this claim.

75 Precisely speaking, it should refer to the relation between an agent and an action that holds if and only if the agents' performing that action would maximize utility. For convenience, I'll typically stick to the simpler formulation.

76 There is more to be said here that I am skipping over. For instance, any property that picks out a subset of the things you ought to do will also have the feature that you are only reason-responsive if applying the concept also inclines you to plan to do the action it is applied to. So nothing I have said would determine that the user of concept C is made more reason-responsive by C referring to what one ought to do than by it referring to any property picking out a subset of the the things one ought to do. To rule this out, we would also have to stipulate that judging that something is not-C tends to incline one to not planning on doing things that bring about that feature. For a discussion of this and some related issues, see Williams (2018, 50-53).

utilitarian concerns, Kantian ones, or something entirely different. So we can see why Williams' metaphysics of representation would potentially allow for referential normativity. The normative role of a concept on its own exerts metasemantic pressure for that concept to refer to a particular thing. So all concepts that share that normative role will face metasemantic pressure to refer to the same thing.

That cannot be the full story, however, since all we have established is this: according to Williams, normative concepts will all have the same reference *ceteris paribus*. An ought-concept will be such that its referring to what one ought to do makes its user more reason-responsive in one respect, but that does not mean that all things considered it makes its user most reason-responsive. For there might be other dimensions of reason-responsiveness along which a different assignment of reference to the concept would increase its bearer's reason-responsiveness. And this is not just an abstract possibility but arguably the case.

Consider a member of the community K, who consistently applies their ought-concept to actions that conform to the categorical imperative on the basis of their so conforming. If what one ought to do as a matter of fact is to maximize utility, then—depending on what the correct account of epistemic reasons is—it may be epistemically unreasonable to infer that one ought to perform some action on the basis of the fact that it conforms to the categorical imperative. That an act conforms to the categorical imperative is clearly not evidence that it maximizes utility. If this claim about epistemic reasons is correct, then assigning the set of actions that conform to the categorical imperative as the referent of this person's concept will maximize their epistemic reasonableness with respect to its application.

We would thus have two sources of metasemantic pressure pulling in opposite directions for members of K. Assigning the property of conforming to the categorical imperative as the

referent of *ought_k* would increase the epistemic reasonableness of how the community members form beliefs about what they ought_k to do (on certain views of epistemic reasons). This applies metasemantic pressure for their concept to refer to the property of conforming to the categorical imperative. On the other hand, assigning the property of maximizing utility as the referent would increase the practical reasonableness of how community members regulate their actions using judgments about what they ought_k to do. This applies metasemantic pressure for their concept to refer to the property of maximizing utility.

Williams responds to this issue by claiming that in most actual cases the metasemantic pressure from considerations of practical reasons outweighs the pressure from considerations of epistemic reasons (2018, 61-64). In fact, falsely taking 'ought' to pick out what conforms to the categorical imperative might not require any epistemic unreasonableness at all. If a person arrives at their judgments about what they ought to do by carefully assessing trade-offs between intuitions about normative cases and discussing matters with capable peers, then they could plausibly be epistemically blameless even if they arrive at the wrong conclusion. After all, surely it is reasonable to form beliefs about what one ought to do by engaging in reflective equilibrium starting from largely accurate normative intuitions. Williams thinks that the normal case of a person having false beliefs about what they ought to do has enough in common with this sort of scenario that there will in general be little or no metasemantic pressure on their ought-concept to pick out what they (falsely according to Williams) believe they ought to do.

But as Williams acknowledges, the above claims only apply to people who form their ought judgments similarly to how we do. Imagine a group whose members immediately and indefeasibly token their ought-concept in response to things that conform to the categorical imperative. Then the case becomes much stronger that assigning the property of maximizing

utility as the referent of that concept renders them epistemically unreasonable. And, depending on how much of an epistemic internalist one is, that case may also seem quite strong if the group simply has highly unreliable normative intuitions that serve as the starting point for their theorizing. Insofar as we want to take that group's ought-concept to refer to the same thing as our ought-concept for the reasons described in the previous section, this is a problem for Williams' account. For it seems that that account could just as easily result in their concept referring to something else (say, actions that conform to the categorical imperative) or it being referentially indeterminate between two options.

How great a problem this is for Williams' account of referential normativity will depend on one's reasons for seeking such an account. If the reasons are ardent realist ones, then the problem is significant. Ardent realists would be at least as concerned that a Bad Guy with unshakable intuitions entirely in line with their morality be objectively incorrect as someone deviating more mildly from us. So Williams' view appears incapable of supplying the ardent realist with the level of normative objectivity they seek. On the other hand, if one thinks normative concepts are referentially normative on the basis of intuitions about moral-twin-earth-like cases, then the above issue is only a problem for Williams if those intuitions apply to the type of agent being discussed as well.

My own take is that this is a significant problem for Williams on both motivations for the view, but I'll set it aside to focus on different issues in the remainder of the paper. This is for two reasons. First, I take the problems raised in the later sections to be more general and to present problems for all accounts of referential normativity even if they lack this feature of Williams' own account. Second, it is possible that a defender of Williams' view could argue that even with there being some metasemantic pressure for ought-concepts to refer to things other than what one

ought to do, that pressure is always overridden by the pressure supplied by one's substantive practical reasons. Of course, the burden would be on that defender to explain why this would be the case.

