Propositions and Nondescriptivism in Metaethics

James Lindsey David Brown
University College London
MPhil Stud Philosophical Studies

Declaration

I, James Lindsey David Brown, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

.....

James L. D. Brown 1 September 2016

Abstract

This thesis seeks to reconcile two theoretical perspectives: nondescriptivist approaches to metaethics, and propositional approaches to semantic content. Metaethical nondescriptivism maintains that moral language is in some important sense nondescriptive or nonrepresentational. Fundamentally, such language does not serve to describe how the world is. Rather, it should be understood in terms of its nondescriptive functional role in our cognitive economy and social intercourse. It is often thought that nondescriptivism implies the denial of moral propositions or propositional contents. That is, it denies that moral sentences have propositions as their contents, that moral propositions are the objects of attitudes such as belief, and that there are moral propositions that are true or false. This denial leads to a number of very serious problems in providing an adequate account of the semantics of moral language. This thesis argues that the nondescriptivist should reject the assumption that nondescriptivism implies denying moral propositions a place in one's theory of moral language. This assumption presupposes a descriptive or representational conception of propositions that is unavailable to the nondescriptivist. The nondescriptivist should therefore accommodate moral propositions by providing a suitably nondescriptive conception of propositions. The thesis examines in detail three recent attempts in the literature to realise this reconciliation. It is argued that none of the approaches are successful. However, it is nonetheless suggested that the discussion highlights how the nondescriptivist has a wide range of theoretical resources with which to argue for propositional nondescriptivism, providing possible grounds for optimism.

Table of Contents

~ 1 ~
Introduction

This thesis is about moral language. It is about words like 'good' and 'bad', 'right' and 'wrong', and the sentences in which they appear. Moral language has puzzled philosophers. As one writer has it, moral words are Janus-faced. Looking to the past, perhaps, these words display their *descriptive* nature. When I claim that it is wrong for politicians to take cash for questions, I seem describe how matters stand. My description of the moral aspect of cash for questions is something that I can believe, doubt, perhaps even know. On the other hand, when I make this claim, I am not merely stating how things stand. Rather, I am expressing my condemnation towards those who take cash for questions, perhaps in the hope that you will come to share my attitude, or even to deter you from engaging in such actions. Moreover, I seem to do so simply by using a word like 'wrong'. I do not need to do anything additional to indicate my stance on the matter, or to say the word in any particular way. So looking to the future, perhaps, moral words display their *nondescriptive* nature.

In a more theoretical vain, we might say that while moral terms seem to have a straightforwardly descriptive *semantics*, they have a distinctively nondescriptive *functional role*. These days, philosophers generally acknowledge that any theory of moral language needs to account for both aspects. But it is often thought that they are in tension. While in the old days one could get away with simply denying one of these aspects, it is now generally accepted that both require explanation. Very crudely, we can distinguish between descriptivist theories that take the descriptive nature of moral language to be fundamental, and nondescriptivist theories that take the nondescriptive nature of moral langauge to be fundamental. Things are in fact a lot more complex than this, but it's a helpful starting point nonetheless.

The focus on this thesis is metaethical nondescriptivism.[1] I argue that nondescriptivism is at least in principle compatible with the standard semantic resources often thought to be available only to descriptivism. The resources in question are propositional and truth-conditional semantics. Many people define nondescriptivism as the thesis that moral terms do not express a propositional or truth-conditional content. However, this is over-simplistic and need not be accepted. The idea

[1] Nondescriptivism can be applied to a variety of different domains. Throughout, I will use the term as shorthand for 'metaethical nondescriptivism' unless otherwise specified. Context should make it clear in any case how it is being used.

that it must presupposes an assumption about the nature of propositional content that the nondescriptivist may well reject. This is the assumption that propositional content is essentially representational, in some substantive sense to be specified. I argue that it is open to the nondescriptivist to adopt a suitably nonrepresentational conception of propositional content. By doing so, she can accept an orthodox semantics for moral terms while retaining a distinctively nondescriptivist metaethics.

One of the themes of this thesis is that one cannot get too far in philosophy by arguing and theorising at such a general level of abstraction—that is, by "hypothesising *into the blue*". Much of the body of the work will therefore be dedicated to examining developed theories that exploit this argumentative move, or at least something like it. While I will not end up recommending any of the theories discussed here, lessons will be learned about which aspects of which theories provide promise, and which do not. Only recently has the idea that nondescriptivism and orthodox semantics might be compatible gained serious traction. As such, the approach is still in its infancy, braving its first steps. While we will see that establishing those first confident steps is a difficult and complex task, I believe that the resourcefulness displayed by the theories examined gives promise to possibility of such a theory being realised.

1.1 Semantic Nondescriptivism

While nondescriptivism is primarily a thesis about moral language, its motivations and aspirations are more holistic. This thesis is primarily concerned with language, but I will use this section to sketch some of the more general attractions of the position. As well as hopefully showing why nondescriptivism is worth the bother in the first place, we will later see that these considerations give rise to criteria of adequacy for any nondescriptivist theory.

Beginning at perhaps the beginning,[2] consider the following statement of nondescriptivism about moral language:

> This peculiar use of 'good' is, we suggest, a purely emotive use. When so used the word stands for nothing whatever, and has no symbolic function. Thus, when we so use it in the sentence, '*This* is good,' we merely refer to *this*, and the addition of 'is good' makes no difference whatever to our reference. When on the other hand, we say '*This* is red,' the addition of 'is red' to 'this' does symbolise an extension of our reference, namely, to some other red thing. But 'is good' has no comparable *symbolic* function; it serves only as an emotive sign expressing our attitude to *this*, and perhaps evoking similar attitudes in other persons, or inciting them to actions of one kind or another. (Ogden and Richards 1923: 125)

[2] "I must confess that I had read *The Meaning of Meaning* some years before I wrote *Language, Truth and Logic*, but I believe that my plagiarism was unconscious" (Ayer 1984: 28)

'Peculiar' here simply refers to the fact that the word 'good' has many other applications outside the moral domain. We are primarily concerned here and throughout with the moral use.

In this passage, there are at least three claims about moral language that we can distinguish. First, we have the negative claim that moral language is not 'symbolic'. In our current terminology, we can say that it is nondescriptive or nonrepresentational.[3] Second, we have the positive claim that moral langauge is expressive, in the sense that its function is to express our attitudes or mental states. Although it is only implicit in the above passage, we can also distinguish a third positive claim that the kind of attitude expressed by moral language is a nonrepresentational conative or desire-like attitude. This is in contrast to a representational attitude such as belief.

Nondescriptivism is standardly understood to be a semantic view about the meaning of moral terms. Whereas the meaning of perceptual terms, such as 'red', are predicative and can be given in terms their extension or of the property they refer to, moral terms, such as 'good', have some other kind of meaning. Call this thesis semantic nondescriptivism. Semantic nondescriptivism tells us nothing positive about what moral terms mean. Given that we use moral terms a great deal in everyday discussion, it seems implausible to say that these words have no significance whatsoever (even if they are not 'literally meaningful')[4]. So nondescriptivism is usually supplemented with a positive thesis about the kind of meaning that moral terms do have. The dominant approach is to give the meaning of moral expressions in terms of the nonrepresentational attitudes or mental states that we use moral words to express. Strictly speaking, semantic nondescriptivism consists only of the negative claim. However, for simplicity, I'll often use the term to capture all three claims.[5]

It seems undeniable that moral terms are generally used to express nonrepresentational states, such as condemnation, approbation, resentment, guilt, pride, and so on. But the descriptivist need not disagree with this. Her disagreement with the semantic nondescriptivist is over the claim that the *meaning* of moral terms consists in their being used to express these attitudes. So what is gained by semantic nondescriptivism? To answer this question, we need to look to metaethics more generally.

Consider the following suggestion of the explanatory scope of metaethics:

> Understanding the commitments of ordinary moral or value discourse and practice would appear to involve accounts of at least the following: the semantics of the language of morals and value; the apparent metaphysical status of moral properties or values; the putative

[3] I will treat 'descriptive' and 'representational' as more or less synonymous throughout.

[4] c.f. Ayer (1946: 107).

[5] Kalderon (2005) emphasises the importance of recognising that these are distinct claims. For the purposes of this chapter, however, we can ignore these issues.

> epistemology of morality or value theory; and the relation of morality or values to practical reasoning. (Darwall, Gibbard and Railton 1992: 127)

As these writers go on to note, these questions are all interconnected. How one accounts for one aspect will have implications for how one accounts for others. I think that semantic nondescriptivism can be helpfully thought of as taking moral semantics as its starting point, and then proceeding to account for the other aspects of moral practice in terms of its semantics.

Nondescriptivism is usually born out of a general philosophical naturalism. This is best seen as a sort of methodological naturalism rather than any single thesis or cluster of theses.[6] On this view, philosophy is constrained by the natural sciences—particularly physics, chemistry, biology—in terms of what kind of ontology it is acceptable to posit in one's theory (i.e., only those entities posited by the natural sciences).[7] Philosophy is also constrained in terms of what sciences tells us about ourselves as natural creatures in a natural environment:

> To be a naturalist is to see human beings as frail complexes of perishable tissue, and so part of the natural order. It is thus to refuse unexplained appeals to mind or spirit, and unexplained appeals to knowledge of a Platonic order of Forms or Norms; it is above all to refuse any appeal to a supernatural order." (Blackburn 1998: 48-9)

Of course, there are many ways of being a naturalist, and no sensible non-naturalist would leave any appeal to a non-natural order as 'unexplained'. But here at least is a starting point for theorising in general.

One of main attractions of nondescriptivism is its metaphysical and epistemological solvency.[8] The characteristic move of nondescriptivism is to argue that a demand for an explanation of the metaphysics and epistemology of morality is premised on a misunderstanding about moral language. Consider the term 'good'. If we use 'good' to refer to a property in the world, then presumably we are owed some account of the nature of this property and how it can be 'placed' or 'located' in the natural world. For goodness is not obviously the kind of thing that is delivered to the senses or revealed by natural science (though one might argue either of these points).

Supposing that it is, by what sort of means or mechanism do we come to know about or even talk about these properties? The cognitive sciences provide detailed explanations of the mechanisms by which we perceive the world through sense perception. However, while it might be commonplace to say that we can sometimes just '*see*' the wrongness of an action or the goodness of person, we do not have a 'moral

[6] c.f. Railton (1989) on between 'methodological' and 'substantive' naturalism, and Price (2013) on 'subject' and 'object' naturalism.

[7] Note that banning non-natural entities in one's explanatory theory is not necessarily the same as banning non-natural entities in other contexts. See below and §3.5.2.

[8] The phrase is Miller's (2013).

faculty' akin to sense perception. More plausibly it is simply another way of saying that we have the capacity to make moral judgments, which is itself the very thing we are trying to explain.

Semantic nondescriptivism, on the other hand, seems to sidestep these awkward questions entirely. For recall that 'good' is not used to refer to a property. Rather, it is used to express our sentiments. As 'good' does not purport to represent anything in the world, it is simply a mistake to try and provide a metaphysics of what 'good' refers to. To think that any such account is owed is to mistake the meaning and function of the moral terms. *A fortiori*, no epistemological account of how we could have cognitive access to the moral domain is required either. This is not to say, however, that our moral practices are in any way mistaken—that slavery is not *really* wrong, for example, because 'wrong' *merely* expresses a noncognitive attitude. Slavery *is* wrong, and we can even say that this is objectively true. However, to say as much is to take an ethical stance, not to describe a way that the world is. If semantic nondescriptivism is correct, then there is simply no 'external' sense in which our moral claims 'hook up with' the natural world.

These reasons also explain the advantages of nondescriptivism over so called error-theories. According to such views, while moral language is descriptive, and so terms like 'good' purport to refer to some property of goodness, no such properties actually exist.[9] As such, all (positive) moral claims turn out to by systematically false. So my claim that slavery is wrong is in fact false, as there simply are no moral states of affairs in which wrongness is instantiated to make my claim true. Like nondescriptivism, error-theories are both naturalistically acceptable and metaphysically and epistemologically solvent. So error-theory is attractive for reasons similar to nondescriptivism.

However, given the choice, it seems preferable to opt for a theory that does not render such a commonplace practice systematically mistaken. Not only does this seem to offend common sense. From a theoretical perspective, if there we have two competing interpretations of a practice, one of which interprets the practice as systematically mistaken, then the principle of charity tells us to accept the interpretation that does not. The theoretical point is all the more strengthened when we see that error-theorists usually explain the existence and purpose of moral discourse in terms of its social function. If moral discourse does indeed have this function, then it seems more plausible to explain the practice primarily in terms of this function rather than primarily in terms of some erroneous representational function.[10]

So nondescriptivism is acceptably naturalistic, metaphysically and epistemologically solvent, and provides a vindicatory explanation of folk moral practice.[11] Though this is

[9] The *locus classicus* is Mackie (1977).

[10] Blackburn (1993a).

[11] Quietist descriptivist approaches such as Dworkin (1996) would also claim to have metaphysical and epistemological solvency. There is a delicate question as to how best to characterise the difference between these positions and nondescriptivism. However, the nondescriptivist might argue for her

more controversial, many nondescriptivists argue that descriptivism (both naturalist and non-naturalist varieties) faces serious problems in explaining the distinctive directive role of moral judgments (the mental states expressed by moral statements). There seems to be a strong connection between moral judgment and motivation to action. For example, say that I sincerely claim that stealing is wrong, yet I often steal things, absent of weakness of will, being a kleptomaniac, etc. You ask me if I really do believe that stealing is wrong, seeing as I have no qualms with stealing. I reply that I do sincerely believe that stealing is wrong, and that this gives me a reason not to steal—however, these considerations do not move me and I feel no compulsion against stealing.

Something seems to have gone wrong here. What is the nature of the mistake? Some have argued that there is a conceptual connection between moral terms and being motivated to act. In the above example, I would have misused or misunderstood the concept of wrongness, and so would have failed to make a genuine moral judgment. Alternatively, one might maintain that making a moral claim rationally commits one to be motivated to act accordingly. So above, I would be irrational by not being motivated not to steal. Weaker still, we might think that it is necessary that normal or typical members of a community be so motivated to act. All of these views are versions of *motivational internalism*: that there is a necessary connection between moral judgment and motivation to action.[12]

Nondescriptivists often argue that descriptivists have a hard time explaining motivational internalism. Descriptivism has it that to make a moral judgment is to accept a moral proposition. However, for any proposition that we might accept, there is another question of what to do about it, or more generally what normative significance to give it: "Even if that belief were settled, there would still be issues of what importance to give it, what to do, and all the rest. For we have no conception of a 'truth condition' or fact of which mere apprehension by itself determines practical issues. For any fact, there is a question of what to do about it." (Blackburn 1998: 70)

Some think that this is the lesson to be learned from Moore's open question argument.[13] For any putative descriptive definition of, say, 'goodness', there is always an open question whether the definiens *really is* good. According to nondescriptivism, this is just a reflection of the fundamentally different nature of descriptive judgment and moral judgment. Semantic nondescriptivism explains this straightforwardly: motivational internalism follows directly from semantic nondescriptivism. To claim that

---

position on the grounds that nondescriptivism provides a thoroughly naturalistic explanation of moral practice that the quietist lacks. On the virtues of nondescriptive over quietist explanations generally, see Price (2015).

[12] See, for example, Darwall (1997) for a survey of the various forms of motivational internalism.

[13] e.g. Darwall, Gibbard and Railton (1992: 116-20).

stealing is wrong is to express a conative attitude towards stealing simply in virtue of the meaning of 'wrong'.[14]

This line of argument is controversial for a number of reasons. First, it arguably begs the question against the descriptivist. On what non-question-begging grounds does the nondescriptivist claim that there is no fact the mere apprehension of which would determine practical issues? Second, it is perfectly open to the descriptivist to accept that desire-like motivational states necessarily play a role in moral judgment. The descriptivist only has to deny that the sentimental nature of moral judgment has any role to play in the semantics of moral language (or is constitutive of the moral judgment itself). Third, motivational internalism is itself a controversial thesis. The descriptivist might instead adopt an externalist position, according to which it is only a contingent fact that people who make moral judgments are motivated to act in certain ways.

So arguments for nondescriptivism from the motivational aspect of morality are inconclusive. It is worth noting, however, that almost all parties to the debate do agree that the connection between moral judgment and motivation stands in need of *some* explanation, whatever this connection exactly is. Nondescriptivism provides possibly the most straightforward explanation of the practical role of moral discourse as nondescriptivism is fundamentally formulated in such terms. For example, semantic nondescriptivism (in the broader sense) involved specifying the directive attitude or mental state in providing the meaning of moral terms. This seems to be in contrast to descriptivism, which requires additional supplementary theory to explain how this particular area of descriptive discourse has such strong ties to motivation.

There are other ways in which the nondescriptivist might argue against descriptivism.[15] However I won't pursue such issues any further here. What I hope is clear from the discussion is the way in which nondescriptivism begins with a single thesis about language and then uses this to explain (or perhaps explain away) other puzzling aspects of moral discourse and practice. While a number of problems have been raised for semantic nondescriptivism, there is one very big problem in particular. This is the Frege-Geach problem, which will be the topic of the next section.

1.2 Frege-Geach

The Frege-Geach problem is not so much a single problem as a whole host of problems that arise for the semantics of nondescriptivism.[16] Semantic nondescriptivism for

[14] Nondescriptivists might also cite the Humean theory of motivation—that beliefs are motivationally inert and always require a desire like state—as defence of nondescriptivism.