With that in mind, we will continue to the main problems I'll be considering for the account just described, starting with some objections briefly raised against Williams by Eklund himself.

4. Ardent Realism and Special Properties

4.1. Eklund's Objection

Eklund objects that it is not enough that we have intuitions of referential normativity and that the internal features of our language use allow for it to establish that normative expressions are referentially normative (2017, 129-134). In addition to this, the world must cooperate as well. Consider his example of *gidplissing*, which people apply to things that are taken to be pleasing to God. We can imagine another community that uses a term that shares this feature with 'gidplissing' but that has a starkly different conception of what is pleasing to God. Would these two terms have the same reference? If God exists and is such that both communities succeed in referring to God with their respective terms, then presumably the terms do have the same reference. They both pick out whatever things are in fact pleasing to the God that actually exists. At least one of the communities will have been applying the term largely to things that don't in fact please God. But these applications of the term bear significantly less metasemantic weight than God's actual preferences, since those are what the term is centrally designed to track.

But what if God doesn't exist? Then it is hard to see why 'gidplissing' and 'gidplissing*' should be taken to refer to the same things. Perhaps we might want to say that both terms will

fail to refer to anything in such a case, so in that sense they would still have the same reference. But it seems more appealing to say that each term would still refer to the thing that best satisfies the theories surrounding the term. A central aspect of that theory is false for both groups, but the specific claims about what objects are gidplissing or gidplissing* still remain and can plausibly be satisfied by some actually existing properties. Such a stance would accord with treating 'gidplissing' as Lewis argues we should treat theoretical terms (1970).

What this example illustrates is that whether certain aspects of a term's conceptual role determine a unique reference can depend on the way the world is. If God exists, then plausibly any term whose conceptual role relates in the manner of 'gidplissing' to what pleases God will have the same reference. If God doesn't exist, then other features of the conceptual role of 'gidplissing'-like terms will determine their reference, allowing different such terms to diverge in their reference. And even if one prefers to interpret this example as one in which 'gidplissing' and 'gidplissing*' fail to refer to anything, then the example would still show something of interest. Namely, it would illustrate the more familiar point that whether a term refers at all depends on features of the world beyond the term's use.

Eklund argues that things might be similar for normative terms. These terms have a conceptual role that includes their normative role. Whether that aspect of their conceptual role determines a unique reference for any other term that shares it may depend on what the world is like, just as in the case of 'gidplissing'. Specifically, the idea is that even if we accept Williams' metasemantic picture, that will only establish the referential normativity of normative terms if robustly normative properties in fact exist.

What does this latter requirement amount to? Eklund's own view is that a property is normative just in case it is capable of being referred to by a referentially normative predicate

(2017, Chapter 5). But this does not tell us particularly much about what the world has to be like in order for it to contain normative properties. This is especially so given that Eklund does not endorse a particular metasemantic view to explain how referential normativity could be possible in the first place. Williams does though, and it is not immediately clear why it would have the metaphysical burdens that Eklund suggests. What it takes for a term to be referentially normative on Williams' view is just that it is used with a normative role and that (according to our toy example of a normative role) there are substantive practical reasons to plan to do certain things (namely maximize utility, according to a separate toy account). It follows from the account that if it is the case that there is a reason to plan to do what maximizes utility, ought-concepts will be referentially normative and pick out whatever course of action maximizes utility.

It is not immediately clear what part of this account requires anything metaphysically weighty to be the case. The core assumption about the world is just that we have a reason to plan to do what maximizes utility. If 'reason' picks out a natural relation, say by being identical to the relation between individuals and the course of action that would maximize utility if taken by that individual, then this assumption does not require a belief in any objects or properties that we are not already all committed to, at least assuming 'utility' can be naturalistically explained. Admittedly, it does require a commitment to substantive practical reasons that people have independently of any of their contingent desires, and this is very controversial. But it is a controversial claim of first-order ethics which does not necessarily commit us to any unwanted metaphysics. Because of this, it is not clear that Eklund's objection has any force on its own. While there might be other reasons to think only a special kind of property could be the referent of a referentially normative term or concept, nothing in Williams' own framework calls for it *prima facie*. Ultimately, I agree with Eklund that it does, but further argument is needed, which I

aim to provide in the next section.

Before moving on to my own arguments for this conclusion, I will say more about what I mean by a property being “special”. We earlier saw that certain metanormative theories would have the result that something's having the property of goodness would not be enough to satisfy the ardent realist's desire for the world to objectively settle what is choice-worthy. After all, that thing might lack the property of goodness*, and we haven't provided a reason why goodness should guide action rather than goodness*.⁷⁷ As a first stab, for a property to be *special* in the relevant sense is just for it to satisfy the desires of the ardent realist for a feature that makes something objectively choice-worthy in the way that those metanormative theories without referential normativity fail to do.⁷⁸

Traditionally in metanormative debates, this specialness has been taken to involve a property's being non-natural, but I am choosing not to frame the issue in those terms. This is partly because it is unclear whether a non-natural property is capable of being special in the required manner.⁷⁹ I also do not want to rule out that a natural property could be shown to be special to the satisfaction of an ardent realist, though I do not know how this could be done. The claim that ardent realist intuitions can only be satisfied by special properties means that securing referential normativity on its own cannot do so, and that additionally the objective choice-worthiness of the proposed normative properties needs to be defended. In short, ardent realists need metaphysics in addition to philosophy of language. What exactly that metaphysics is lies outside the scope of this paper.

⁷⁷ I am here ignoring the ineffability worries discussed in section two.

⁷⁸ This isn't to say that there being a special property is the *only* way to satisfy these ardent realist desires. The goal of this paper is to determine whether instead non-special properties referred to by referentially normative concepts would also suffice to satisfy them.

⁷⁹ See Dasgupta (2017) for an argument that non-natural properties could not play this role.

4.2. Can Only Special Properties Be Referred to by Referentially Normative Concepts?

We can ask two questions about whether we should think normative properties must be special for Williams' account. First, could any non-special property play the metasemantic role that Williams posits for normative properties? Second, if so, would the resulting account of normative language and metaphysics satisfy ardent realist intuitions? In this section, I suggest a negative answer to the first question, though I'll highlight some reasons my arguments are inconclusive. The section after this one will present arguments for a negative answer to the second question.

Let's look more closely at Williams' account of what grounds the reference of ought-concepts. The grounds are that this reference maximizes the reason-responsiveness of the concepts' users which is in turn grounded in the facts that (1) these concepts have the ought normative role and (2) what we have most practical reason to plan to do is to maximize utility (along with the fact that there is no countervailing metasemantic pressure). We saw in section three how (1) and (2) together establish that assigning what maximizes utility as the referent of ought-concepts makes their bearers most reason-responsive (ignoring the issues that have been set aside).

Note that the concept of reason that figures in this explanation is itself a normative concept. When one judges that one has most reason to P, that just as plausibly tends to cause one to plan to P as when one judges that one all-things-considered ought to P. And one could say similar things for other possible notions of normative role. From this, it immediately falls out of Williams' account that the concept of practical reason is referentially normative and that specifically it refers to what maximizes utility.⁸⁰ When we substitute what it is to have the most

⁸⁰ As before, to be more careful the concept refers to a relation that holds between individuals and actions when that action is what maximizes utility. Also, presumably the concept of having most practical reason to P does not

practical reason to P into (2), it becomes: what maximizes utility is to plan to maximize utility. With that in mind, in this section I will be arguing that facts (1) and (2) seem too trivial to ground the reference of ought-concepts.

Now, there is a sense in which these facts are not trivial at all. We already saw in section three that it is far from trivial that there will be no countervailing metasemantic pressure to override the metasemantic pressure of a concept's normative role. But we are setting aside that issue for now. It is also far from trivial that planning to maximize utility itself maximizes utility. Perhaps utilitarianism is too flexible a theory in human hands and people who adopt it will inevitably twist their assessments of what maximizes utility to fit immoral ends. In such a case, it may be that planning to follow strict deontological rules would maximize utility. Or perhaps people have a tendency to deviate from their plans in a self-serving direction when the moment to act arrives, and so it better maximizes utility to plan to be even more self-sacrificing than would actually maximize utility. Similar points could be made for other ethical theories and for other accounts of normative role. It is easy, for instance, to come up with reasons why it might not maximize utility to always blame others for failing to maximize utility.

But the point is that for Williams' account to work, this issue and the earlier one must not arise. There must be no deviation between what we have most reason to do and what we have most reason to plan to do. Otherwise the account would simply get the extensions wrong. To spell this out, recall that we are assuming for concreteness that what one ought to do is to maximize utility. So Williams' account is supposed to have the result that ought-concepts refer to what maximizes utility, and to do so because we have most practical reason to plan to maximize

have its content grounded all in one step, but instead has its content grounded first in the concepts of most and of practical reason, which in turn would each be grounded in the manner described by Williams' account. It should be clear from how my argument works that these details do not affect matters.

utility. But as we just saw, we have most practical reason to plan to maximize utility just in case planning to maximize utility maximizes utility. So unless planning to maximize utility maximizes utility, ought-concepts will not refer to what, by stipulation, one ought to do. And just as Williams' account needs that to be the case, it also needs there to be no countervailing metasemantic pressure.

Once we set aside these issues that would apply to any view of what ought-concepts refer to, we are left with nothing but the claim that we have most reason to maximize utility as the grounds for ought-concepts referring to what maximizes utility. *Prima facie*, this seems a bit odd. How can the fact that what maximizes utility maximizes utility ground the reference of ought-concepts? The fact seems too trivial to do that work. Now, one might note that the ground is actually that we have most practical reason to maximize utility. It happens that the concepts of having most practical reason to P and of P being the utility maximizing action are coextensive. But that leaves open the possibility that the properties ascribed by these concepts are distinct. And it is plausible that for distinct though coextensive properties, P_1 and P_2 , important facts can be grounded in the fact that whenever an object has P_1 it has P_2 . If so, there would not necessarily be much oddness in the reference of ought-concepts being grounded in the fact that we have most practical reason to maximize utility.