[15] For example, one might also argue that the descriptivist has a hard time explaining how the moral supervenes on the natural—i.e., how there can be no moral difference with a natural difference. See, for example, Blackburn (1993b) and Horgan and Timmons (1992).

[16] Geach (1960, 1965).

language *L* is the thesis that the sentences of *L* do not have propositions as their contents.[17,18] Compare:

(1) 'Stealing is a crime'
(2) 'Stealing is wrong'

Let's assume that (1) is descriptive and (2) is nondescriptive. The meaning of (1) can be given in terms of the content that it expresses, namely the proposition *that stealing is a crime*, which provides the truth-conditions for (1).[19]

As 'stealing' is a descriptive term, we can assume that it has the same semantic function in (2) as in (1). In (2) 'stealing' picks out the object of the noncognitive attitude that one has by accepting (2). (We'll see how this can be made more precise in the next chapter, but this simple formulation will suffice for present purposes.) Now, however, consider the following sentence:

(3) 'If stealing is wrong, then tax avoidance is wrong'

Here, (2) appears in an embedded or 'unasserted' context, *viz*., the antecedent of a conditional. Notice, however, that to accept a conditional, one need not accept its antecedent. For example, I do not need to believe that Elvis is alive to believe that if Elvis is alive, then he is keeping a lot of secrets. So in order to accept (3), I should not have to accept (1). This means that I can accept (3) without having any negative attitudes towards stealing.

However, according to semantic nondescriptivism, to say that something is wrong is to express one's (say) condemnation towards it—this is because of what 'wrong' *means*. So by its very meaning, my acceptance of (3) should involve my having this attitude towards stealing. However, we have just seen that this is false. No such attitude is required to meaningfully use 'wrong' in (3). So if the meaning of 'wrong' in (2) is given in terms of the attitude one has in accepting (2), then 'wrong' cannot have the same meaning in (3). However, it clearly does have the same meaning. So semantic nondescriptivism provides an incorrect account of the meaning of wrong.

---

[17] I use 'sentence' as shorthand for 'declarative sentence' throughout.

[18] I will treat 'meaning' and 'semantic value' as more or less synonymous, as with 'propositional meaning' and 'content'. Yalcin (2014) argues that a theory of semantic value should be sharply distinguished from a theory of content. I ignore this complication here.

[19] The semantic sense of 'expression' is not the same as the act-type sense of 'expression'. This terminology is in some ways unfortunate, as the sense of 'express' in the semantic nondescriptivists semantics is more plausibly the latter, whereas the sense in which a sentence expresses a proposition is clearly the former. See Bar-On and Chrisman (2009) for discussion of this issue.

Moreover, it is not simply a matter of providing some other kind of meaning for 'wrong' in embedded contexts. For observe that anyone who accepts (2) and (3) seems also entitled to infer:

(4) 'Tax avoidance is wrong'

If 'wrong' were to have a different meaning in embedded contexts, the inference to (4) from (2) and (3) would commit the fallacy of equivocation. Both sentential and subsentential expressions are compositional in how they combine together with other expressions to make more complex ones. However, they retain a common content across all such contexts. The descriptivist, by positing propositions as the contents of sentences, faces no such problem, as propositional content remains stable and constant across contexts. The nondescriptivist, by contrast, fails to account for this.

The Frege-Geach problem not only highlights a problem about meaning, but also about logic. The reason why we are entitled to the above inference is because it is logically valid: (roughly) the premises cannot be both true and yet the conclusion false. If we take these sentences to be descriptive, then the entailment and inconsistency relations between them can be explained in terms of the alethic properties of the propositions they express. No such explanation is available to the semantic nondescriptivist, however. Indeed, if the nondescriptivist allows this explanation in the descriptive case, then it is unclear how truth and valid reasoning are even meant to apply in the moral case.

Of course, no one holds this sort of semantic nondescriptivism anymore, and much progress has been made developing sophisticated nondescriptive semantics that can answer these problems. I won't consider these here[20], but it is worth noting the enormous undertaking that any such approach involves:

> In fact, and this cannot be emphasized enough, *every* natural-language construction that admits of descriptive predicates admits of moral predicates, and seems to function in precisely the same way: tense; conditionals; every kind of modal—alethic, epistemic, or deontic; qualifiers like 'yesterday'; generics and habituals; complement-taking verbs like 'proved that' and 'wonders whether'; infinitive-taking verbs of every class, including 'expects to', 'wants to', and 'compels to'; binary quantifiers like 'many' and 'most'; and more. It is crucially important to understand that the embedding problem for noncognitivism is not simply a problem about the validity of *modus ponens*, or even simply about logic. *Every* construction in natural languages seems to work equally well no matter whether normative or descriptive language is involved, and to yield complex sentences with the same semantic properties. (Schroeder 2008b: 5)

By appealing to propositions, however, we provide a straightforward and elegant account of how the sentences compositionally interact across a range of contexts while

[20] See Schroeder (2008a) and the references therein for an overview.

retaining a stable content. Moreover, we also provide a straightforward and elegant account of the logic of the sentences in question. We might therefore hope for a nondescriptivism that is compatible with propositionalist approaches to semantics and logic. This is the subject of this thesis.

1.3 Propositional Nondescriptivism?

The proposal put forward in this thesis is that nondescriptivism and propositional or truth-conditional semantics are compatible. The basic thought is very simple. Propositions or propositional contents are posited in semantic theory to play a certain theoretical role. Particularly, they are the contents of sentences, the objects of propositional attitudes, and the primary bearers of truth and falsity. By postulating propositions in one's semantic theory of some domain, one can account for the compositional and logical properties of that domain of discourse. In and of itself, such an approach does not imply anything about the *nature* of propositions or propositional content. Therefore, there is no reason, at least in principle, why a suitably nondescriptive or nonrepresentational conception of propositions is not available to the nondescriptivist. If propositional content is not substantively representational, then a commitment to moral propositions need not imply a commitment to a metaphysically and epistemologically problematic domain of morality.

The 'representationalist assumption' implicit in the standard set up of the debate between descriptivism and nondescriptivism is that propositions are essentially representational in some robust sense.[21] This robustness might be thought to consist in an ontological commitment to moral properties that figure in states of affairs that correspond to true moral propositions via some metaphysically substantive, perhaps naturalistic, word-world relation. As natural as this assumption might be, it is not compulsory. Seeing that it is not compulsory creates the conceptual space required to develop a nondescriptivism that allows for moral propositions or propositional contents in its semantics. For want of a more attractive name, I will call such a position propositional nondescriptivism.

At this level of abstraction, I don't think that there is much that can be said about exactly what propositional nondescriptivism consists in. As we shall see throughout our inquiry, there is no one single way in which to develop this thought. Clearly, however, there are at least two criteria of adequacy for any such theory. First, the account must be acceptably naturalistic. That is, it must explain our moral discourse and practice in a way that is consonant with our being natural creatures with natural capacities. Second, the account must be metaphysically and ontologically solvent. That is, it must explain moral discourse and practice in such a way as to sidestep the need for any metaphysics or epistemology of morality. It should go without saying too that one must also reject

[21] c.f. Price (2013: 9).

semantic nondescriptivism, as this is clearly incompatible with propositional nondescriptivism.

Given more orthodox theories about the nature of propositions, it is perhaps not surprising that this conceptual space has gone largely unnoticed. Consider a Fregean account, in which propositions are the senses of sentences, and senses are understood as modes of presentation that determine the reference of the sentence, where this is a state of affairs. The conception of reference here is clearly representational, as it commits one to the state of affairs presented by propositions. Or consider a Russellian view, in which the constituents of a proposition are the very worldly individuals that the proposition is about. On such a view, if a sentence has a proposition as its content, then it's hard to see how we could have anything but descriptivism for those sentences.

If propositions are the primary bearers of truth and falsity, another reason for thinking that propositions are essentially representational is if one had some robustly representational conception of truth as some sort of substantive metaphysical relation between propositions and the world (Bar-On, Chrisman and Sias 2014: 235-6). Alternatively, if one takes possible worlds conceptions of propositions at face value, then propositions might naturally be associated with metaphysical commitment. For example, if one takes an ersatz view of possible worlds according to which possible worlds—maximally specific ways the world might be—simply are maximally consistent propositions (Schroeder 2013: 420). The present point is not that nondescriptivism is consistent with standard conceptions of truth and propositions. It is simply that one need not be committed to standard conceptions simply by using propositions in one's semantic theory.

The remainder of this thesis is divided into two main chapters. Chapter 2 examines respective attempts by Mark Schroeder and Michael Ridge to carry out this project within a broadly expressivist framework, where semantic content is derived from mental content. Both writers distinguish between propositions, which are not associated with metaphysical noncommitment, and representational contents, which are. Whereas both moral and descriptive belief are understood to involve propositions, only descriptive belief involves representational contents. Despite sharing this fundamental distinction, each view develops it in very different ways. I argue, however, that the distinction in neither case can be made to work.

Chapter 3 examines a suggestion that propositional nondescriptivism might be best formulated as a thesis in metasemantics. The focus of the chapter will be Matthew Chrisman's propositional nondescriptivism, which argues for a semantics for 'ought' as a special kind of modal operator, which in turn is used to argue for a nondescriptivist metasemantics. In addition to its detailed development of a complementary semantics and metasemantics for moral language, Chrisman's account is also noteworthy for its inferentialist rather than expressivist formulation of nondescriptivism. It will be argued

that Chrisman's nondescriptivism is problematically incomplete. The chapter concludes with a tentative suggestion for a way forward.

To lesser and greater extents, the writers examined here have broader theoretical and explanatory aims than those of the present inquiry. At times I have adapted certain aspects of the views in order to better suit the purposes of this thesis. I try to point out where I have done this. I don't think that I ever stray too far from any of the original views. However, if the authors themselves would not wish to own the change, the adapted views nonetheless provide interesting and fruitful materials, and of course owe themselves to their originals.

~ 2 ~
Expressivism and Representational Content

This chapter focuses on a strategy for accommodating moral propositions within metaethical nondescriptivism within an expressivist framework. According to expressivism, the contents of sentences are accounted for in terms of the states of mind that we use them to express. Expressivists further hold that the states of mind expressed by moral claims are nonrepresentational or noncognitive. This is in contrast to descriptive sentences that are used to express representational beliefs. The distinction between the two kinds of mental state is standardly characterised in terms of propositions. Whereas representational beliefs have propositional content, moral beliefs do not. To reconcile expressivism and propositional approaches to moral content, we need some other way to distinguish descriptive and moral thought. The suggestion explored in this chapter is that the expressivist should distinguish between propositions and representational contents. Whereas the former are the objects of our attitudes and contents of our sentences, the latter play a representational role of marking metaphysical distinctions in reality. Nondescriptivism can then be understood as the claim that there are moral propositions but there are no moral representational contents. Descriptive thought is then distinguished from moral thought in terms of the relation that descriptive thought bears to representational contents. This chapter examines two attempts to provide a propositional nondescriptivism along these lines.

2.1 Expressivism, Propositions, and Representational Contents

In the sense under discussion in this chapter, the essence of expressivism can be summed up in the following slogan: "to explain the meaning of a term, explain what states of mind the term can be used to express." (Gibbard 2003: 7) This raises two questions. What about a term's meaning is being explained? What it is to express a state of mind?

Perhaps surprisingly, expressivists and their opponents have not always been clear on what they take the answers to these questions to be.[1] Probably the most prominent answer to these questions today is that expressivism provides a psychologistic compositional semantics for moral language, or natural language more generally. The

[1] We will explore the variety of possible explanatory questions in some depth in the next chapter.

idea of a semantic theory for moral expressions is to assign contents (meanings, semantic values) to moral terms in such a way as to systematically explain the semantic contribution that moral terms make to the contents of the claims in which they figure. Orthodox approaches frame this in terms of the contribution that moral terms make to the truth-conditions of moral propositions. By contrast, expressivism frames this in terms of the contribution that moral terms make in determining what mental state is 'expressed' by moral claims.[2] The expression relation can then be cashed out in terms of assertability conditions for moral claims, where the conditions make reference to what mental state the speaker is conventionally in when uttering the claim. This is in contrast to a truth-conditional approach, where what is expressed by a claim is not a kind of mental state but a proposition.[3]

Importantly, expressivism takes mental content to be more fundamental than linguistic content. The contentfulness of sentences derives from the contentfulness of the states of mind they express. In order to know whether a claim is descriptive or nondescriptive, it must be known whether the mental state expressed by the claim is descriptive or nondescriptive. As a semantic program, expressivism does not imply nondescriptivism. For all that has been said, the mental states expressed by moral claims might be robustly representational beliefs. However, it is more or less universal in metaethics for expressivists to also be nondescriptivists about moral thought. As such, I will henceforth use 'expressivism' to mean nondescriptive expressivism, unless stated otherwise.

In virtue of what is moral thought nondescriptive? In virtue of what are other kinds of thought descriptive? While expressivists have many positive things to say about the nature of moral thought, the sense in which it is nondescriptive is ultimately a negative claim—that it is not descriptive. Standardly, descriptive thought is explained in terms of propositions. For example, a descriptive belief is defined as an agent's bearing a particular relation to a corresponding descriptive proposition. The truth-conditions for descriptive claims are then derived from the corresponding proposition. On this picture, moral belief is nondescriptive in virtue of its not consisting in bearing this relation to a moral proposition. Rather, moral belief has some other nondescriptive functional role.

While the positive story that the expressivist gives of the nondescriptive role of moral thought is vital in explaining moral thought, it should noted that the descriptivist need not disagree with this account. For the descriptivist will presumably agree that moral thought and discourse does indeed have a distinctive practical role in our lives. This is not inconsistent with a descriptivist conception of moral thought and discourse, though as we saw in the previous chapter, this might require some explanation. In any case, stating the distinctive functional role of moral thought is not alone sufficient to

[2] This is explored in detail in Schroeder (2008b: ch.2).

[3] Plausibly, there are two distinct senses of 'expression' here—see Chapter 1 n.18.

distinguish descriptivism from nondescriptivism. Rather, the nondescriptivist needs to deny some positive claim that the descriptivist is committed to.

If one wishes to advance a propositional nondescriptivism, then the above distinction between descriptivism and nondescriptivism will not do. How else might this distinction be drawn? Note that the main expressivist thought is that moral thought is distinctively different to descriptive thought *qua* attitude. It does not follow from this basis idea that the distinction must be given in terms of propositions or propositional content. Of course, if one already has a representational conception of propositions, this way of drawing the distinction seems obvious enough. However, there is nothing in the expressivist framework that forces this conception. So one might hope to find some other characteristic of descriptive thought that moral thought lacks.

Here is a suggestion. Expressivists claim that moral thought is nondescriptive because descriptive thought consists in being appropriately related to an essentially representational entity. Suppose that propositions are not essentially representational. One might therefore posit some other kind of entity that is essentially representational that only descriptive thought is appropriately related to. Call such an entity a *representational content*. If both descriptive and nondescriptive thought are propositional, one might say that what makes descriptive thought descriptive is its being appropriately related to representational content; what makes moral thought nondescriptive is its not being appropriately related to representational content. This is because there are no moral representational contents, though there are moral propositions.

Propositions play the role of being the objects of our attitudes, the contents of our sentences, and the bearers of truth and falsity. What sort of role do representational contents play? Nondescriptivism rejected moral propositions because of the supposed metaphysical import of (true) propositions for the moral domain. So we might say that representational contents are those entities that "serve to carve up the world," that "correspond to distinctions in reality," and that "are associated with metaphysical commitment of some kind" (Schroeder 2013: 418). While these representational-*cum*-metaphysical roles are often associated with propositions, the expressivist might hope to pull them apart from the other roles played by propositions in order to make room for a nondescriptive conception propositions.

This leaves open the question as to what kind of things propositions are, what kind of things representational contents are, and how the two are related. However, one might hope that by drawing this distinction, we open up the conceptual space required for a propositional nondescriptivism. The remainder of this chapter examines two expressivist theories that utilise this distinction. By 'expressivism', I just mean a commitment to explaining semantic content in terms of mental content, rather than the narrower sense of the term sketched above. As will be seen, the two views develop the basic thought in quite different ways. Both, however, will be rejected.

2.2 Schroeder on Propositions and Representational Contents

The first view to be examined is due to Mark Schroeder (2013). Schroeder develops an expressivist view in which both moral and descriptive beliefs have moral and descriptive propositions as their respective contents, but only descriptive beliefs involve a representational content. The key move is to show exactly how moral and descriptive beliefs are structured. Schroeder suggests that whereas all beliefs consist in bearing a certain relation to a proposition, descriptive beliefs also involve bearing a relation to a representational content. If propositions and representational contents are distinct entities, the task for a propositional nondescriptivism is to explain how a single descriptive belief-state can consist in both a relation to a proposition and a distinct relation to a representational content. After providing a highly account of how moral and descriptive beliefs are structured, Schroeder then implements the general framework to provide a more fully worked out propositional nondescriptivism based on his (2008b) biforcated attitude semantics.

2.2.1 Schroeder's Theory: Structured Beliefs and Being For

As Schroeder sees it, the key to understanding how a nondescriptivist can allow for moral propositions is to understand the difference in *structure* between the mental states characteristic of moral thought and those characteristic of descriptive thought. It is orthodoxy to understand belief in terms of a particular relation that an agent bears to a proposition. For any agent *A* and any proposition *P*, *A* believes *P* just in case *A* bears the belief-relation to P. All beliefs necessarily exhibit the following structure:

(1) $A(P)$ (2013: 422)

All beliefs have this structure. If I believe that Meredith ought to seize the means of production, then I stand in the belief-relation to the proposition *that Meredith ought to seize the means of production*. Moral beliefs involve no representational content, and so the basic structure of moral belief is given simply by (1).

The suggestion was that in descriptive cases, belief also consists in a relation to a distinct representational content in addition to a proposition. So the complexity of the structure of descriptive belief needs to be increased in order to accommodate this. Schroeder suggests the following. For any agent *A*, any descriptive proposition $B(C)$ and any representational content *C*, descriptive beliefs necessarily exhibit the following structure:

(2) $A(B(C))$ (2013: 422)

As $B(C)$ denotes a proposition, all instances of (2) will be instances of (1). However, descriptive propositions are necessarily structured so as to contain a representational content as a propositional constituent. Furthermore, (2) can also be 'carved up' to get the relation $A(B(\_))$, which is a relation an agent bears to a representational content.[4] Thus, descriptive beliefs consist in two separate relations to two distinct (though related) entities.

To get a firmer grip on how this last point works, Schroeder makes the following analogy. Suppose I am about to go on holiday to Paris. We might say that I am in the state of *being about to go to Paris*. What does being in this state consist in? On the one hand, it involves me being appropriately related to the city Paris, where the relation is something like *being about to go to x*. However, it also involves me being appropriately related to the act-type of going to Paris, where the relation is something like *being about to φ*. Schroeder make the following observation:

> Since the relational action of *going to* is part of both the action of *going to Paris*, and of the relation of being *about to go to*, and since Paris figures in the action of *going to Paris*, there is no puzzle about how the state of being about to go to Paris can be carved up in each of these distinct ways. (2013: 421)

This highlights how the state of being about to go to Paris is not simply a conjunction of two distinct states that I happen to be in. Rather, it consists in these two states. How one decides to analyse the complex state depends on one's theoretical interests.

In an analogous way, descriptive belief states can be analysed into an agent bearing two distinct relations to two distinct entities. First, it consists in one's bearing a certain relation to a descriptive proposition, which necessarily contains a representational content as a constituent. But it also consists in one's bearing a different relation to the representational content. Having a moral belief, by contrast, consists only in one's bearing a relation to a proposition. To lend support to the view, Schroeder notes that, if true, it would explain why a distinction between propositions and representational contents is so easy to overlook (2013: 422). Many paradigm cases of belief are descriptive beliefs that involve both a proposition and representational content, this fact only revealed theoretical analysis. As such, it would be natural to conflate the two.

At such an abstract level, it is hard to assess this general framework. Indeed, it is hard to see exactly what a more concrete example would look like. For example, no clue

[4] It's actually unclear exactly how the relations are represented on this picture. While Schroeder expressly states that '$A(\_)$' and '$A(B(\_))$' denote relations, strictly speaking, the relations are actually denoted by '$\_(\_)$' and '$\_(\_(\_))$'. Given that the structure of descriptive propositions is $B(C)$, this seems to commit Schroeder to the claim that propositions necessarily involve some binary relation. Confusingly, Schroeder goes on to talk about the relation denoted by 'B'. This is confusing as the other relations are given by brackets, not letters. We will see below that this creates problems for Schroeder and is not merely a matter of presentation.

has been given to what sort of thing 'B' designates.[5] (We will return to this in the next section.) Furthermore, there is nothing about the framework itself that is inherently nondescriptivist. True, we have a stipulated distinction between propositions and representational contents. However, nothing at all has been said about what kind of thing either of these things are. Even accepting that there is a distinction, propositions might still turn out to be unacceptably representational for the nondescriptivist. I submit that the only way of answering these questions is by looking to the implementation of the general framework, to which we now turn.

Schroeder develops a variant of his biforcated attitude semantics (BAS), developed in his (2008b), as an instance of the general framework outlined above. The original semantics were developed as an expressivist alternative to truth-conditional semantics for natural langauge. On the present picture, it aims to accommodate moral propositions rather than reject them. As such, the variant semantics are not really an alternative to truth-conditional semantics. Rather, it might be instead seen as an alternative picture to the standard view of what truth-conditional semantics is about. Indeed, Schroeder takes "the primary lesson" of his variant account to be "that the differences between descriptivist and nondescriptivist semantics can be fruitfully thought of not as a dispute between very different ways of doing semantic theory (one with propositions and one without), but rather as a dispute about what propositions—the objects of the attitudes and bearers of truth and falsity—are *like*." (2013: 424) The account presented here is a somewhat simplified version of Schroeder's own account. As well as for readability, this is because there are a number of technical details that are irrelevant for present purposes. Nothing should hang on this.

Schroeder's BAS was originally developed in an attempt to provide a formally adequate solution to the Frege-Geach problem. It takes as its starting point the problem as it arises for negation, but the solution is fully general.[6] What is important for present purposes is Schroeder's diagnosis of the problem. As he sees it, the problem arises due to the lack of structure in the attitudes that moral sentences are purported to express. This is in contrast to the attitudes expressed by descriptive sentences. Consider:

(3) 'Snow is white'

The presence of the descriptive predicate 'is white' indicates that the speaker of (1) expresses a descriptive kind of mental state, namely *belief*. Further, the predicate also corresponds to a certain content of the belief, namely the propositional constituent *is white*. Moreover, any descriptive predicate used assertorically will indicate both (i) a

[5] See n.4 above.

[6] For the Frege-Geach problem as it arises for negation, see Unwin (1999, 2001).

uniform attitude (belief) and (ii) some particular descriptive content that figures in the overall content of the attitude.

Now consider:

(4) 'Stealing is wrong'

Standardly, expressivists claim the moral predicate 'is wrong' indicates a certain attitude, say *disapproval*, that the speaker has towards stealing. Unlike descriptive predicates, there is no particular kind of moral content that moral predicates pick out that appears in the overall content of the attitude. Whereas descriptive predicates indicate *a general structured* attitude, moral predicates indicate *particular unstructured* attitudes.

Schroeder argues that it is precisely this lack of structure in the kind of mental states expressed by moral claims that gives rise to the Frege-Geach problem. After all, we would not say that the descriptive predicate 'is white' in (3) indicates an unstructured attitude, *believes-white*, that the speaker has towards snow (2008b: 56-7). But this is precisely what expressivism says about moral predicates. The solution, therefore, is to posit some basic and general structured noncognitive attitude that moral sentences express. On this picture, moral predicates contribute to the content of the general attitude and not just indicate that attitude (which they do also). Schroeder suggests we call this attitude *being for*.[7]

Being for is a practical attitude that agents bear towards properties. The relevant kind of properties are things that we can do, actions (broadly construed) that can be expressed in English using gerunds. The idea is that for every predicate, there is a corresponding property of the relevant sort that an agent who assertorically uses that predicate 'is for'. As what the agent is for is something the agent can do, we can think of the properties involved as those things that an agent is disposed to do when in the mental state expressed by the use of the corresponding property.

So suppose that moral discourse is constitutively bound up with the practice of blaming. This is of course over-simplistic, but supposing it were right, we might then say something like the following. The moral predicate 'is wrong' corresponds to the property *blaming for*; the speaker of (4) expresses the mental state of *being for blaming for stealing*. More generally, for any predicate F, there is a corresponding relation R_F such that 'x is F' expresses the mental state of being for bearing R_F to x; that is, $F(a)$ expresses FOR(bearing R_F to a) (2008b: 58).[8] As with belief, we now have a uniform structured

[7] Hence bi*for*cated attitude semantics.

[8] More precisely, the properties are given in terms of lambda-abstractions rather than gerunds, so for any predicate '$F(x_1,...x_n)$', there is a corresponding relation denoted by $R_F(z,x_1,...x_n)$ so that if '$a_1,...a_n$' are singular terms denoting $o_1,...o_n$, then $F(a_1,...a_n)$ expresses FOR($\lambda z(R_F(z,o_1,...o_n))$), where 'z' might naturally understood to denote an agent (2008b: 78).

attitude type that is expressed by all moral predicates. Moreover, each moral predicate will correspond to some content, where this is the kind of property given above.[9]

Note, also, that the content of the attitude in question corresponds to a natural property, which is picked out by a descriptive term (e.g. blaming). So the account avoids any metaphysical commitment to moral properties. Counter-intuitively, the content of moral thought on this picture is a sort of descriptive content. This is not to say that we have reduced moral terms to descriptive content in any way, however. While moral terms contribute to the content of the thought that they express, they are not referential, and so do not *refer* to these contents. As we saw above, they express a structured attitude. The nondescriptive nature of moral thought is then captured by the practical attitude of being for—moral thought "is tied to action, in the broadest possible sense. When you are for something... that is what you do." (2008b: 84) In other words, the functional role of a state of *being for* is to lead the agent to acquire the property that the state is 'for'.

Schroeder then goes on to show how an account along these lines can explain the compositionality and logic of moral thought up to the complexity of predicate logic. The logical relations between moral sentences are explained (inter alia) in terms of the inconsistency relations between the contents of the thoughts involved. As the contents of moral thoughts are descriptive, all that is needed is for gerunds to be governed by the relevant logical properties (2008b: 68). Intuitively, this seems to be so. For example, kissing and telling entails kissing and entails telling, stealing and not stealing are inconsistent, and so on.[10] In the next section we will see there is reason to doubt this, but let's assume that it's right for now.

BAS was first motivated by contrasting being for to belief. But the contrast now gives rise to a problem. This is the problem of mixed sentences, which include both moral and descriptive predicates. It is unclear what sort of state the expressivist will say that this expresses—is it belief or being for? Moreover, given that belief and being for have quite different contents (properties and propositions), it is unclear how one state could have both kinds of content.

The solution is to analyse one attitude in terms of the other so that both moral and descriptive sentences express a single type of mental state. Schroeder thinks that to analyse being for in terms of belief is just to abandon nondescriptivism, and so the expressivist must analyse belief in terms of being for. As being for takes properties as its

[9] 'Correspond' not in any referential or representational sense, but in the sense that the use of any such predicate *expresses* a mental state with the corresponding content. Schroeder uses the terminology 'semantic value' to denote this relation.

[10] Once quantification and lambda-abstraction have been introduced, Schroeder provides a more rigorous treatment of the logic of the properties that being for takes as its objects (2008b: 80-2). There are also complications that arise due to the possibility or otherwise of moral dilemmas (2008b: 71-4). I ignore these issues here.

object, what is needed is a general kind of property that is distinctively involved in descriptive belief. To this end, Schroeder introduces the notion of *proceeding-as-if*:

> I say that believing that *p* is being for proceeding as if *p*. What we need, of course, in order to analyze believing that *p* as being for something, is some relation to *p*, so that we can construct a property out of *p* that is something that the believer can be for. This property is something that the believer *does* with respect to *p*, or at least is *for* doing with respect to *p*, when she believes that *p*. And so I call this relation *proceeding as if*. On my best gloss, to proceed as if *p* is to take *p* as settled in deciding what to do. So being for proceeding as if *p* is being for taking *p* as settled in deciding what to do. Assuming that being for has the motivational property that someone who is for α will tend to do α, other things being equal, it follows that someone who believes that *p* will tend to proceed as if *p*, other things equal. That is, she will tend to treat *p* as settled in deciding what to do. (2008b: 93-4)

So the claim 'grass is green' expresses the attitude of *being for proceeding-as-if grass is green*. While Schroeder is not entirely confident with this analysis of belief, in order to give it some plausibility, he notes that it is what believers in fact do, "so this is not, after all, a crazy thing to say about belief." (2008b: 94)

Here ends the simplified sketch of Schroeder biforcated attitude semantics.[11] In the original account, moral sentences are nonpropositional, as they express mental states that are nonpropositional. Despite the fact that descriptive sentences also express states of being for, they remain propositional, as the states of being for take properties that involve propositions—'proceeding-as-if' takes propositions in its complement position. The account therefore must be adapted in some way if we are to accommodate moral propositions.

Propositions are the objects of our attitudes. If moral and descriptive belief consists in states of being for, then the possible objects of our attitudes are given by the class of properties that being for can take as its object. So why not simply identify propositions with these objects? Propositions can then be individuated in the following way:

> to think about what properties are involved in the proposition that grass is green, that murder is wrong, that Max is in Albequerque, or that if she says yes, we'll be engaged, we have to think about what someone who believes that grass is green, that murder is wrong, that Max is in Albequerque, or that if she says yes, we'll be engaged, is motivated to do, other things being equal. (2013: 425)

[11] It turns out that a new negation problem can be created for descriptive belief on this proposal. To solve it, states of being for have to take pairs of properties as their objects (hence *bi*forcated attitude semantics), where these are entailing properties with a major property at least strong as the minor. So the descriptive belief that *p* = <FOR(pai *p*)*, FOR(¬pai ¬*p*)>, where * denotes the major attitude (2008b: 99).

As the properties in question are descriptive, they will plausibly have the right kind of compositional and logical properties distinctive of propositions (though, again, see below).

Moreover, note that the propositions for any domain of discourse will be able to explain in part what is distinctive about that domain. This was seen with the example of morality and blaming above. This is a welcome aspect, as it is common to think that what is distinctive about a domain of discourse is explained by what that discourse is *about*, i.e., in terms of the propositions of that domain. However, this thought has now been given a nondescriptivist twist. For 'what the discourse is about' is not explained referentially in terms of a basic representational attitude (belief), but rather in terms of the distinctive properties that the practical attitude of being for takes as its objects for domain in question.

Importantly, this still leaves room for descriptivism. Any descriptive belief will consist partly in a relation to an essentially representational entity, *viz.*, a representational content. Schroeder leaves the question of what a representational content might be open. However, he suggests that many of the traditional conceptions of propositions might be suitable to play this role. For example, representational contents might be understood to be sets of metaphysically possible worlds, or as structured Russellian propositions (2013: 426n).

2.2.2 Some Problems For Schroeder

One of the key notions introduced in BAS that was that of proceeding-as-if. This notion, together with that of representational content, was vital for differentiating descriptive beliefs and propositions from other kinds of beliefs and propositions. In the presentation of the general framework, these two notions correspond to 'B' and 'C' respectively. Thus, proceeding-as-if is not an inessential detail, but essential in differentiating descriptivism and nondescriptivism under BAS. However, the notion gives rise to a number of difficulties that stem from how the notion interacts with representational contents in order to comprise descriptive propositions. So the problems lie at the very foundation of Schroeder's approach to propositional nondescriptivism.

The main problem with using proceeding-as-if in order to pick out the class of descriptive beliefs is that the notion seems to apply just as much in the moral case as it does in the descriptive. It was seen above that whatever we think of proceeding-as-if as an *analysis* of belief, it does look like a plausible *description* of what believers do. For example, if I believe that there is gin in the fridge, then it seems true to say that I am for proceeding-as-if there is gin in the fridge. This roughly means that I take this as settled in deciding what to do. So if when deciding whether to have a drink, I take it as settled that there is gin in the fridge; when deciding whether I need to buy any drinks, I take it as settled that there is gin in the fridge; and so on.

Suppose now that I believe that drinking alcohol is morally wrong. Is it not true to say that I am for proceeding-as-if drinking alcohol is morally wrong? It certainly seems like something that I can take this as settled in my decision making. For example, if I am deciding whether to do something wrong, then I will take it as settled that drinking alcohol is wrong. Moral beliefs and descriptive beliefs share a common functional role in connection with how beliefs and desires can combine to create intentions. For example, my belief that the gin is in the fridge together with my desire to have a drink may result in my forming an intention to drink a glass of gin. The previous example shows that this is also true of moral beliefs; indeed, Schroeder himself makes this point (2013: 414). However, proceeding-as-if was meant to be the property distinctive of descriptive belief—definitionally so, in fact. As it plausibly applies to all kinds of belief and not just a particular class of beliefs, either moral belief must be descriptive, or we have failed to pick out what is distinctive of descriptive belief.

In his original BAS, Schroeder acknowledges that the locution 'being for proceeding-as-if' can be used to truly describe to all kinds of belief, not just descriptive. To maintain the descriptive/nondescriptive distinction, he makes the following suggestion:

> [So] anyone could pick and choose which sentences, 'P', [i] qualify as ordinary descriptive sentences, on the grounds that it is possible to understand what it would be for it to be the case that P, and [ii] which require instead an analysis of what it would be to proceed as if P... On this picture, the basic expressivist idea is an idea about *what it is* to proceed as if murder is wrong. (2008b: 156 emphasis added)

So the idea is that whether a claim is descriptive or nondescriptive depends on what 'proceeding-as-if p' *consists in* for the claim in question. Thus one might maintain that while it is true that believing that murder is wrong is being for proceeding-as-if murder is wrong, "[t]o proceed as if murder is wrong... *just is* to blame for murder." (2008b: 155)

Rather than coming to rescue the amended BAS, however, this response further confounds it. In the earlier account, moral discourse is still being understood as nonpropositional. So where Schroeder claims that an analysis of descriptive belief consisted in understanding "what it would be for it to be the case that P", this presumably means to apprehend a proposition, where this only applies in the descriptive case. But this cannot be used to draw the distinction between descriptive and nondescriptive belief on the adapted account, as we are allowing for nondescriptive propositions. What this shows is that it is highly intuitive to treat proceeding-as-if as taking *propositions* for its object. Indeed, this is what the original account of proceeding-as-if seems to maintain. However, this cannot be the case in the adapted account, as it applies just as much in the moral case as it does in the descriptive case. And this just means that we have failed to pick out any distinctive characteristic of descriptive belief that distinguishes it from nondescriptive belief.