However, the point brought up in the above paragraph cannot help Williams, as will become clear after rehearsing some earlier points. Because what ardent realists are in the first place concerned with is that what one ought to do in a given circumstance is the same as what one ought_x to do in it for any other ought-concept, I've generally framed issues in this paper in terms of all ought-concepts having the same extension. But, as I noted in an earlier footnote, Williams' account is meant to explain not just the extensions of concepts but what properties they

pick out. So his account will have the result not just that practical-reason-concepts and ought-concepts are coextensive but that they pick out the very same property. The property of having most practical reason to P just is the property of P being what one ought to do which just is the property of P being the act available to one that maximizes utility. So while it is true that the fact that all P₁'s are P₂'s can do explanatory work even if P₁ and P₂ are coextensive, that point will not help here.

I do not wish to frame this as a fully knockdown argument, as certain metaphysical assumptions could help Williams escape these issues. For one thing, one could adopt the view that the fact that x has P₁ can be distinct from the fact that x has P₂ even if P₁ and P₂ are identical. And one could further claim that, although the fact that what has P₁ has P₁ cannot ground anything interesting, the fact that what has P₁ has P₂ can in certain cases do so even when P₁ is P₂.⁸¹ These assumptions allow for the possibility that the fact that what we have most reason to do is maximize utility grounds the fact that ought-concepts refer to what maximizes utility.

We could try to get around this issue by arguing that grounding-like claims should not be expressed as relations between facts. Instead we can make use of a sentential operator, 'because', and express the above claim as being that ought-concepts refer to what maximizes utility because we have most practical reason to do what maximizes utility.⁸² This latter claim makes no mention of facts and so the proposal about facts considered in the last paragraph is not immediately relevant to it. That said, one could make a comparable proposal to the one made above about our new claim and about the 'because' operator. This would be the claim that it can be the case that A because all P₁'s are P₂'s even though it cannot be the case that A because all P₁'s are P₁'s, for some

81 See Rosen (2010) for a view of facts and grounding that allows for this. Thanks to Alexander Kocurek for the pointer and to Harold Hodes for emphasizing the importance of this issue.

82 See Sider (2012) for a discussion of how one can avoid ontological commitment to entities such as facts by using operators in this fashion.

interesting fact, A, and for identical properties P_1 and P_2 . Or one could simply insist that grounding-like claims should express relations between facts, and so we should not replace the earlier claims with these new ones using sentential operators.

It should be clear from the preceding discussion that whether or not Williams' account succeeds will depend on a number of difficult-to-adjudicate claims about the nature of facts and grounding. I cannot claim to have settled those metaphysical disputes against Williams, but at the very least I have demonstrated that there are some additional costs to Williams' view in that accepting it commits one to certain metaphysical theses that one might otherwise wish to have avoided. This at least partially undermines the idea that Williams has shown that non-special properties can be the referents of referentially normative concepts.

4.3. Is Referential Normativity Without Special Properties Good Enough?

Now let us turn to the second question from above: would Williams' account satisfy the ardent realist assuming non-special properties are the referents of ought-concepts? In this section, I argue that the account would not do so, because Williams' account does not contain anything that would justify the normative significance of the referents of ought-concepts. Specifically, I'll claim that considering an alternate metasemantic framework to Williams' should raise worries for ardent realists that are similar to those raised by considering a community whose normative expressions refer to different things from our own.

To support this claim it will help to consider an “alternative (alt) Williams” that proposes the same general metaphysics of representation as Williams while believing that what we have most reason* to do is to conform to the categorical imperative.⁸³ Kantians, after all, have just as

⁸³ *Reason** is just alt Williams' concept with the normative role of our reason-concept, which for the purposes of this paper one cannot assume is the same concept or has the same reference as our reason-concept. For

much reason as utilitarians to want a metasemantic account that avoids paradoxes of widespread referential indeterminacy. Alt Williams will propose the alternative theory of reference that concepts refer to what make their users most reason*-responsive. In particular, he will conclude that ought-concepts refer to what conforms to the categorical imperative because we have most reason* to plan to conform to the categorical imperative.

According to Williams, alt Williams has unknowingly proposed the exact same account as Williams. Reason-concepts are referentially normative, as was discussed above, so we have reason* to do what we have reason to do. That means that the alternative theory of reference and the original one would both result in ought-concepts referring to what maximizes utility; it is just that alt Williams would think incorrect things about what his account says. So if Williams is correct, alt Williams does not really provide an alternative account of the content of normative concepts. By the same token, if alt Williams is correct, then both accounts will still be the same, but with the result that ought-concepts refer to what conforms to the categorical imperative.

We can get around this issue by explicitly proposing a theory of content where what determines reference is what alt Williams thinks determines reference. When it comes to structural epistemic and practical reasons as well as substantive epistemic reasons, the view would be the same. That a concept's referring to something would increase the reasonableness of its user along any of these dimensions would still count in favor of that concept having that reference. But where Williams claims that substantive practical reasons play a similar metasemantic role to these other types of reason, our new theory would deny this. In place of this, the theory proposes that all else equal a concept will refer to whatever would make it the

convenience, I am here treating the actual Williams as committed to the toy account of what we have most reason to do being used in this paper, which he is not. You can think of both Williams and alt Williams as described here as being alternatives to the actual Williams that proposed the theory under discussion.

case that people are planning to follow the categorical imperative or blaming others for failing to follow it, or so on. That is, this new theory explicitly states that if an expression has the normative role of an ought-concept, then there is metasemantic pressure for the concept to refer to what conforms to the categorical imperative.

This is similar to how Williams' view puts metasemantic pressure on ought-concepts to refer to what maximizes utility. The difference is that Williams arrives at this through the intermediary of what we have most reason to do. In the first place the metasemantic pressure on ought-concepts is that they refer to what we have reason to plan to do. It is only through the further putative fact that we have most reason to plan to maximize utility that the fact that an act does so has relevance to the reference of normative concepts.

Now that we have found a way of spelling out a genuine alternative to Williams' theory of content that would conflict with it even if Williams' theory were in fact correct, we can ask which theory is correct. The two theories would do an equally good job of explaining the reference of non-normative concepts, since they agree on what mechanisms are relevant to determining the reference of these concepts. In particular, both theories could equally solve skeptical challenges that allege widespread referential indeterminacy for these terms, at least if Williams is correct that his own theory does so. And it is not clear why the fact that Williams' theory entails that normative expressions refer to what maximizes utility whereas the alternative does not should be taken to be an advantage for it.

One could argue that it is an advantage because what maximizes utility *is* what we have most reason to do, so his theory offers a more unified account of reference. All concepts refer to what would best rationalize the actions of their users on his account, whereas the other account is oddly disjunctive in that sometimes what matters is what would rationalize the user's actions and

other times what matters is the relationship between the user's actions and the categorical imperative. But the discussion from earlier in this section should cast doubt on the strength of this response. The claim that maximizing utility is what we have most reason to do is just an expression of the self-identity of the property of maximizing utility. Why should this fact make one theory of reference superior to another? It is no less a fact that acts that conform to the categorical imperative conform to the categorical imperative. Why is this fact in any worse of a position to explain why a theory of content is correct than the one that does the work on Williams' theory? The upshot here is that it is hard to see what facts could make it the case that Williams' theory of content is correct and the proposed alternative false. So by the same token it is hard to see what would make it the case that ought-concepts be referentially normative concepts referring to what maximizes utility rather than ones referring to what conforms to the categorical imperative or any number of possible alternatives.

The preceding paragraph is in part targeted at the concerns of the last section: are the grounds that Williams proposes capable of explaining the reference of ought-concepts? But they raise even greater issue for our second question of whether referentially normative concepts referring to non-special properties suffice for the ardent realist. For just as we have a hard time finding a reason why one of these theories of reference would be correct and the other false, if we just stipulate that Williams' is correct we have an even more difficult time seeing how that helps the ardent realist. There simply does not appear to be anything in Williams' account that can justify the idea that the world itself favors maximizing utility over conforming to the categorical imperative.

You might object that I am placing unreasonable burdens on the ardent realist, requiring

that they establish an unrealistic amount to support their view.⁸⁴ My own perspective is that ardent realism has these commitments latent within itself. As I mentioned when introducing ardent realism, we have a handle on this view by virtue of having a handle on certain hard-to-precisely-define intuitions about the objectivity of normativity. When I reflect on those intuitions, I find that they lead to the rejection of the scenario described above—the scenario of Williams' view vindicating the objective significance of maximizing utility *contra* alt Williams—in much the same way that they lead to the rejection of non-referentially-normative accounts of ought-concepts. That latter rejection, recall, amounts to the rejection of the view that the mind-independent truth of ought-claims vindicates their objective significance in the face of alternative communities whose ought-concepts refer to different things.

Now, one can make a flat-footed response to the objections I have raised to a metaphysically conservative version of Williams' view on behalf of ardent realists. This response would go:

Look, if this view works, then it grants us not just mind-independently true normative claims, but normative claims that will be expressed by anyone using concepts with the same normative roles as our own. What more can a realist want?

But there is also a flat-footed response to the original ardent realist concern, one that I suspect many naturalistic realists would be happy to endorse. That response would go:

Look, this view [one on which normative concepts are not referentially normative] grants us mind-independently true normative claims. And your only objection to it is that some other groups fail to even talk about normative properties because their seemingly normative concepts refer to distinct properties from these. But why does that matter?

⁸⁴ Claims along this line have been suggested in comments by Matti Eklund and Harold Hodes.

There is no serious question about whether we should use our own concepts or theirs for moral deliberation. Our normative concepts are almost by definition the ones that are correct to use, and you yourself grant that one cannot even formulate the question of which concepts to use without it having such a trivial result. And nothing could be less arbitrary than that one should use the normative concepts that are correct to use. What more can a realist want?

There is something appealing in each of these flat-footed responses, and also, I find, a lingering concern underneath that appeal. My goal here is not to reject or accept either one of them. But it does seem to me that as a matter of consistency the intuitions that lead to the rejection of the second response also should lead to the rejection of the first one.