Furthermore, if it did take propositions for its object, this would actually result in there being two distinct belief states in descriptive cases where intuitively there is only one. For any state of being for proceeding-as-if *p*, '*p*' would be ambiguous between the proposition and the representational content that figures in that proposition. As it is part of our hypothesis that these are distinct, there would therefore two distinct beliefs corresponding to any descriptive claim.

Perhaps one might make the following reply. In the everyday sense of the notion, 'proceeding-as-if' does apply to all beliefs, and does take propositions as its object. However, 'proceeding-as-if' is here being introduced as a piece of theoretical terminology to pick out whatever property is distinctive of descriptive belief. Therefore, it is irrelevant that there is a sense that applies to moral beliefs, for this is not the sense in question. Moreover, let it be defined by stipulation that this property takes only representational contents as object. So while the exact nature of this property is yet to be specified, nothing has been said to rule out there being such a property that is distinctive of descriptive belief in the right sort of way.

There is no inconsistency in maintaining such an account. The problem, however, is that we have been given no reason to think that any distinctive property actually exists. *A fortiori*, we have no reason to accept that there is any real distinction to be drawn along these lines between descriptive and nondescriptive belief. Any recognisable sense of 'proceeding-as-if' applies just as much to moral belief as it does to descriptive belief, and it cannot simply be assumed that there is some sense that does mark a real distinction. Thus, the account fails to provide the resources to maintain the distinction between descriptive and nondescriptive beliefs other than by (unmotivated) fiat.

Does the foregoing discussion tell us anything about the prospects of the general framework itself? As already mentioned, I think that it is far too abstract to draw any strong conclusions one way or the other concerning its plausibility. However, the onus is surely the proponent of such an approach to show that it is viable. It cannot simply be assumed that there is a real distinction between descriptive and nondescriptive beliefs in the form of (1) and (2). The adapted BAS in the end failed not simply because it contained false or implausible assumptions (though this may also be the case), but because it failed to properly distinguish between descriptive and nondescriptive beliefs. So in the end, the account failed not just because it was false, but because it was not a genuine instance of the general framework. At least, except by implausible stipulation. One might take the nature of this failure as grounds for pessimism that the approach can ultimately be made to work.

What the failure came down to was providing an adequate account of what takes the place of *B* in (2), the analysis of the structure of descriptive belief. From the general account alone, it is not at all obvious what sort of thing this could be. (Although in the end it fails, it does seem to fit BAS rather nicely; one suspects that, despite express

statements to the contrary, the framework was advanced with BAS in mind.) Notwithstanding the above objection, however, we see that there are very strong constraints as to the kind of entity that could be descriptive propositions on this picture.

First, notice the constraints that are given by (2) itself. It has already been seen how it is unclear what kind of thing is meant to be denoted by 'B'. Further, consider the constraints imposed by any account of propositions. Particularly, propositions must have the right logical and compositional structure. It was seen that Schroeder captures this by the logic of gerunds. In the descriptive case, this requires the assumption that proceeding-as-if *p* and proceeding-as-if not-*p* are inconsistent. This is a plausible assumption, but it is worth highlighting that it is an assumption nonetheless, as even this fundamental feature of the account might be challenged. Moreover, the constraints are stronger than simply this, as we need to guarantee not just any inconsistency, but *logical* inconsistency for conflicting beliefs with contradictory *pai*-propositions.

Further, it turns out that the kinds of properties posited as propositions by BAS do not have the right logical properties in many contexts more complex than predicate logic, such as embedding under tense and modal operators (2008b: 169-72). Recall that the state of *being for* motivated one to have the property one is for. If we understand tense-operators to operate on propositions, they will operate likewise on the relevant properties. However, this would implausibly result in being motivated to, say, have a property in the past. The general point is that propositions have logical and compositional properties across a wide-range of cases, and any account of propositions will need to meet these constraints.

Traditional accounts of propositions are structured exactly to have these properties. Schroeder's account, however, begins with another kind of constraint—the structure given in (2)—and then proceeds to accommodate the other properties of propositions from thereon. The only attempt, BAS, manifestly fails in this respect. It seems that the logical and compositional properties of propositions are the sorts of features that we should be concerned to accommodate from the outset, not downstream from a more fundamental concern, such as introducing the requisite structure of descriptive and nondescriptive mental states.

I see no reason to rule out the actualisation of an adequate theory based on Schroeder's general framework. However, the aforementioned constraints do pose a serious obstacle to any such theory, and little reason has been given to think that the approach can ultimately be made to succeed. Given that there are other candidate approaches to developing a propositional nondescriptivism other than Schroeder's, I submit that time will be better spent examining these other approaches, rather than puzzling over what possible sort of entity could play the required role of descriptive propositions in Schroeder's framework.

In the remainder of this chapter, another expressivist approach that distinguishes propositions and representational contents will be examined. This approach advances

both a different conception of what propositions are, as well as a rather account of how descriptive beliefs (and, as we shall see, nondescriptive beliefs) are related to representational contents. However, I will argue that this approach likewise fails to adequately maintain the required distinction between the descriptive and the nondescriptive case.

2.3 Ridge's Ecumenical Expressivism

Michael Ridge's Ecumenical Expressivism is a 'hybrid' theory of moral thought and language that combines elements of descriptivism and nondescriptivism. The theory is nondescriptivist in that the contents of moral sentences and beliefs are not representational contents. More particularly, it is an expressivist theory that explains the contents of moral claims in terms of the mental states expressed by such claims. What makes the view hybrid is the claim that the mental states expressed by moral claims are relational states, that have both a descriptive and nondescriptive component. In particular, the state consists in a representational belief being appropriately related to a nonrepresentational 'normative perspective'. (It is also possible for hybrid theories to be descriptivist, and so it remains to be seen in exactly what sense Ecumenical Expressivism is a version of nondescriptivism.)

As with Schroeder, Ridge hopes to provide a nondescriptivist account of moral propositions. Rather than beginning with a nondescriptivist conception of propositions and building the account from there, Ridge begins with a general account of moral discourse and incorporates moral propositions into that account. So while our main concern here is to examine and assess Ridge's account of moral propositions, it will be necessary to first provide a brief outline of the theory as a whole. Also like Schroeder, Ridge aims to accommodate moral propositions by distinguishing between propositions and representational contents. However, as moral beliefs are hybrid and so always involve a descriptive belief component, it follows that moral beliefs always involve a representational content. Whether a belief is descriptive or not depends on the way in which the representational content interacts with the belief. Summarily put, if the content of the belief just is a representational content, then it is descriptive; if it is not, it is nondescriptive.

While Ridge manages to avoid some of the problems that Schroeder's approach encountered, it will ultimately be rejected. Similarly to Schroeder, however, Ridge's account of moral propositions fails to utilise the distinction between propositions and representational contents to adequately distinguish between descriptive and nondescriptive thought. Unlike Schroeder, the problems for Ridge arise from the nondescriptive rather than the descriptive case.

2.3.1 Ecumenical Expressivism and Moral Propositions

Before outlining Ecumenical Expressivism, it should be noted that the version presented here is actually a variant on Ridge's own account. Ridge argues that the central nondescriptivist thesis should not be a semantic thesis about the contents of moral claims or thoughts, but rather a metasemantic thesis that explains why moral claims and thoughts have the contents that they do. The presentation here, however, will forego formulating the theory in metasemantic terms. While this might seem like an unjustifiably large change to make, I believe that there is good reason for doing so.

First, the distinction does far less work in Ridge's theory than it might initially seem. Indeed, once the commitments of the theses are worked out in more detail, the formulations of the main claims of Ecumenical Expressivism put forward by Ridge and the variants put forward below differ little in substance.[12] Second, perhaps more importantly, there a number of issues that arise from Ridge's metasemantic formulation that can easily be avoided if formulated in more standard terms. Moreover, the issues that arise have nothing much to do with Ecumenical Expressivism, but with placing metaethics at the metasemantic level in general. These issues will be explored in depth in the next chapter. While the discussion there will not focus on Ecumenical Expressivism, the same arguments apply *mutatis mutandis*.[13] So it is both possible and desirable to avoid these problems in our discussion of Ecumenical Expressivism. For readability, however, I will refer to the presented account as if it is Ridge's own presentation.

With these caveats in place, we can now turn to the theory itself. Ecumenical Expressivism aims to account for not just morality, but practical normativity more generally. Vaguely stated, practical normative judgment aims to settle "the thing to do" and the "thing to intend", and normative claims (or better: assertoric utterances of moral normative sentences) express these judgments (2014: 19).[14] There are a number of predicates that are distinctive of practical normativity, such as evaluatives ('good', 'bad'), directives ('ought', 'must'), and reason claims ('x is a reason to φ'). To provide a unified and fully general account of normative thought and discourse, Ridge suggests that all normative claims can be analysed into an equivalent claim about where the action stands on any acceptable standard of practical reasoning (2014: 40). A standard is broadly speaking a rule or policy which can be used as the basis of judgment or decision; 'acceptable' is a primitive normative term.

[12] For reasons of space, I cannot show that this is the case here. I simply ask for the reader's trust on this matter.

[13] Some worries about this distinction that are more specific to Ridge's account can be found in Alwood (2016).

[14] 'Judgment' here is being used as a theoretically neutral term to describe whatever mental state is expressed by normative claims.

For example, the moral predicate 'good' can be analysed in terms of being ranked highly by any acceptable moral standard. So to claim that 'Giving to charity is (morally) good' is to claim that giving to charity is highly ranked by any acceptable moral standard. While this is provides a general analysis of moral claims, it is important to emphasise that the analysans is still itself within the normative domain. This is because it employs the normative notion of an *acceptable* moral standard. The idea at this stage is simply to provide a unified account that 'locates' normative thought and discourse. It is not to provide a nonnormative reductive analysis.[15]

The standards-based account of morality is silent on whether moral discourse is descriptive or nondescriptive. As an expressivist theory, Ecumenical Expressivism proceeds to ask what kind of mental state is expressed by the conventional use of moral sentences. Its answer is that moral sentences express moral judgments. Moral judgments are relational states consisting of (i) a (noncognitive) normative perspective, and (ii) a representational belief. Further, (i) and (ii) must be connected in the right way.

All moral claims can be paraphrased into claims about acceptable moral standards. Standards are understood as rules or principles used as a basis for judgment. As agents, we have a basic nondescriptive attitude of *accepting* rules or principles, where this is to be disposed to issue the relevant prescriptions (2014: 111). Moreover, being *treated* by an agent as a standard is more basic than *being* a standard of reasoning (2014: 40). A normative perspective is "a set of relatively stable self-governing policies about which standards to reject and accept." (2014: 115) So to adopt a particular normative perspective is to take a particular practical stance with respect to decision making. Moreover, the relevant standards are understood as 'ultimate', which means they are not based on any more fundamental standards and provide a complete guide to action (2014: 116-17). As normative perspectives are understood maximally relative to an agent at a particular time, we can understand an agent's moral perspective as the subset of their normative perspective, *viz*., the subset concerning distinctively moral concerns. This accounts for the noncognitive component of moral judgment.[16]

The cognitive component consists in a descriptive or representational belief, the contents of which are related in a particular way to the agent's normative perspective. Ridge provides the following example (2014: 119). Consider the following claim:

---

[15] The standard-based context-sensitive analysis is presented first as a *semantics* for moral claims. Expressivism is then understood as a metasemantic thesis about what grounds the semantics. However, insofar as the analysis aims to provide a fully general and unified account that locates the practically normative, we need not be too concerned about the exact nature of the analysis, as in any case, the paraphrase will still hold whatever account we give. The important point is that the analysans is itself normative. It is instructive to compare this to Schroeder (2007) who suggests that a reduction of the normative should consist in two stages: first, a reduction of all normative notions to a fundamental normative notion (for Schroeder, reason claims; for Ridge, acceptable-standards claims); second, a reduction of the fundamental normative notion to a nonnormative notion.

[16] On norm-acceptance and plan-acceptance as noncognitive, see Gibbard (1990, 2003).

(5) 'X is good as an end'

This can be analysed into the following equivalent claim:

(6) 'X would be highly ranked as an end by any acceptable ultimate standard of practical reasoning'

This judgment is analysed as a hybrid state constituted by the following pair:

(6_N) A normative perspective.
(6_R) The belief that X would be ranked highly as an end by any admissible ultimate standard of practical reasoning.

These two states are connected in virtue of 'any admissible ultimate standard' being indexed to the agent's normative perspective. So the standards referred to in (6_R) are those standards not ruled out by the agent's normative perspective. It follows that normative judgments can be multiply-realised, as there need be no uniquely fixed content of the representational belief involved in the judgment (at least in abstraction from any particular normative perspective). Furthermore, an agent must necessarily have (6_N) and (6_R) in order to be in the mental state expressed by (6), as it is constitutive of moral judgment to be a relational state in this way.

Given the descriptive component of moral judgment, one might naturally raise the following question: in what sense is Ecumenical Expressivism a form of nondescriptivism? The difference between descriptivism and nondescriptivism can be stated as a difference in how the content of a moral claim is related to the content of the representational belief component of moral judgment. It was noted above that there are also descriptivist hybrid theories, and it is helpful to first examine what would make a hybrid view descriptivist. Ridge makes the following suggestion. A hybrid view is descriptivist just in case the content of the claim '*p*' is identical to the content of the representational belief component of the relational state expressed by '*p*' (2014: 80). So it might still be the case that moral judgments necessarily involve a noncognitive attitude, or even constitutively involve a noncognitive component.

What is required for descriptivism is that the content of the claim expressing a moral judgment is for its content to be identical to the representational content of the descriptive belief component of the judgment. For this would have the consequence of being committed to moral representational contents. And if we associate representational contents with metaphysical commitment, then this draws the line between descriptive and nondescriptive belief just where it should do.

As Ecumenical Expressivism does not meet this requirement, it is not a form of descriptivism. Consider claim (6). One way in which we can gloss the content of the

claim is in a deflationary manner. Here, we can talk about the truth-conditional content of (6), as speakers will be able to grasp the truth-conditions of (6) in the minimal sense that they can grasp the relevant platitudes. Grasping the relevant platitudes, however, does not amount to knowledge of the descriptive conditions under which the claim is true, i.e., grasping a representational content. The analysis of moral claims yielded a normative analysans. So while the representational belief is a necessary component of the mental state expressed by (6), the content of (6) is not identified with the descriptive content of (6_R). Rather, (6) has an irreducibly moral content.

So while the primitive normative notion of an 'acceptable standard' might indicate the presence of a descriptive belief involving the (representational) content of an admissible standard, acceptability does not in any way reduce to admissibility. Rather, what one takes to be acceptable is determined by one's normative perspective; admissibility then refers to the standards deemed acceptable by the agent (hence is representational). Moreover, because 'acceptable' is explained nonreferentially, the theory avoids all of the problems associated with descriptivism about moral terms outlined in the previous chapter. In other words, Ecumenical Expressivism does not commit us to there being any moral representational contents. So how can we understand moral propositions?

Here, Ridge appeals to Scott Soames' theory of propositions as cognitive event types. The theory is largely motivated by a number of problems that arise for traditional theories of propositions. We need not dwell on such issues here. What is important is the contrast in explanatory priority between traditional theories and Soames' alternative. Consider that, on the one hand, we describe *propositions* as representing the world as being a certain way. However, we also describe *agents* as representing the world as being a certain way, such as when we *judge* that *p*. These two ideas seem intimately connected. The question is whether our cognitions are representational in virtue of their bearing a relation to propositions, or whether propositions are representational in virtue of our representational activities.

Traditional approaches maintain that the representational nature of propositions is explanatorily prior. On this picture, propositions are somehow intrinsically representational. When an agent represents things as thus or so in thought or language, its representationality is derived from the proposition or propositions that the thought or sentence is related to. Soames argues that this commitment is the source of many of the traditional problems that arise for theories of propositions (2010). In light of this, he urges that we should reverse the order of explanation. That is, we should understand the fact that an agent represents things as thus and so as being the conceptual basis of representation. In other words, propositions are representational in virtue of their connection to agential representational activity (Soames 2014: 96).

On this picture, the primary instances of representation are concrete events, where these are certain cognitive performances or actions or operations. The basic building block is the notion of entertaining a proposition. What entertaining amounts to in each case depends on the nature of the proposition in question. The simplest case of entertaining a proposition is predication, understood as a cognitive act. For example, consider the concrete act of an agent predicating redness of a ball. With these concrete events taken as basic, propositions are then understood as the minimal event type that corresponds to such concrete events. It is minimal in the sense that it is what is representationally common to any arbitrary agent predicating redness of a ball. Thus, a proposition derives its representationality from the representationality of concrete events. We then say that a proposition is true just in case things are the way that the proposition represents them as being (2014: 96).