At this stage, defenders of Williams' view have two options. They can propose that there genuinely is something special about the property of maximizing utility that explains its central importance in the theory of content. Or, they can deny that such explanation is needed. Both options have troublesome consequences for those who claim that referential normativity can on its own satisfy important motivations for realism.

First consider the proposal that there is nothing special about the property of maximizing utility that explains its central role in the theory of content. It is simply a brute fact that it plays such a role. This may be acceptable simply as a matter of coming up with the best theory of content and in particular focusing on satisfying intuitions of shared content about moral-twin-earth-like cases. The ultimate grounds of referential normativity facts may seem odd on this view, but any fundamental theory of reference is going to need some primitives and perhaps these are the best we can get. However, even if this response succeeds in defending Williams' metaphysics of content, it should be clear that it does not work for the ardent realist. If there is

nothing to be said as to why maximizing utility is a better ground for a theory of reference than the categorical imperative, then we are in no better of a position than when there was nothing we could say to justify guiding our actions by what we ought to do rather than what we ought_k to do. So an ardent realist will be no happier with the resulting theory of content than they would be with one that resulted in non-referentially normative concepts.

On the other hand, suppose that there must be something special about maximizing utility for it to play such a role in a theory of content. Then first of all, we have to justify that there is something special about it, or, leaving the toy model, whatever property ultimately ends up being the referent of ought-concepts. This would return us to the situation described by Eklund by analogy with the concept of *gidplissing*. Williams' theory of content on its own would not have the result that normative concepts are referentially normative, but would only do so supposing certain contentious metaphysical theses were true. I will not speculate on the exact nature of those theses and whether or not they are in fact true, but clearly such a theory would not achieve the goal of finding a metaphysically conservative theory that satisfies the ardent realist.

5. Conclusion

This all might seem somewhat disappointing. The baseline metanormative situation involves two types of normative realists. One is metaphysically conservative, and the other is generally driven to propose additions to a naturalistic worldview by their conviction that the theories of the first group don't meet the highest standards of normative objectivity. Reframing the debate in terms of referential normativity bore some promise of breaking through this stalemate and providing us with a metaphysically conservative view that can face up to the most stringent realist demands. But if the arguments of the preceding section succeed, this promise was an illusion. Ardent

realists remain committed to non-conservative metaphysical theories. Either referential normativity is not sufficient for ardent realism or referential normativity can only be attained if there are normatively special properties.

So far what I have said mostly leaves untouched the non-realist motivations for referential normativity, which are in fact Williams' motivations for his own view. It may in the end be acceptable for a non-special property to play the fundamental role Williams' view requires of normative properties, at least for those who are not ardent realists. But I think the arguments of this paper may also undermine these semantic intuition-based motivations as well to a lesser extent. Eklund's arguments show that certain realist intuitions prompt the need for there to be shared content in cases like moral twin earth. Perhaps such intuitions are what ultimately generate the semantic intuitions about shared contents that have driven debates surrounding this thought experiment. If so, then Eklund's analogy with gidplissing would apply, and we should only be as confident in our shared content intuitions as we are in the metaphysical theses that underpin them.

To expand on this point, recall that there are two versions of the semantic intuitions argument. The first version relies on our having the intuition that our ought-terms or concepts and the twin ones have the same content. The problem with this argument is that the notions of semantic and mental content are technical ones whose nature is determined in large part by the explanatory roles to which they are put. As such, we should not expect people in general to have accurate intuitions about how these notions apply. Perhaps domain experts can develop such intuitions over time, but the moral twin earth thought experiment is clearly designed to be assessable by people without such expertise. Philosophers often expect undergrads with no significant philosophy of language background to be able to assess these intuitions, for instance.

For this reason, basing the argument on intuitions about content is not the right idea.⁸⁵

Instead, that argument should be made based on intuitions about disagreement. Disagreement is not a technical notion, and so we can expect just about everyone to have potentially worthwhile intuitions about when two statements or mental states are in disagreement with each other.⁸⁶ And further, facts about when two such states disagree with each other are central examples of what the technical notions of semantic and mental content are meant to explain. So non-expert intuitions in this area are plausibly relevant to the questions about content with which we are concerned. And these intuitions have the content not merely that there is a superficial or potentially verbal disagreement, but that the two groups genuinely disagree.

With this as the best version of the semantic-intuition-based argument, the case that these underlying intuitions are influenced by inchoate ardent realist intuitions is bolstered. For when we first introduced those latter intuitions, they were in part defined in terms of one being drawn to the idea that our “difference with Bad Guy [is] more than a mere disagreement in taste” (Eklund 2017, 1). And this being a core feature of the ardent realist intuitions makes sense. If the world itself favors a certain way of acting, then those guiding their actions in accordance with it are right and those greatly diverging are wrong. The two groups must genuinely be in disagreement if just one of the groups is to be objectively correct.