Of course, propositions are often more complex than subject-predicate form. So more complexity needs to be introduced. Generally, the various ways of entertaining a proposition can be given as follows:

> The simplest are those in which properties are predicated of objects. Complex propositions may involve other operations such as conjoining, disjoining, and negating properties or propositions, as well as operating on, for example a two-place relation R to form the reflexive, one-place property *self-R-ing*. They may also involve applications of functions to objects, or to properties (or propositional functions). In addition, some complex propositions involve the ascription of higher-order properties to lower-order properties (or propositional functions) as in quantification. Propositions of any sort may also be arguments of further predications, which we find in modal propositions and attitude ascriptions[...] (2014: 99)

And so on. The basic idea is the same as with predication. For whatever way a proposition represents something as being, we find the corresponding mental act that represents the world as being that way.

With the basic cognitive activity of entertaining a proposition in place, propositional attitudes can be introduced as attitudes that involve entertaining a proposition in some way (2014: 97). For example, to judge that x is F is to affirm or endorse that predication in thought and in reasoning. To believe that x is F is to be disposed to judge that predication. To assert that x is F is to commit oneself and to treat oneself as entitled to the predication through a communicative linguistic act. To know that x is F is something like to be justified in believing the predication when it is true. And so on.

It is not a huge step to see how this account of propositions might dovetail with an expressivist approach to moral discourse. If one identifies propositions with cognitive event types, then one might think that to provide an account of the content of a claim in terms of the type of mental state it expresses more or less is just to provide an account of the content of a claim in terms of the proposition it expresses. This is what Ridge endorses. The only difference is that, as with the notion of belief, 'cognitive event type'

should be understood in a similarly broad sense, to cover both moral beliefs and descriptive beliefs (2014: 128).[17] In this way, Ridge hopes to identify types of moral thoughts as instances of entertaining moral propositions.

Granting Soames' view of propositions, simply recommending that the relevant kinds of cognitive event type comprising propositions should be broadened is not itself sufficient to guarantee the propositionality of moral thought and language. This is because the cognitive event type needs to display the right sort of structure in order to adequately play the role of propositions. In the descriptive case, it is easy to see that entertaining representational propositions will meet this requirement. This is because for any proposition *p*, there is an isomorphism between the cognitive event type of entertaining *p* and *p* itself. In the nondescriptive case, however, it cannot simply be assumed that the cognitive event type corresponding to the relevant kind of judgment will possess the properties required of it. It is in answer to this question that the motivation for developing a hybrid account becomes apparent. The basic strategy in accounting for the propositionality of moral discourse is to 'offload' the work required onto the representational component of moral judgment.

Beginning with compositionality, any account of moral thought needs to provide a fully general and recursive account of how any arbitrarily complex moral sentence gets its content. Ridge argues that this work can be offloaded to the representational belief component of a proposition in the following way (2014: 120). Consider the claim, 'If pleasure is good as an end, then Socrates sought pleasure'. Like other moral judgments, this is a relational state involving a noncognitive normative perspective related appropriately to a representational belief. For an agent who makes this judgment, the normative perspective is the same as that which would be involved in an atomic moral judgment (normative perspectives are maximal at a time). The content of the representational belief is then 'If pleasure would be highly ranked as an end by any admissible standard of practical reasoning, then Socrates sought pleasure', where 'admissible standards' refers to those not ruled out by the agent's normative perspective. More generally:

> Take any logically complex sentence S in which a normative predicate is used (and not just mentioned)[...] S expresses (a) a normative perspective, and (b) the belief s*, where s* is what one gets when one takes 'S' and replaces all occurrences of normative predicates in 'S' with the obviously corresponding phrases about what any admissible standard would be like in the relevant way. (2014: 145)

[17] I follow Ridge here in using 'cognitive' to be more or less synonymous with 'mental'. One might understand 'cognitive' here to simply *mean* representational; perhaps Soames does. However, I have opted to stick with using 'cognitive' for readability, though nothing substantive turns on this. If the reader so wishes, she may replace each use of 'cognitive' with 'mental'.

In this way, we "thereby get a fully general and recursive" account "of how each claim gets its propositional content." (2014: 130) This is achieved by having systematically explained the kind of mental state type expressed by any arbitrary complex moral claim.

It also needs to be shown that moral judgments have the right sort of logical properties. For example, it needs to be shown how logical validity applies to moral discourse. As before, Ridge explains these features of by offloading the work to the representational aspect of moral propositions. All that is then required is to 'psychologise' the orthodox approach to so that it can be situated into the broader expressivist framework. To this effect, Ridge defines validity as follows:

> An argument is valid just in case any possible believer who accepts all of the premises but at one and the same time denies the conclusion would thereby be guaranteed to have inconsistent beliefs, where this remains true on any acceptable substitution of the non-logical terms of the argument. (2014: 156)

An 'acceptable substitution' is a substitution of terms of the 'same semantic kind'. Note that Ridge requires an account of validity in terms of the commitments of a believer, because the representational beliefs contained within moral judgments necessarily contain the indexical predicate 'any admissible standard'. So without indexing validity to a particular speaker in this way, there would be no guarantee of validity, as different people will judge according to different standards.

With validity so defined, it can then be shown how validity can be applied to moral arguments. Consider the following argument:

(P1) If one lives in a glass house, one ought not to throw stones.
(P2) John lives in a glass house.
(C) So, John ought not to throw stones.

To see whether the argument is valid, we need to see what results from accepting (P1) and (P2) while denying (C). More particularly, we need to see whether accepting the premises and the negation of the conclusion results in having inconsistent representational beliefs.

(P1) and (P2) are normative claims that involve a representational belief indexed to a normative perspective. Thus, someone who accepts the premises but denies the conclusion is committed to (something like) the following three claims:

(P1′) For any agent *A*, if *A* lives in a glass house, then *A*'s throwing stones is ruled out by any admissible standard of practical reasoning.
(P2) John lives in a glass house.

(¬C) John's throwing stones is not ruled out by any admissible standard of practical reasoning.

(P1′), (P2) and (¬C) are all representational beliefs. As such, we can simply look to their contents to see whether they can together be consistently maintained by any possible agent. It seems clear that they cannot, for any two of the beliefs held together rule out the remaining belief. As such, the original modus ponens argument comes out valid.

Note how it seems to avoid one of the problems Schroeder faced, concerning the logic of propositions. His problem, recall, was that by identifying propositions with properties, he needed the relevant properties to have the same logical and compositional relations as propositions. It was argued that he failed in this respect, and that this aspect of his account placed severe constraints on what the right kind of properties could be. However, representational contents here have exactly the right kind of form. Indeed, in the case of representational belief, the representational content just is the propositional content. As such, no parallel problem arises.

In summary, Ridge aims to accommodate moral propositions within Ecumenical Expressivism by identifying the proposition *p* with cognitive event type of entertaining the proposition that *p*. In the simplest case, to entertain *p* is to predicate *F* of *x*. However, as moral thoughts are relational states containing a noncognitive aspect, Ecumenical Expressivism must explain how states other than simple representational states can play the role required of propositions. Ridge's general strategy is to offload this work to the representational beliefs that in part constitute moral judgment. In this way, it hopes to explain the propositionality of moral language from a nondescriptivist viewpoint.

2.3.2 A Lack of Entertainment: A Problem for Ecumenical Expressivism

Soames' account takes representational cognitive acts to be the building blocks from which to construct a theory of propositions. He is explicit in understanding propositions as essentially representational. Indeed, propositions are *individuated* in terms of what is minimally representationally common to all possible instances of the proposition (Soames 2014: 96). Ridge, however, cannot accept this. If he did, then moral judgments would be wrongly identified with the representational belief component of a moral judgment. This is because the hybrid state and the representational belief have exactly the same representational content—the hybrid state has no representational content over and above that of the descriptive belief component. So the criterion of identity for propositions must differ depending on whether the proposition is descriptive or nondescriptive.

This might not seem too worrying. In many ways, this simply mirrors what has already been said about moral belief. While there are robustly representational beliefs, there are also deflationary beliefs. The former are individuated by their representational

content. The latter are individuated by specifying the kind of mental state expressed by the corresponding claims. So the way to individuate moral propositions must be to identify the minimal event type that corresponds to the various concrete instances of particular moral judgments.

Recall that for Soames, the most basic building block in his theory of propositions is the notion of *entertaining* a proposition. To entertain a proposition is to perform is a cognitive achievement, it is something one *does*. In the simplest case, to entertain the proposition that *x* is *F just is* to *predicate F*-ness of *x*. We thereby represent *x* as being *F*. Here, the notion of entertaining a proposition is explained in terms of our representational activity—the latter is prior in the order of explanation. Moreover, entertaining a proposition does not here simply mean the propositional attitude that one has when one consciously considers a proposition in consciousness without judging, doubting, wondering, or whatever. Entertaining a proposition *just is* the representational activity one engages in. This is what makes entertaining a proposition the foundational cognitive activity from which to understand all other propositional attitudes. Entertaining a proposition is not contrasted with judging a proposition. Rather, judging a proposition *involves* entertaining a proposition. For example, judging that snow is white involves predicating whiteness of snow, which is to say, involves entertaining the proposition that snow is white.

So a representational proposition is identified with the cognitive event type of entertaining that proposition, which is in turn identified with a cognitive representational act (e.g. predication). This event type is what representational beliefs have as their contents. So what according to Ridge is the kind of cognitive event type that moral propositions are to be identified as being? The relevant class of cognitive event types needs to be broadened to include moral propositions; predicating using a moral predicate cannot be of the same kind of thing as predicating using descriptive predicates. Soames' gives us an account of the latter. What kind of thing plays the analogous role in the moral case?

Ridge provides an account of what it is to 'merely' entertain a moral proposition. Drawing from simulation theory, he suggests that to merely entertain a moral proposition is to have a merely simulated normative perspective paired with the corresponding descriptive belief (2014: 128). However, this fails to pick out the cognitive activity of entertaining a moral proposition in the relevant sense. Recall that entertaining a proposition is something we do when we judge, for example. What is needed is a general cognitive activity that is present across *all* propositional attitudes, and which our propositional attitudes can be defined *in terms of*. Clearly the simulation account will not work here, as it would involve defining judgment in terms of accepting or affirming a simulation of judgment in thought. This gets things the wrong way around. Judgment is conceptually prior to simulation, not vice versa. So simulation of judgment cannot be the cognitive event type that moral propositions can be identified as being.

To be fair to Ridge, it is not clear he intends simulation to play this role. Rather, simulation is suggested as a plausible candidate for what we are doing when we 'merely entertain', say, that stealing is wrong. However, Ridge simply fails to identify *any* cognitive event type that can do the work of entertaining a proposition in Soames' sense. Rather, vague talk is given about 'ways in which' one might entertain a moral proposition (2014: 128). For example, judging, hoping, or 'merely entertaining' (simulating) that stealing is wrong are all ways in which one might entertain that proposition. However, this tells us nothing about what cognitive event types moral propositions actually are. Moreover, given that they are, to at least some extent, different in kind to descriptive propositions, we have not been given any reason to think that there actually *are* any such nondescriptive cognitive events from which to construct a class of cognitive event types corresponding to moral propositions.

So Ridge fails to identify any cognitive event type that moral propositions could be. The situation is worse, in fact, when one considers that the only natural candidate would be entertaining the representational proposition of the belief component of moral judgment. After all, it is this component of moral judgment that explains the propositionality of moral thought. So if one accepts Soames' account of propositions together with a hybrid theory of moral judgment, then there is great pressure to identify moral propositions with the representational contents of the belief component. However, this would be to embrace descriptivism.

Perhaps one might argue against this move by claiming that the theoretical benefits of not making this identification (in other words, the benefits of nondescriptivism) justify positing some other cognitive event type to be moral propositions. Particularly, that cognitive event type that is common across all instances of the proposition being cognised. However, insofar as (i) we have yet to see any possible, let alone plausible, nonrepresentational candidate to play this role, and (ii) there is an obvious representational candidate that can play this role, any such move is implausibly ad hoc and cannot be accepted.

The pressure to identify moral propositions with the representational contents of the belief component not only arises from there being an absence of other candidates. There is some positive reason for making this identification, insofar as one accepts Soames' theory and one agrees with the general account of moral judgment according to Ecumenical Expressivism, modulo its commitment to nondescriptivism. This has to do with the distinction between judgment and belief. In metaethics, 'judgment' is often used as a theoretically neutral term to describe the mental state characteristic of moral thought. The dialectic purpose of this is to not beg any questions as to whether moral judgments are beliefs or some other attitude. This is the sense in which Ridge generally uses the term. However, this terminology is somewhat unfortunate, as there is another sense in which judgments are distinct from beliefs. This is the sense in which while a

belief is a state, a judgment is an event. One has beliefs over time, whereas one judges at a time. A judgment is something that one *does*—it is active in a way that belief is not. Belief is more like a dispositional state, whereas a judgment is a sort of cognitive achievement. The two are intimately connected—as was seen above, one rough way of glossing their relation is that to believe that *p* is to be disposed to judge that *p*. Whatever their precise connection, it is plausible that belief and judgment in this sense belong to two distinct metaphysical categories.

Consider now the following. Ecumenical Expressivism defined moral judgment as a relational *state*. What, then, is the *event* that constitutes moral judgment in the active sense? Ridge suggests the following: "Normative beliefs, then, are dispositions which give rise to normative judgments, where these are the relevant descriptive judgment/normative perspective pair." (2014: 128) This is puzzling as the only active component in moral judgment is the descriptive judgment. Normative perspectives do not have active counterparts in the same way that beliefs do. However, if judgments are events, then how could a moral judgment be pair of which one of its constituents is a state? Perhaps two events might plausibly be conjoined to make a single event, but it is unclear that an event could consist of an event and a state. If anything, this looks like a form of Ecumenical Cognitivism, in which the presence of a normative might be necessary, but the content of the moral claim—the moral proposition—just is the content of the representational belief.

One might make the following reply. While normative perspectives are synchronically static, they are diachronically fluid. Indeed, they must be if we are ever to change our minds about moral matters. Say that after engaging in moral reasoning, I come to acquire a new moral belief that *p*, where 'belief that *p*' is understood as deflationary belief. When I acquire this belief, I make a descriptive judgment that corresponds to the representational belief component. However, I have done more than this: my normative perspective has also changed. So could a normative judgment consist in the occurrence of these two events?

The answer is no, and the reason is simple. While some of my judgments may occur this way, many others do not. It is simply not a necessary condition for moral judgment that my normative perspective changes. Say that it is a long standing belief of mine that investing in nuclear arms is morally wrong. I can in principle make this judgment at any point in which I have this belief. So most of my judgments occur without any change in normative perspective. In fact, it seems possible that I could come to have this belief before making a judgment, for we form many of our beliefs unconsciously. So the reply fails.

Let's take stock. Ridge aimed to adapt Soames' theory of propositions to Ecumenical Expressivism. At the heart of Soames' theory is the notion of entertaining a proposition, which is a technical term to capture the distinctively representational activity that we engage in representing the world in thought. Propositions were then identified with

such cognitive event types. As entertaining in this sense is essentially representational, the notion needed to be broadened to include the cognitive event type that is distinctive of moral judgment as per Ecumenical Expressivism. However, Ridge fails to tell us what the relevant cognitive event type involved in our moral attitudes is. Therefore, he fails to provide an account of what moral propositions are. Furthermore, it was argued that it is unclear what the cognitive event type for moral proposition could be under Ecumenical Expressivism other than the representational belief component of moral judgment. However, if this were correct, then our account of moral thought and language would be descriptivist, as the contents of moral claims and beliefs would be a representational content. Particularly, this would yield a form of subjectivism. This is because the normative term 'acceptable standard' would reduce to the descriptive term 'admissible standard'. So the truth-conditions for moral claims would be a matter of what standards the speaker as a matter of fact does and does not accept.

2.4 Conclusion

In this chapter, we examined whether propositional nondescriptivism could be achieved by distinguishing propositions from representational contents in an expressivist framework. It was argued that Schroeder's attempt failed as he did not provide a plausible account of descriptive propositions, and that Ridge's attempt failed as he did not provide a plausible account of nondescriptive propositions. It seems that we cannot adequately account for one kind of proposition except at the expense of the other. I therefore conclude that this approach should be rejected.

~ 3 ~

Metasemantics and Inferentialism

In this thesis, we have been exploring the suggestion that metaethical nondescriptivism and propositional semantics are compatible. In the previous chapter, we examined two approaches that provided nonrepresentational accounts of what moral propositions *are*. In this chapter, we take a step back to examine in a more holistic way the relation between one's semantic theory and one's metasemantic theory for moral language. The thought is that propositional nondescriptivism might be achieved by formulating nondescriptivism as a metasemantic theory. It is argued that if metaethics is properly conducted at the level of metasemantics, then semantic concerns about propositionality are orthogonal to the descriptivism/nondescriptivism debate. After motivating this claim generally, I examine the semantics and metasemantics for 'ought' developed by Matthew Chrisman. While I argue that Chrisman fails to successfully make the case for metasemantic nondescriptivism, I tentatively suggest a way in which this approach might be vindicated.

3.1 Semantics and Metasemantics

What is the difference between semantics and metasemantics? There are a number of places in which this distinction can, and indeed has, been drawn. My aim in this section is not to argue for a definitive distinction between these two domains of inquiry. Rather, I hope to bring out two quite general but distinct kinds of concern that are characteristic of each domain respectively. This will serve to show how nondescriptivism might plausibly be thought to concern the metasemantics of moral language rather than the semantics. This provides a way of arguing for the compatibility of propositional or truth-conditional semantics and nondescriptivism.

It should be emphasised that neither semantics nor metasemantics displays clear boundaries anyway as to its scope, method, and explanatory goals. Moreover, many issues may be to a greater or lesser extent the concern of both. So there is some scope as to where one chooses to draw the line. In any case, whether the distinction pays dividends or not will depend on the work one puts it to in one's theory. As I will be using the terms, I take neither semantics nor metasemantics to be exhaustive of the kind of inquiry that might go by those names.

Recall Gibbard's dictum and the question it raised. *To explain the meaning of a term, explain the state of mind the term can be used to express*. Yes, very well, but explain *what*?

Philosophical literature is rife with claims to explain the meaning of various expressions. The problem is that there are a whole number of things that one might wish to explain about an expression's meaning. Some aspects about meaning will probably not be relevant to philosophical inquiry—etymological explanation, for instance.[1] For example, perhaps we want to know why a word has a certain meaning rather than another. Or perhaps we want to know why it has any meaning at all. Or perhaps we want an explanation that tells us what it is for an expression to have a certain meaning, or to be meaningful generally. Or perhaps there is some peculiar aspect of some particular expression that calls for an explanation, perhaps in light of other things we know about it. Or perhaps all that is wanted is an explanation of what a term actually means.

So there are in principle any number of things that we might wish to explain in relation to the meaning of a term. Here is a general distinction that one might make between two classes of explanations:

> Insofar as linguistic semantics aspires to 'specify' or 'report' the meanings of sentences and sub-sentential expressions (in some systematic way), philosophy of language could on to tell us *how* or *why* these symbols come to have those meanings—perhaps unearthing more basic or fundamental facts in virtue of which such semantic states of affairs obtain. We take this second sort of inquiry to be paradigmatic, if not exhaustive, of metasemantics. (Burgess and Sherman 2014a: 1-2)

In this way, we can distinguish between 'linguistic' semantics, which is an empirical inquiry into the semantic properties of expressions (compositionality particularly—more on this shortly), and metasemantics, which is a metaphysical inquiry about linguistic semantics and the results that it supplies.[2]

We can further distinguish two basic sorts of metaphysical explanation that might be relevant. First, there are *grounding* explanations, which explain *that in virtue of which* some semantic theory is true of a language. Second, there are *constitutive* explanations, which explain what the facts or properties of some semantic theory *consist in*.[3] Both kinds of explanation can either be directed to explain the meanings of *particular* expressions, or *meaningfulness* more generally. Talking loosely, we might speak of the metaphysics of meaning in terms of the *nature* of some language or fragment. I don't intend to commit to any particular conceptions of what these kind of explanations consist in. However, I hope the sense in which each is metaphysical, at least in a broad sense of the term, is clear.

---

[1] Not to say that etymologies are always irrelevant as explanandum or explanans. The etymological *mode* of explanation, however, does not seem to be in the remit of philosophy.

[2] Burgess and Sherman take their terminological lead from Kaplan (1989: 573-4).

[3] I'll speak loosely throughout of semantic properties and facts. I don't intend to take on any particular commitments by such talk, as will become clear by the end of the chapter.

Most philosophers who use the above distinction or something like it are fairly agreed that grounding explanations (broadly construed) are in the scope of metasemantics. However, there is divergence as to whether constitutive explanations ought to be included. For example, Brandom, distinguishes 'formal semantics' which is concerned "with computing the semantic values of some expressions from those of others", from 'philosophical semantics' which is concerned with "what *kind* of semantic values expressions should be taken to have," as well what grounds particular meanings and meaning in general (2010a: 342 emphasis added).[4] On the other side, Stalnaker distinguishes 'descriptive semantics' from metasemantics, where the former includes saying "what *kinds of things* the semantic values of expressions of various categories are" in addition accounting for compositionality; metasemantics exclusively concerns grounding explanations (1997: 540 emphasis added).[5]

I take it that the main reason for wanting constitutive explanations of meanings to be in the scope of one's metasemantics resides in the metaphysical nature of the question. It seems conceivable that we could agree on all the semantic facts while disagreeing over what these facts consist in. I take that the main reason against including it is the fact that constitutive explanations involve *identity* conditions of semantic facts. We'll return to this question in further detail in the discussion of Chrisman's metasemantics in the later section of this chapter, so I'll leave the issue hanging for now. Contrary to appearances, we will see that this is not merely a terminological dispute.

By contrast to metasemantics, semantics is an empirical inquiry into the semantic properties of natural language. The particular branch of semantics that concerns us here is compositional semantics. The basic idea is to show how the meanings of sentences are a function of the meaning of their parts and their mode of composition. While compositional semantics is likewise explanatory in ambitions, the explanation is in this case not metaphysical. Rather, it aims to explain our ability to speak and understand a language based on only a limited knowledge of vocabulary. It is generally thought that if meaning were not compositional in this way, then the language would be unlearnable, as we would lack any explanation of how we can produce and consume an indefinite number of novel sentences.[6] In this way, compositional semantics explains "how an infinite aptitude can be encompassed by finite accomplishments. For suppose that a language lacks this feature; then no matter how many sentences a would-be speaker learns to produce and understand, there will remain others whose meanings are not given by the rules already mastered." (Davidson 1984a: 8)

So the meanings or semantic values of complex expressions are given as a function of the semantic values that have been assigned to the primitive expressions of the language or fragment under consideration. One prominent way in which to model the semantic

[4] See also Burgess and Sherman (2014a) and Pérez Carballo (2014).

[5] See also Ridge (2014) and Yalcin (2014).

[6] This thought goes back at least to Frege (1963: 1).

value of expressions is in terms of how they contribute to the truth-conditions of the propositions expressed by the sentences in which they figure.[7] Whether the model provides a correct representation of the compositional properties under examination can be tested by seeing if the intuitive truth-conditions of sentences match those predicted by the model. In the sense of compositional semantics being discussed, it is important to understand exactly what the predictive power of the semantics consists in: "Empirical power in such a theory depends on success in *recovering the structure* of a very complicated ability—the ability to speak and understand language." (Davidson 1984b: 24, emphasis added)

That this is the primary explanandum is important to keep in mind because it can otherwise obscure what a truth-conditional semantics need and need not be committed to. For example, suppose we understand sentences to express propositions, the composition of which determines the truth value of the sentence. Using propositions in this way allows us to track commonality of content across a wide variety of linguistic contexts, such as embedding and attitude ascriptions. However, propositions are here understood simply as elements in a formal model. What matters here are their *formal* or *structural* properties.[8]

In other words, for the purposes of compositional semantics, we simply understand propositions in whatever way best represents the formal properties of the compositional properties of sentences, whether this be as a structured entity, a set of possible worlds, a function to possible worlds, or whatever. In and of itself, this saying nothing about the nature or properties of the expressions in question over and above their formal compositional structure. Compare Yalcin:

> It is a platitude that modeling a class of properties is, generally speaking, simply not the same as offering some kind of translation procedure operating on the bearers of those properties, and neither is it the same as giving an interpretation procedure in some intentional sense, or in the sense familiar from artificial languages. This platitude should be respected, even when what is to be modeled are the semantic properties of natural language, and even when the modeling proceeds using formal tools developed originally for stipulating model-theoretic interpretations for artificial languages. One can use model-theoretic tools to model meaning properties without assuming that in doing so, one must be associating expressions with their supposed referents or representational contents. Again, while that kind of gloss may be natural when one is using such tools to *stipulate* referents or meanings in an artificial language (for, say, logic or mathematics), it is question-begging, and must be justified, when the tools are deployed in connection with the empirical study of natural language. (2014: 45)

[7] For simplicitly, I'll ignore the distinction between propositional and nonpropositional truth-conditional semantics.

[8] c.f. Davdison: "we must be able to specify, in a way that depends effectively and *solely on formal considerations*, what every sentence means." (1984a: 8, emphasis added)

The point is that, as an empirical study of natural language, compositional semantics should be able to use model-theoretic tools to model the compositional properties of language without taking on any particular metaphysical commitments—for example, some particular conception of propositions. This is as it should be. As Pérez Carballo notes: "semanticists are not in the business of pronouncing on metaphysical issues. It would be incredible if the viability of current semantic theory as we know it depended on the outcome of a controversial metaphysical dispute." (2014: 136)

The relevance to debates about descriptivism and nondescriptivism should by now be apparent. It is often assumed that nondescriptivism is incompatible with propositional or truth-conditional semantics. However, if semantics is metaphysically neutral in the way being suggested, then it is *prima facie* unclear why these must be incompatible. This is in fact not a new point, and was observed by Davidson in his early work on semantics:

> If we suppose questions of logical grammar settled, sentences like "Bardot is good" raise no special problems for a truth definition. *The deep differences between descriptive and evaluative (emotive, expressive, etc.) terms do not show here*. Even if we hold there is some important sense in which moral or evaluative sentences do not have a truth value (for example, because they cannot be "verified"), we ought not to bottle at "'Bardot is good' is true if and only if Bardot is good"; in a theory of truth, this consequence should follow with the rest, keeping track, as must be done, of the semantic location of such sentences in the language as a whole—of their relation to generalizations, their role in such compound sentences as "Bardot is good and Bardot is foolish," and so on. What is special to evaluative words is simply not touched: the mystery is transferred from the word 'good' in the object-language to its translation in the meta-language. (1984b: 31, emphasis added)

There are two key points to take from this. First, compositional semantics as empirical theory should be able to proceed fairly autonomously, using whatever theoretical tools best predict the semantic values of sentences. Second, propositional and truth-conditional compositional semantics are in principle compatible with nondescriptivism.

'In principle' is an important qualification here. For it might turn out that no nondescriptivist theory could be made to fit the results of the semantics. This depends *both* on what the results of the semantics are, and how the nondescriptivist theory is formulated. The point is that there is no reason to rule this out from the outset. Let us now turn, therefore, to examine Chrisman's semantics and metasemantics for 'ought'.

3.2 Chrisman on the Meaning of 'Ought'

As well as being perhaps the most developed nondescriptivist theory that utilises the distinction between semantics and metasemantics, Chrisman's (2015) theory is also noteworthy for its inferentialist rather than expressivist articulation of nondescriptivism. Chrisman develops a truth-conditional semantics for 'ought', in

which it is modelled as a special kind of necessity modal. These semantics are then used to argue for a nondescriptivist inferentialist metasemantic explanation of that in virtue of which the semantics for 'ought' are true. (Thus, Chrisman sides with Brandom—more than once, as we shall see.) The general approach of this thesis has been to see what does and does not follow *given* nondescriptivism. We have not generally been concerned with arguments *for* nondescriptivism. Nonetheless, it will be helpful to outline Chrisman's arguments as it will help to show the nature of the relation between semantics and metasemantics, which is the central theme of this chapter.

3.2.1 Semantics for 'Ought'

Whereas we have been more concerned with moral language generally, Chrisman is primarily concerned with the term 'ought'. While many of the contexts in which we make ought-claims are moral, there are a number of other contexts in which we use the term. For example, we make prudential claims about what we ought to do to promote our welfare; we make teleological claims about what we ought to do given some end; we make evaluative claims about what states of affairs we think ought to obtain; and we make epistemic claims about what ought to be the case given what we know or believe (2015: 32-3). The same ought-sentence might be used in a number of these contexts, and no doubt there are other contexts as well. The point is simply that 'ought' can be used in various different sorts of context.

Chrisman thinks that this places fairly strong constraints on any semantics for 'ought'. Initially, one might think that the term is many-ways ambiguous. Like the word 'bank', we have a single sign that we use to express different expressions with distinct meanings.[9] Unlike 'bank', however, 'ought' is not given distinct dictionary entries depending on which kind of 'ought' is meant. Moreover, the different kinds of 'ought' are not translated into distinct words in foreign languages, as with 'bank'. Rather, like other modal auxiliary verbs, such as 'may' and 'must', "there appears to be a deep intra- and interlinguistic systematicity to the semantic contribution of 'ought' to the sentences in which it figures." (Chrisman 2015: 42)

If 'ought' were ambiguous, then these linguistic facts would be a huge unexplained coincidence. Thus, any plausible semantics for 'ought' must be sensitive to the wide variety of ought-claims, not just moral ones. Chrisman argues that such constraints give us good reason to reject analyses of ought in terms of other normative notions, such as obligations, reasons, and values.[10] Rather, 'ought' should be understood as a primitive or

[9] c.f. Harman (1973: 235).

[10] On obligations (2015: 36-7); on reasons (2015: 43-51); on values (2015: 52-8). A central theme is that no analysis successfully captures every "flavor" of 'ought'. It should probably be noted that Chrisman does not take these arguments to be conclusive. Rather, they are meant to show that there is *prima facie* good reason to treat 'ought' as a primitive normative term.

fundamental normative notion. So we need a uniform semantics that does not reduce 'ought' to any other normative notion.

We are after an account of 'ought' that provides a perspicuous representation of how its meaning of systematically contributes to the meaning of the sentences in which it figures. Chrisman thinks this is best done by representing the meaning of the term as a function to truth-values (relative to a context and world of evaluation). However, it should be kept in mind that at this stage of the inquiry the semantic notions are used only as tools to represent the compositional properties of expressions. The framework is assumed to be theoretically neutral as to whether 'ought' is descriptive or nondescriptive. We can therefore associate ought-sentences with propositions, which determine the truth-conditions for those sentences, without (yet) making any commitments as to the fundamental nature of such language.

As we have already observed, 'ought' is a modal auxiliary verb, similar to 'can' and 'must'. We might therefore expect it to function semantically in the same way as alethic modal operators. Looking first to modal logic, the semantics for necessity operator '□' and possibility operator '◊' can be given respectively in term of universal and existential quantification over a set of possible worlds. These expressions operate as functions from a 'prejacent' proposition to a truth value, relativised to a world of evaluation and accessibility relation. Truth *simpliciter* is truth at the actual world, and different accessibility relations can be defined to model different kinds of alethic modality. For example, nomological necessity can be represented as universal quantification over all possible worlds that are consistent with the laws of nature that hold at the world of evaluation.

At a first pass, we can treat 'ought' as a necessity modal by defining the appropriate accessibility relation. Generally, the semantic value of 'ought' could be given as follows:

(1) $[\![\text{ought}]\!]^{R;w} = \forall p : [\![\text{ought}(p)]\!]^{w} = 1$ iff $[\![\text{p}]\!]^{w} = 1$ in all worlds R-accessible from w. (2015: 72)

We then define the relevant accessibility relation for each kind of ought-claim. For example, this might be defined in the moral case as follows:

R_{moral} $=_{\text{def}}$ the relation that holds between two worlds w and w' iff what is morally ideal in w is true in w'. (2015: 69)

This gives us the following semantic value for the moral 'ought':

(2) $[\![\text{ought}]\!]^{R\text{-moral};w} = \forall p : [\![\text{ought}(p)]\!]^{w} = 1$ iff $[\![p]\!]^{w} = 1$ in all worlds R_{moral}-accessible from w. (2015: 70)

In other words, 'ought' is a unary function that maps propositions to the semantic value *true* just in case that proposition is true at all morally ideal worlds at w. We can now derive the truth-conditions for ought-claims as follows:

(3) 'ought(p)' is true in w just in case p is true in all worlds R_{moral}-accessible from w.

So, for example, consider the sentence:

(4) 'The Levite ought to have helped the wounded traveller.'

This sentence is true just in case the Levite (or his counterpart) helped the wounded traveller in all morally ideal worlds—i.e., all worlds in which what ought to be the case (in the actual world) is the case.

Plausibly, the Levite ought to have had mercy on the traveller, and so (4) is intuitively true. However, under plausible assumptions, (2) predicts that (4) is false. In a morally ideal world, the traveller would not have been robbed and wounded in the first place. So it is not the case that in any morally ideal world, the Levite ought to have helped the traveller—there would be no need. These are examples of non-ideal contexts. Another kind of problem case for (2) are moral dilemmas. If two moral ideals issue incompatible demands on an agent, the set of morally ideal worlds might simply be empty, trivialising all moral ought sentences. While it might be questioned whether moral dilemmas actually exist, this is a substantive moral question that should not be ruled out by one's semantics (Chrisman 2015: 79-80).

Drawing on Kratzer (1981, 1991), Chrisman suggests that we replace R_{moral} with a contextually determined modal base and ordering source to which 'ought' is relativised (2015: 84-6). Very roughly, the modal base is the set of possible worlds that are consistent with a contextually determined set of background conditions. So, for example, in the good Samaritan case, the relevant background conditions include the fact that the traveller was robbed and wounded. Secondly, rather than selecting the ideal worlds in the modal base, we instead have a contextually determined ordering source that sorts the worlds in the modal base into a partial ranking. In the moral case, the ordering source is determined by something like the set of propositions that state what is considered best by moral norms. As the ordering is partial, this allows for worlds to be ranked equally.