To sum up, the best version of the first type of argument for referential normativity is one based on intuitions about genuine disagreement. These intuitions overlap with the intuitions supporting ardent realism, and additionally it is clear why the intuition that the world itself favors

85 For more along these lines, again see Dowell (2016).

86 The way I've framed this, in terms of content-bearing items disagreeing, is not quite as everyday as the framing in terms of people disagreeing. But I think it's clear that people take people to disagree in virtue of their statements and mental states even on the more everyday framing. Feel free to substitute “intuition that two people disagree in virtue of one having mental state A and the other having mental state B” for “intuition that mental state A disagrees with mental state B” if that seems more natural.

a certain course of action would be accompanied by intuitions about such disagreement. On the basis of this, I find it plausible that the intuitions behind the traditional moral twin earth argument, construed in its best form, are largely due to inchoate ardent realist intuitions. The case is not knockdown, but it is a reasonable one.

If the above claims are correct, they have potentially important consequences for the traditional moral twin earth arguments, which are primarily directed against naturalistic forms of normative realism. Of course, the claim that most people have intuitions that commit them to normative properties being special is not an immediately welcome one for a naturalist. But if the naturalist is correct in thinking that such properties do not exist, then the situation might end up turning out in the naturalist's favor. Without any special properties, no perfect satisfier of people's normative intuitions would exist, and plausibly as a result normative expressions would refer to the best of the imperfect satisfiers of those intuitions, which would likely be a natural property of some kind. The naturalist can thus accept the existence and linguistic relevance of the moral twin earth intuitions as consistent with their view. This would be much harder to do on the original account of those intuitions as springing from the expressivist character of normative discourse.

REFERENCES

- Alwood, Andrew. 2016. "Should Expressivism Be a Theory at the Level of Metasemantics." *Thought* 5: 13–22.
- Anscombe, G.E.M. 1964. *Intention*. Cambridge: Harvard University Press.
- Ayer, Alfred J. 1952. *Language, Truth, and Logic*. New York: Dover Publications.
- Bach, Kent. 1994. "Conversational Implicature." *Mind and Language* 9 (2): 124–62.
- Bezuidenhout, Anne. 2002. "Truth-Conditional Pragmatics." *Philosophical Perspectives* 16: 105–34.
- . 2006. "The Coherence of Contextualism." *Mind and Language* 21 (1): 1–10.
- Borg, Emma. 2004. *Minimal Semantics*. Oxford: Oxford University Press.
- Boyd, Richard. 1988. "How to Be a Moral Realist." In *Essays on Moral Realism*, edited by Geoffrey Sayre-McCord, 181–229. Ithaca, NY: Cornell University Press.
- Cappelen, Herman, and Ernest Lepore. 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.
- . 2006. "Response." *Mind and Language* 21 (2): 50–73.
- Carballo, Alejandro. 2014. "Semantic Hermeneutics." In *Metasemantics: New Essays on the Foundations of Meaning*, edited by Alexis Burgess and Brett Sherman. Oxford: Oxford University Press.
- Carston, Robyn. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Charlow, Nate. 2013. "The Problem with the Frege-Geach Problem." *Philosophical Studies*, 1–31.
- Chrisman, Matthew. 2012. "On the Meaning of 'Ought.'" In *Oxford Studies in Metaethics*, edited

- by Russ Shafer-Landau, 7:304–32. Oxford: Oxford University Press.
- . 2013. “Attitudinal Expressivism and Logical Pragmatism in Metaethics.” In *Pragmatism, Law, and Language*, edited by Graham Hubbs and Douglas Lind. London: Routledge.
- . 2016. *The Meaning of “Ought”: Beyond Descriptivism and Expressivism*. New York, NY: Oxford University Press.
- Dasgupta, Shamik. 2017. “Normative Non-Naturalism and the Problem of Authority.” *Proceedings of the Aristotelian Society* 117 (3): 297–319.
- Davis, Wayne. 2003. *Meaning, Expression, and Thought*. Cambridge: Cambridge University Press.
- Dorr, Cian. 2002. “Expressivism and Wishful Thinking.” *Nous* 33 (4): 558–72.
- Dowell, Janice. 2016. “The Metaethical Insignificance of Moral Twin Earth.” In *Oxford Studies in Metaphysics*, edited by Russ Shafer-Landau. Vol. 11. Oxford: Oxford University Press.
- Dreier, James. 1996. “Expressivist Embeddings.” *Philosophical Studies* 83: 29–51.
- Eklund, Matti. 2017. *Choosing Normative Concepts*. Oxford: Oxford University Press.
- Fine, Kit. 2012. “Guide to Ground.” In *Metaphysical Grounding: Understanding the Structure of Reality*, edited by Fabrice Correia and Benjamin Schnieder, 37–80. Cambridge: Cambridge University Press.
- Frost, Kim. 2014. “On the Very Idea of Direction of Fit.” *Philosophical Review* 123 (4): 429–84.
- Geach, Peter. 1965. “Assertion.” *Philosophical Review* 71 (4): 423–32.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings*. Oxford: Oxford University Press.
- . 2003. *Thinking How to Live*. Cambridge: Harvard University Press.
- Goodman, Nelson. 1983. *Fact, Fiction, and Forecast*. 4th ed. Cambridge: Harvard University

- Press.
- Grice, H.P. 1957. "Meaning." *Philosophical Review* 66 (3): 377–88.
- Groenendijk, Jeroen, and Martin Stokhof. 1991. "Two Theories of Dynamic Sentences." In *JELIA '90, European Workshop on Logics in AI*, edited by Jan van Eijck, 478:55–64. Lecture Notes in Computer Science. Berlin: Springer.
- Hare, R.M. 1952. *The Language of Morals*. New York: Oxford University Press.
- Heim, Irene, and Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.
- Horgan, Terence, and Mark Timmons. 1992. "Troubles on Moral Twin Earth: Moral Queerness Revived." *Synthese* 92 (2): 221–60.
- Jackson, Frank. 1998. *From Metaphysics to Ethics*. Oxford: Clarendon Press.
- Kaplan, David. 1989. "Demonstratives." In *Themes from Kaplan*, edited by J. Almog, J. Perry, and H. Wettstein, 481–563. Oxford: Oxford University Press.
- Kratzer, Angelika. 2012. *Modals and Conditionals*. Oxford: Oxford University Press.
- Larson, Richard, and Gabriel Segal. 1995. *Knowledge of Meaning*. Cambridge: MIT Press.
- Lewis, David K. 1970. "How to Define Theoretical Terms." *The Journal of Philosophy* 67 (13): 427–46.
- . 1974. "Radical Interpretation." *Synthese* 27 (3): 331–44.
- . 1983. "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61 (4): 343–77.
- Mahon, James Edwin. 2016. "The Definition of Lying and Deception." In *The Stanford Encyclopedia of Philosophy*, edited by Edward Zalta, Winter 2016. <https://plato.stanford.edu/archives/win2016/entries/lying-definition/>.
- Millikan, Ruth. 1984. *Language, Thought, and Other Biological Categories*. Cambridge: MIT

- Press.
- Neale, Stephen. 2004. "This, That, and the Other." In *Descriptions and Beyond*, edited by Anne Bezuidenhout and M. Reimer, 68–182. Oxford: Oxford University Press.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Clarendon Press.
- Railton, Peter. 1986. "Moral Realism." *Philosophical Review* 95 (2): 163–207.
- Recanati, Francois. 2004. *Literal Meaning*. Cambridge: Cambridge University Press.
- Ridge, Michael. 2014. *Impassioned Belief*. Oxford: Oxford University Press.
- Rosen, Gideon. 2010. "Metaphysical Dependence: Grounding and Reduction." In *Modality: Metaphysics, Logic, and Epistemology*, edited by Bob Hale and Aviv Hoffmann, 109–36. Oxford: Oxford University Press.
- Sayre-McCord, Geoffrey. 2015. "Moral Realism." In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. <http://plato.stanford.edu/entries/moral-realism>.
- Scanlon, T.M. 2007. "Wrongness and Reasons." In *Oxford Studies in Metaethics*, edited by Russ Shafer-Landau. Vol. 2. Oxford: Oxford University Press.
- Schroeder, Mark. 2008. *Being For: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.
- . 2009. "Hybrid Expressivism: Virtues and Vices." *Philosophical Studies* 160: 257–309.
- Sepielli, Andrew. 2012. "Normative Uncertainty for Non-Cognitivists." *Philosophical Studies* 160 (2): 191–207.
- Sider, Ted. 2012. *Writing the Book of the World*. Oxford: Oxford University Press.
- Silk, Alex, and Russ Shafer-Landau. 2013. "Truth-Conditions and the Meanings of Ethical Terms." In *Oxford Studies in Metaethics*, 8:195–222. Oxford: Oxford University Press.
- Soames, Scott. 2010. *What Is Meaning?* Princeton, NJ: Princeton University Press.

- Sperber, Dan, and Deirdre Wilson. 1986. *Relevance Theory*. Oxford: Blackwell.
- Stalnaker, Robert. 1984. *Inquiry*. Cambridge, MA: MIT Press.
- Stanley, Jason, and Zolton Szabo. 2000. "On Quantifier Domain Restrictions." *Mind and Language* 15 (2): 219–61.
- Street, Sharon. 2009. "In Defense of Future Tuesday Indifference: Ideally Coherent Eccentrics and the Contingency of What Matters." *Philosophical Issues* 19: 273–98.
- Szabo, Zolton. 2006. "Sensitivity Training." *Mind and Language* 21 (1): 31–38.
- Travis, Charles. 2006. "Insensitive Semantics." *Mind and Language* 21 (1): 39–49.
- Unwin, Nicholas. 1999. "Quasi-Realism, Negation, and the Frege-Geach Problem." *The Philosophical Quarterly* 49 (196): 337–52.
- Wedgwood, Ralph. 2001. "Conceptual Role Semantics for Moral Terms." *Philosophical Review* 110 (1): 1–30.
- Williams, J. Robert G. 2016. "Representational Scepticism: The Bubble Puzzle." *Philosophical Perspectives: Metaphysics*, Supplement to *Nous*, 30 (1): 419–42.
- . 2018. "Normative Reference Magnets." *Philosophical Review* 127 (1): 41–71.
- Yalcin, Seth. 2014. "Semantics and Metasemantics in the Context of Generative Grammar." In *Metasemantics: New Essays on the Foundations of Meaning*, edited by A. Burgess and B. Sherman. Oxford: Oxford University Press.