Where c denotes a context, $f(w)$ denotes the modal base, and $<_{g(w)}$ denotes the ordering source, we can now give the semantic value of 'ought' as follows:

(5) $[\![\text{ought}]\!]^{f,g;w} = \forall p : [\![\text{ought}(p)]\!]^{c;w} = 1$ iff $[\![p]\!]^{c;w} = 1$ in all worlds $v \in \cap f(w)$, for which there is no $v' \in f(w)$ such that $v' <_{g(w)} v$. (2015: 86)

In other words, 'ought' is a function from a proposition that gives the semantic value true just in case the proposition is true in all the worlds in the modal base for which there are no higher ranked worlds by the ordering source.[11]

As well as dealing with non-ideal contexts, (5) also solves the problem of moral dilemmas. Two worlds that are morally ideal according to conflicting norms may be jointly ranked highest by the ordering source, where the semantic value of 'ought' makes no reference to morally ideal worlds. However, Chrisman argues that (5) is inadequate because it fails to capture the agentive nature of certain kinds of ought-claim, such as in moral contexts (2015: 108-12). There is an intuitive distinction between 'ought-to-be' and 'ought-to-do', in the sense that while some ought-claims are used to say what ought to *be* true, many are used to say what an agent should *do*.[12] As the prejacent in (5) is a proposition, one might worry that this does not adequately capture the active agency of 'ought-to-do' claims.

Sceptical of attempts to capture such agency within the prejacent proposition or by adding an 'agent parameter' to index ought-claims to, Chrisman suggests 'ought-to-be' takes 'practitions' rather than propositions for prejacents (2015: 136-9). Originally developed by Castañeda (1975), practitions are nonpropositional imperative-like contents with a structure and satisfaction conditions analogous to standard subject-predicate propositions. For example, the sentence:

(6) 'Ludwig, put it down!'

expresses the structured practition in which 'put it down' is 'predicated' of Ludwig. (6) is correct or incorrect relative to a norm in the form 'if in *C*, do φ, don't ψ...'. For example, this might be the norm expressed by the sentence, 'If listening to a visiting lecturer, do not threaten them with a poker...'.

While Chrisman thinks that the 'challenge of ought-to-do' poses a substantial problem for (5), I think he overstates the case. The challenge is largely motivated from the idea that there is intuitively an important conceptual distinction to be made between the two kinds of ought-claim. However, given that the point of the compositional semantics is simply to model the compositional properties of expressions, conceptual considerations are not obviously relevant. This only becomes a real problem for (5) if it fails to correctly predict the correct truth-conditions. In any case, we need not dwell on such questions here. The key point for Chrisman regarding nondescriptivism is that 'ought' is best understood semantically as a modal operator. For simplicity, therefore, I

[11] Things are actually more complicated due to the contrast between 'ought' and 'must', the latter being intuitively stronger. Chrisman suggests that the ordering source might have two tiers of propositions, those which are 'necessitated' relative to an ideal (must) and those 'expected' (ought). I ignore this complication here. See Chrisman (2015: 90-6) for discussion.

[12] c.f. Schroeder's account of propositions as things we can do (§2.2).

will assume for the remainder of this chapter that (5) provides the correct semantics for 'ought',

3.2.2 Metasemantics for 'Ought'

Chrisman believes that the task of metasemantic is to provide general grounding explanations of the results of semantics in terms of more fundamental nonsemantic properties and facts (2015: 15). If we take (5) to specify the meaning of 'ought', then what we are after is an explanation of that in virtue of which (5) is true of 'ought'. I will henceforth use 'metasemantics' to refer exclusively to this sense of the term.

The relevant grounding base for semantics consists of the psychological, sociological, and ontological facts to do with the actual and possible linguistically meaningful *use* of phonemes and graphemes (2015: 14). Very generally, we can distinguish between three types of use that might be thought to ground meaning:

> Representationalists treat the use of language to talk about reality as fundamental. So in the final analysis, they see something like asymmetric dependency or functional relations between language and extralinguistic reality as generating meanings. Ideationalists, by contrast, take the use of language to express our minds or thoughts as fundamental. So in the final analysis, they see the expression of ideas as conferring meanings to arbitrary sounds and scribbles, thereby making them part of a language. Finally, inferentialists take the semantically fundamental use of language to be making moves in an inferentially structured practice; for example, to commit to some claim that can provide reasons for other commitments and can itself stand in need of inferential legitimization. So in the final analysis, they see inferentially articulated commitments as the fundamental level of explanation of how and why our language has the semantic structure it has. (Chrisman 2015: 16)[13]

These views need not be mutually exclusive. Moreover, they can be understood either as global claims about all language, or local claims about some language fragment. As Chrisman sees it, metaethical descriptivism is broadly aligned with representationalism, and expressivism with ideationalism. As Chrisman argues for an inferentialist version of nondescriptivism, I will for the most part ignore expressivism and ideationalism in this chapter.

On this view, descriptivism about 'ought' (roughly) comes to the claim that the truth of (5) is grounded in how we use 'ought' to represent ways that reality could be (2015: 161). While this notion is left somewhat vague (presumably to allow for different ways of spelling out the view), I take it as essential to Chrisman that this is an *ontologically committing* notion of representation. By ontologically committing, I do not mean that representationalism is itself an ontological thesis. Rather, representational uses of language function to describe some way that ontological reality is or might be. It is then

[13] Chrisman does not consider the view that the semantic facts might be fundamental and have no grounds—c.f. Boghossian (1989). I ignore this complication here, though see §3.5.2 and Chapter 1: n.12.

another question whether reality is as the language represents it as being. So by using ethical language representationally, we describe moral properties, relations and states of affairs, whether or not these in fact exist. I will henceforth use the qualifier 'robust' to mean ontologically committing in this sense.

Intuitively, this draws the boundary in roughly the right place. Recall that one of the general motivations behind nondescriptivism encountered in the first chapter was to avoid any commitment to a metaphysics of morality without thereby attributing some kind of systematic error to all of our moral judgments. A natural thought would then be to define nondescriptivism as the negation of representationalism. As we will see in the next section, there are problems with this suggestion. However, for the moment let's assume that this suggestion is on the right lines.

Not only does Chrisman think that there are general reasons to be sceptical of attempts to incorporate moral properties and states of affairs into our overall ontology. He also argues that his semantics for 'ought' provide a particular difficulty for representationalism. If (5) is true of 'ought' in virtue of how we use the term to represent reality as being, what ontological commitments do we incur by making ought-claims or judgments?

> According to the [representationalist] view, the truth-conditions of ought-sentences should be interpreted as representing a complex way reality could be, which includes things like the truth of a proposition relative to possible worlds and the legitimacy of practitions relative to possible norms. This involves ontological commitment to the existence of such possible worlds and possible norms; and obviously that commitment coheres with the realist view in metanormative theory that things like values, obligations, and norms are part of the fundamental fabric of reality. (2015: 167)

Chrisman argues this position is an "ontologically prolifigate position, in the sense that it is an affront to common sense about the difference between what is real and what is imaginary/virtual/fictional." (2015: 170) Chrisman takes this to follow from a commitment to the existence of (nonactual) possible worlds. One is reminded here of the 'incredulous stare' objection to modal realism. The problem, however, is that a commitment to the existence of possible worlds does not imply a commitment to modal realism. For example, possible worlds might be identified with sets of propositions or sentences. Existence to these kinds of entities are hardly ontologically proliferate. Moreover, Chrisman seems to suggest that a commitment to the existence of possible worlds implies a commitment to the existence of the contents of those possible worlds, such as "dragons"—hence (presumably) the "affront to common sense". However, this simply does not follow on ersatz views of possible worlds.[14]

[14] It should be noted that Chrisman does acknowledge this objection, though he offers no reply to it (2015: 169n). Perhaps one might argue that the ontological conservativeness of nondescriptivism at

Chrisman also argues that nondescriptivism is supported by the semantic function of 'ought'. He notes how in truth-conditional semantics, 'ought', like the logical operators more generally, gets a different treatment than other kinds of terms, such as predicates and indexicals. One might argue that 'ought' therefore does something fundamentally different in language than representing (2015: 169). However, this seems like an odd argument for Chrisman to make. The whole point of distinguishing semantics from metasemantics was to isolate the inquiry into the logico-semantic properties of expressions from broader philosophical debates about what such language fundamentally does or is fundamentally about. However, Chrisman seems to be doing here exactly what he says we should not do, *viz.*, inferring metasemantic conclusions directly from semantic premises. However, perhaps only the weaker point is being made that it is less obvious that sentential operators are used to represent compared with predicates.

There is also a question as to what we are to say about other moral terms, which are not obviously parsed as modals. Chrisman does make some suggestions as to how one might analyse other notions in terms of ought and then take ought as fundamental (2015: 209-14). It would take us too far afield to examine these here. I will note, however, the burden that this seems to place on arguing for moral nondescriptivism generally. Chrisman sometimes remarks that, in the worst case, nondescriptivism could still be argued for 'ought', even if not for other moral terms. However, this also relies on his earlier arguments against analysing ought in terms of other normative notions. So there is a worry that the final position might be incredibly weak.

These problems notwithstanding, one might argue that Chrisman's arguments have bite insofar as there is an attractive nonrepresentational alternative, which of course he thinks there is. So let us turn to that now. Metasemantic inferentialism is the idea that the contents of sentences are grounded in our inferential practices. Thus stated, it implies neither descriptivism nor nondescriptivism. It all depends on the nature of the inferential practices in question. Theses practices might be articulated in terms of an inferential commitment to (robustly) representing reality as being a certain way. So not just any inferential practice will do for the nondescriptivist.

Fortunately, a nonrepresentational theory of inferential practice is ready to hand in Robert Brandom's (1994, 2000, 2008) Sellarsian socio-normative theory of discursive practice. Chrisman presents a variant on this account in which to situate a nondescriptivist metasemantics for 'ought'. While Brandom's inferentialism is often presented as an alternative to truth-conditional semantics (meaning *as* use), Chrisman

---

least is a *pro tanto* reason for favouring it. Regardless, this has nothing specific to do with 'ought' being a modal operator.

argues that we are better to construe it as a metasemantic account of the kind of linguistic practice that grounds the contents of linguistic expressions.[15,16]

The basic idea is that meaningful sentences have the content they do in virtue of how we use those sentences in a social-normative discursive practice ('the game of giving and asking for reasons'). By using a sentence to make an assertion, one makes a commitment that can both serve as a reason or justification for other commitments by way of a premise, or itself stand in need of justification by way of a conclusion:

> the sense of endorsement that determines the force of assertional speech acts involves, at a minimum, a kind of *commitment* the speaker's *entitlement* to which is always potentially at issue. The assertible contents expressed by declarative sentences whose utterance can have this sort of force must accordingly be inferentially articulated along both normative dimensions. Downstream, they must have inferential *consequences*, commitment to which is entailed by commitment to the original content. Upstream, they must have inferential *antecedents*, relations to the contents that can serve as premises from which entitlement to the original content can be inherited. (Brandom 2000: 193-4)

The relevant kind of inferential relations are understood primarily as semantic or 'material' implications, rather than logical entailments. So, for example, the propriety of the inference from 'Pittsburgh is to the west of Princeton' to 'Princeton is to the east of Pittsburgh' fundamentally follows from the content of the claims, particularly the meaning of 'west' and 'east', rather than from a logical deduction. The suggestion is that all meaning is determined by these sorts of inferences.

What kind of inferentially articulated commitment does one acknowledge when one makes an 'ought' claim? Brandom argues that both modal and normative vocabulary should be understood as broadly logical vocabulary. What kind of commitment does one acknowledge when one makes a claim using logical terms? Brandom argues that our use of logical vocabulary is fundamentally expressive rather than descriptive. However, this is a different sense of 'expressive' than that of the kind discussed in previous chapters. For Brandom, logic is expressive in the sense that it *makes explicit* or *codifies* certain aspects of the inferential structure of our discursive practice. The paradigm example is given by the conditional:

> Prior to the introduction of such a conditional locution, one could *do* something, one could treat a judgment as having a certain content (implicitly attribute that content to it) by

[15] Brandom often flounders on whether his inferentialism is incompatible with truth-conditional semantics—see Price (2011a) for discussion. Price also argues that we should construe Brandom's theory as a 'sideways-on' explanation of meaning, rather than an account of meaning itself.

[16] Brandom (1994) distinguishes 'normative pragmatics', which seems to provide something similar to a metasemantic explanation in Chrisman's sense, from 'inferential semantics', which seems to be about the nature content itself and not just its determination. Chrisman helps himself to both aspects in his metasemantics, but it is not obvious he is entitled to the latter part of Brandom's account. The next section will address these worries more generally.

> endorsing various inferences involving it and rejecting others. After conditional locutions have been introduced, one can *say*, as part of the content of a claim (something that can serve as a premise and conclusion in inference), *that* a certain inference is acceptable. One is able to make explicit material inferential relations between an antecedent or premise and a consequent or conclusion. Since, according to the inferentialist view of conceptual contents, it is these implicitly recognized material inferential relations that conceptual contents consist in, the conditional permits such contents to be explicitly expressed. (Brandom 2000: 60)

By making explicit the otherwise implicit inferential relations that determine propositional content, the inferences can themselves become the subject of assertion and judgment, thereby standing in potential need of justification.

Importantly for our purposes, this is a thoroughly nondescriptive account of how logical content is determined. No representational notions come into the account at all. Logical content is explained in terms other than describing some aspect of reality (the logical part). So we need not appeal to logical properties or facts in the fundamental explanation of logical content. Rather, logical vocabulary is fundamentally used to articulate the conceptual or inferential framework in which all other meaningful language use occurs.

In an analogous way, Chrisman urges that we should understand the fundamental content-determining role of 'ought' as codifying or making explicit 'second-order' inferential relations between more basic items of content. However, in the moral case, the kind of inferential relations made explicit are those to do with *practical* reasoning. As Brandom puts it, "normative vocabulary (including expressions of preference) makes explicit the endorsement (attributed or acknowledged) of *material* proprieties of *practical* reasoning. Normative vocabulary plays the same expressive role on the *practical* side that *conditionals* do on the *theoretical* side." (2000: 89) So, for example, perhaps the claim:

(7) We ought to give more to charity,

makes explicit material inferences such as:

(8) There is suffering in the world which we can alleviate by giving to charity at no unreasonable cost to us → We shall give to charity.

However, as (7) is itself a propositional claim, it not only makes explicit inferential relations such as (8), but it can also serve as a commitment whose entitlement can be called into question.[17] So, for example, one might justify (7) with:

(9) We ought to alleviate suffering in the world if it is at no unreasonable cost to us.

[17] Thus ought-claims are not "*mere* "inference tickets" in a "game" of giving and asking for reasons" (Chrisman 2015: 194), as they themselves express a propositional content.

To conclude, we have the beginnings sketched of a metasemantic inferentialism for 'ought'. As the linguistic practice that grounds the semantics for 'ought' does not involve any robustly representational notions, the account is nondescriptivist. Moreover, as 'ought' was characterised as a "metaconceptual device" for "manipulating more basic items of content", this gives us a natural place to draw the descriptive/nondescriptive boundary in language use. For the more basic items of content can be understood as descriptive contents, which keep track of things in our environment, rather than directly regulating our discursive practices. One might cash out descriptive commitments in terms of the inferential commitment to reality being a certain way (2015: 191-2), or one could be a local inferentialist about moral vocabulary and a local representationalist about descriptive vocabulary (2015: 170-1). Either way, we have a clear contrast between language use that is ontologically committing (descriptive) and language use that is not ontologically committing (nondescriptive).

3.3 Nondescriptivism Lost: A Problem With Grounding

Although it has not been the topic of discussion in this chapter, inferentialism looks in certain ways more promising than expressivism as a form of nondescriptivism. Seeing as the content of a sentence is determined by its inferential role rather than by the mental state expressed by canonical uses of the sentence, it is not obvious that the Frege-Geach problem arises for inferentialism. In the same way that expressivism explains motivational internalism by including directive mental states in its fundamental characterisation, inferentialism explains moral reasoning by specifying the inferential role of moral terms in its fundamental characterisation.

There are other respects, however, in which one might think that inferentialism comes at a cost. One reason has to do with the holistic nature of content or content determination under inferentialism. Chrisman only suggests a few illustrative examples of the kind of inferences that ground the content of expressions. While this may enable us to get the general picture, one might worry that the account is left rather vague and imprecise. The question is to determine the scope of implication relations that determine the meaning of 'ought'. There are innumerable inferences that one might make using such term. Surely not all of these are of equal importance in determining the meaning of the term. This is in fact a problem for inferentialism generally. Chrisman acknowledges the general worry, but provides no answer (2015: 188-9).

Not only is there a question of scope, however, but there also seems to be a lack of systematicity to the kind of inferences that 'ought' codifies. Metasemantics is ultimately about our linguistic behaviour, and so we aren't in the business of providing necessary and sufficient constitutive conditions for ought-claims. However, we might still wish to provide necessary and sufficient conditions, perhaps constitutive, for making ought-claims, perhaps in terms of the rules that it is necessary and sufficient to follow in order

to make ought-assessments.[18] However, if inferentialism is correct, then there's a real question of how much we can expect in this way. If content is radically holistic in the way that Brandom seems to think, then there might simply be no systematic and precise account of the inferences that determine the content of 'ought', or any other expression for that matter. Perhaps we can only give the general metasemantic function together with canonical examples. This is in contrast to expressivism which takes a more atomic approach to explaining the meaning of expressions.

It is probably worth noting also that it is not without controversy that a Brandomian inferentialism has the resources to deal with subsentential compositionality. The basic item of content for Brandom must be a proposition, as asserting a sentence is the simplest move that one can make in the game of giving and asking for reasons. Nothing can be *said* except with a proposition. Very roughly, Brandom accounts for the compositionality of subsentential expressions in terms of the inferential proprieties of the substitution of such expressions across a certain class of sentences. It is not the place to get into such matters here, but it is interesting that Chrisman does not so much as mention the issue, considering his emphasis on compositionality generally.[19] Perhaps he thinks that these issues simply don't arise at the level of metasemantics. Or perhaps as 'ought' is a sentential operator, predicates can be rescued by a local representationalism. In any case, something needs to be said here.

So there are a number of issues that need clearing up if one were to follow Chrisman's line. However, I think that these are more invitations for clarification and development rather than knock-down objections. In any case, there is a more fundamental problem with Chrisman's account: metasemantic nondescriptivism does not imply metaethical nondescriptivism.

Define descriptivism for language *L* the claim that *L* is robustly representational in the ontologically committing sense. Define nondescriptivism about *L* as the negation of descriptivism about *L*. Suppose Chrisman's semantics and metasemantics for 'ought' are both true. So the metasemantics for 'ought' are nondescriptive in the defined sense. This is perfectly consistent, however, with supposing that the *semantics* for 'ought' are descriptive in the defined sense. It was claimed that (5) itself does not imply descriptivism. But it does not it rule out either. So while the metasemantics for 'ought' might be nondescriptive, the semantics for ought might still be descriptive, because the *content itself* is robustly representational. However, if the content of 'ought' is robustly representational, then this fails to meet the criteria of adequacy for nondescriptivism. So Chrisman's metasemantic theory does not constitute metaethical nondescriptivism.

In fact, *any* metasemantics will fail to secure metaethical nondescriptivism in the absence of auxiliary claims. It has nothing much to do with inferentialism, but rather with how the traditional metaethical debate has been relocated to the level of

[18] This kind of approach is carried out, for example, by Zalabardo (ms) for alethic assessment.

[19] Fodor and Lepore (2001a, 2007) particularly press this point against Brandom, who replies in (2010b).

metasemantics. Once we observe that a nondescriptive metasemantics is compatible with a robustly descriptive semantics, however, it begins to look less like a relocation and more like a change of topic. Perhaps one could incorporate one's metasemantic theory into one's overall metaethical theory. Moreover, perhaps it might play an important role in explaining how moral discourse is nondescriptive, if indeed it is. But it cannot explain this by itself. We need an additional account of how the content itself is nondescriptive.

To make the point a little more concrete, consider the conceptual role semantics for 'ought' set out by Wedgwood (2007). Like Chrisman, Wedgwood claims that "the essential conceptual role of normative concepts consists of a certain *regulative* role that these concepts play in *reasoning*—including *practical* reasoning. In that sense, normative judgments about what actions one ought to perform are essentially connected to motivation and practical reasoning because they involve a concept whose essential conceptual role is its role in practical reasoning." (2007: 80) Also like Chrisman, Wedgwood takes 'ought' to be a propositional operator, though now indexed to agent and time parameters. The conceptual role of ought is then given in terms of how one who accepts an ought-claim is thereby committed to make the prejacent proposition part of one's ideal plans about what to do at the specified time (2007: 97). This glosses over many important details. However, both accounts maintain that the role of 'ought' in reasoning fundamentally determines its semantic content.

Wedgwood differs from Chrisman in assuming that the conceptual role of 'ought' determines a robustly representational content for 'ought'. This is largely due to semantic assumptions held by Wedgwood that Chrisman would reject. Particularly, Wedgwood assumes that "the nature of a concept consists purely in the contribution that it makes to the nature of the thoughts in which it appears; and such thoughts are nothing more than ways of representing some possible state of affairs." (2007: 81) The semantic value of 'ought' is given as the state of affairs that would make all the rules of inference that constitute the conceptual role of 'ought' valid (2007: 86). Particularly:

> the semantic value of the practical 'ought'-operator '$O_{\langle A,t\rangle}(p)$' will be the weakest property of a proposition p that makes it the case that it is correct for A to make the proposition p part of her ideal plan about what to do at t, and incorrect for A to make the negation of p a part of her ideal plan. (2007: 100)

Of course, Chrisman would reject Wedgwood's assumption that one needs to posit any such state of affairs to play the role of the semantic value of 'ought' given Wedgwood's compositional semantics. However, this is beside the point. For the present objection to Chrisman's account to go through, it is sufficient that Wedgwood's account is a possibility.

Importantly for our purposes, the functional role that determines the representational content of 'ought' for Wedgwood can be given without any mention of

representational notions. All that is required is acceptance, commitment, and ideal planning.[20] We might say that Wedgwood provides a nondescriptive metasemantics for 'ought', and that this only determines a robustly descriptive semantics for 'ought' given other substantive assumptions. Perhaps Wedgwood would not agree to putting it in this way, but his account could easily be adapted. Regardless, we have here a counterexample to the idea that a nondescriptive metasemantics for moral language implies metaethical nondescriptivism.[21]

The trouble stems from a confusion on Chrisman's part regarding his 'semantic conservativism'.[22] Semantic conservativism is the idea already encountered that one should be able to engage in compositional semantics using model-theoretic tools (or whatever else) without getting embroiled in metaphysical debate.[23] Here, the use of the particular semantic notions is justified insofar as they perspicuously model the semantic properties under consideration. However, once we have decided on our semantics in this metaphysically neutral way, we *do* need to get embroiled in metaphysical debate to settle the descriptivism/nondescriptivism question. This is a question about the *nature* of moral language. In addition to *grounding* explanations of one's semantics, one can ask for *constitutive* explanations of one's semantics. By engaging in a metaphysically neutral semantics and then going on to provide metasemantic grounding explanations, we skip pass constitutive explanations of what it is for an expression to have the meaning that it does.

Notwithstanding express statements that metasemantics is in the business of providing grounding explanations, Chrisman often shifts to the language of constitutive explanation in his discussion of metasemantics. For example, there is talk of "truth-conditions *as* ways reality can be", and "truth-conditions *as* positions in a space of implications" (2015: 161, 184 emphasis added). Moreover, in motivating his semantic conservativism, he makes the following remark: "I am skeptical that anyone who rejects descriptivism must reject the standard truth-conditionalist approach to compositional semantics... because I think one can take a less committal view about what *constitutes* truth in "truth-conditions"" (Chrisman 2015: 9, emphasis added). So while the official line is that Chrisman's metaethical nondescriptivism is a metasemantic view, it is unclear where he stands on the constitutive question.

What is clear, I submit, is that grounding alone cannot secure metaethical nondescriptivism. A full blown metaethical theory of language cannot remain silent on

[20] While 'acceptance' for Wedgwood is belief in a robustly representational sense, this is a consequence of 'ought' being robustly representational rather than being an essential feature of the conceptual role as such.

[21] Perhaps other aspects of Wedgwood's view would rule him out as a counterexample, such as his Platonism about normativity. This example is only meant to be illustrative, however, and so the main point would remain unaffected. A quite different sort of counterexample might come from the kind of Aristotelian realism associated with Wiggins and McDowell.

[22] The phrase is from Bar-On, Chrisman, and Sias (2014).

[23] §3.1.

the nature of the content of 'ought', or what it is for 'ought' to have the content that it does. Supplementary theory is needed. However, remaining silent is not the same as claiming that there is not much to say. In the final section of this chapter, I suggest how an inferentialist metasemantics might be combined with a deflationary theory of content to achieve metaethical nondescriptivism. His use of constitutive language notwithstanding, one might think that this is implicit in his account anyway, and elsewhere he comes close to endorsing some kind of semantic deflationism.[24] After motivating the position, I discuss some of the tensions that might arise from combining deflationism and nondescriptivism.

3.5 Nondescriptivism Regained: Semantic Deflationism

3.5.1 A Way Ahead?

Nondescriptivist grounds for some semantics do not imply constitutive nondescriptivism about those semantics. To save Chrisman's account, we therefore need a suitably nondescriptive answer to the constitutive question. There are two ways to go here. On the one hand, one might provide an inferentialist account of what content consist in—that truth conditions *just are* positions within a space of implications. Alternatively, one might provide a deflationist or minimalist answer to the constitutive question—that there is no deeper answer to what truth conditions consist in, generally speaking. As noted above, I don't think it's entirely clear from Chrisman's own writing which, if any, he would opt for. However, it seems to me that a deflationary approach is better suited to provide for his wider theoretical aims. Let me briefly explain why.

As well as arguing for his particular version of nondescriptivism, Chrisman wants to convince us more generally that debates between descriptivists and nondescriptivists should be conducted at the level of metasemantics and not semantics. By focusing on compositionality and truth-conditions, we obscure the debate, as the standard resources to deal with these features of language form a common treasury for all. If descriptivism or nondescriptivism are grounding explanations, then the troublesome semantic constraints do not directly apply, as these only apply to semantic theory. However, if either view were formulated instead as a constitutive explanation, then the semantic constraints would directly apply. Whereas grounding explanations invoke some class of facts in the explanans distinct from the semantic facts, constitutive explanations tell us about the *identity* conditions of the explanandum. As such, the explanandum and the explanans pick out the same class of facts. *Ipso facto*, any constraints that apply to one's

[24] Chrisman (2014). Again, however, it's not fully clear whether he endorses the approach himself.

compositional semantics apply directly to one's account of what one's semantics consist in.[25]

Chrisman argues by interpreting expressivism in terms of an ideationalist metasemantics, expressivists can embrace truth-conditional semantics and thereby "deploy the standard truth-conditionalist explanations of semantic phenomena that critics accuse them of being unable to explain." (2015: 172) If this kind of ideationalism is a grounding claim, then this leaves room for a constitutive descriptivist construal of moral sentences, and so fails to secure nondescriptivism. If we understand ideationalism as a constitutive explanation, however, then one's truth-conditional semantics just is a psychologistic-expressivist semantics. And so all the problems that arise for semantic versions of expressivism will also arise for one's truth-conditional semantics, if what the truth-conditions consist in are just correctness conditions of the thought one should have when uttering certain sentences.[26]

This need not be taken as an argument against any such view. For example, one could first provide a truth-conditional model for moral expressions and then go on to interpret the elements of the model in psychologistic terms. This would allow one to see exactly the kind of structure required of the mental states that we use such expressions to express. Something like this is argued by Alex Silk (2015), who develops a possible worlds semantics for deontic normative claims, where normative sentences are evaluated relative to a set of ranked possible worlds. Logical relations in this model are then argued to be isomorphic to coherence constraints on the kind of preference attitudes appealed to in decision theory. So there is scope for nondescriptivism to make use of truth-conditional semantics without the move to metasemantics.

In essence, however, such an approach does not avoid the hard toil of providing an nondescriptivist alternative to truth-conditional semantics. At best, it has the methodological benefit of allowing us to get clear on the exact structure of whatever answers the constitutive question. This is at odds with Chrisman's general picture for two reasons. First, Chrisman clearly has something stronger than this in mind. He clearly thinks, for example, that the Frege-Geach problem (in its semantic guises) *isn't really* a problem for expressivism, which it would be on this model. Second, if one did manage to successfully provide (say) an inferentialist constitutive explanation, then it is unclear what explanatory work the grounding explanation is left to do. Perhaps the grounding explanation might help to explain the constitutive explanation in some way. However, if we asked in virtue of what our account of moral language was

[25] One might deny that compositionality does in fact impose any serious constraints—see Horwich (1998a: ch.7), and Fodor and Lepore (2001b) and Horwich (2005: ch.8) for discussion.

[26] Chrisman does acknowledge that the Frege-Geach problem might remain a problem for a metasemantic expressivism due to difficulties specifying the kind of mental state expressed by mixed sentences (2015: 177-80). The argument being presented here, however, is much stronger: *any* problem that arises for an expressivist semantics will arise for a truth-conditional semantics with an expressivist-constitutive explanation.

nondescriptivist, it would seem sufficient simply to cite the nondescriptive constitutive explanation.[27] I think these reasons undermine Chrisman's overarching aims, and so I suggest that he should reject this approach, and instead opt for semantic deflationism.

What is semantic deflationism? Recall that we can roughly distinguish a demand to specify *what* the meaning of some expression *e* is from a demand to explain what *e*'s having this meaning *consists in*. According to the deflationist, there simply is no answer to the second question. Or rather, there is no answer to the second question over and above the first. In other words, it is sufficient to specify what the meaning of an expression consists in simply by specifying the meaning of the expression. While semantic notions, particularly truth, may be used to *express* the specification of the meaning of an expression, truth itself does not consist in any substantive or robust word-world relation. Rather, there is nothing more to say about truth over and above the trivial platitudes governing our use of the truth-predicate.

There are many ways in which to develop this thought. It is not the place to survey these here. For simplicity, I will focus on the semantic deflationism expounded by Paul Horwich (1998a, 1998b). According to Horwich, our use of the truth-predicate can be explained solely by the fact that we are inclined to accept instances of the following equivalence schema:

(ES) The proposition *that p* is true iff *p*.

Moreover, our acceptance of (ES) is both necessary and sufficient to explain the reason for having a truth-predicate in the first place—viz., the role that it plays in generalisations, whether in logic, ordinary language, philosophy, or semantics (Horwich 1998a: 106). These generalisations are achieved by virtue of the logical relations that they stand in to (ES). Consider the claim made in ordinary language:

(10) What Oscar said was true.

Imagine that we think that Oscar has great culinary taste, and what Oscar asserted was that eels are good (1998b: 3). However, we didn't quite catch his remark. If we wish to infer what it is that Oscar said, then we need a proposition equivalent to:

(11) If what Oscar said is *that eels are good* then eels are good, and if he said that *milk is white* then milk is white, ... and so on.

And it is exactly (10) that supplies this. Consider next the case from logic:

[27] Assuming that one wouldn't hold a nondescriptive constitutive explanation together with a descriptive grounding explanation. That would just be weird.

(12) Every proposition of the form *everything is F or not F* is true.

Without the truth-predicate, we would have great difficulty expressing (13), other than by asserting an infinite conjunction of every proposition of that form (1998b: 4). This expressive role exhausts the reason for having a truth-predicate in a language.[28] Other basic semantic notions can be explained in a similar fashion.[29] The idea is that by 'deflating' semantics in this way, the truths of semantics do not commit us to some robustly real domain of facts that such facts must be identified with.[30]

Is semantic deflationism open to the objection that a nondescriptive account of the *use* that determines some content does not rule out a robustly descriptive constitutive account of that content? This is a delicate issue which requires more attention than I can give it here. The objection is right to say that there is nothing in the deflationist's account of our use of, say, the truth-predicate, that strictly rules out the possibility that there might be some substantive property of truth. After all, the theory says nothing positive about *truth* at all (other than trivial platitudes such as (ES)). It is not better to call deflationism a theory of the truth-predicate or of truth ascriptions?[31]

Well, yes and no. Deflationism aims to undermine any motivation for inflating truth in the first place. Why, it asks, do we need to seek some metaphysically substantive notion if our use of the notion can be explained perfectly well without it? Note that this move is in essence the same as the nondescriptivist argument for metaphysical and epistemological solvency.[32] If we explain the concept in question in terms of its nondescriptive practical function, we can sidestep the need for any metaphysical or epistemological account of the domain of the concept in question, even if we do not rule out the possibility of providing one. However, it still might be argued that the perceived *need* to provide one is based on a misunderstanding of the concept in question. So it is still perfectly in order to talk about semantics truths. But we need to take a quietist approach to what these truths consist in. Or so the deflationist argues.

I don't take myself to have shown that this move is sound. Of course, semantic deflationism itself would need to be defended. I suggest this, however, in order to sketch a picture of what a deflationary conception of semantics might look like. If ought sentences express propositions, and propositions are given in terms of truth-conditions, then there is nothing to say about the truth-conditions of ought-claims other than what is stated by our semantics. This is to take a quietist approach to propositional content,

[28] This is not to say that we might use the truth-predicate pragmatically for many other purposes—see Kukla and Winsburg (2015: 27).

[29] For example, see Horwich (1998: ch4, ch5) for 'meaning' and 'reference' respectively.

[30] For arguments that semantic deflationism is compatible with truth-conditional semantics, see, for example, Burgess (2011) and Williams (1999).

[31] Devitt is fond of making this point (2010).

[32] §1.1.

where no constitutive question arises (even if it is not strictly ruled out).[33] So there is no deep answer as to what the meaning of 'ought' consists in. We can still take (5) to provide a representation of how 'ought' compositionally relates to other expressions. Moreover, it might have other uses, such as in translation, for it provides necessary and sufficient conditions for recognising 'ought' in an alien language. However, as truth conditions as such do not have any deep nature, there is no need to understand (5) as having any deep nature either. While (5) may have many explanatory uses, it says nothing deeper than 'ought' means ought.

However, the question naturally arises as to why (5) is true of 'ought' and not some other semantic rule? Why does 'ought', or any term for that matter, mean anything at all? If (5) is an irreducible semantic fact, then how as natural beings are we meant to grasp a rule like (5)? Semantic deflationism provides no answers to these questions. Moreover, there is a worry that the account begins to look mysterious, appealing to an unexplained realm of non-natural semantic facts. If this is left unexplained, then this goes against one of our criteria of adequacy for naturalism.[34] Moreover, principles like (ES) use the notion of a proposition, and so seem to require an antecedent understanding of a meaningful content to which semantic notions can be applied to in the first place.

It is at this point that we can appeal to Chrisman's inferentialism to answer these questions. In Chrisman's case, sentences are meaningful in virtue of how they can be used to acknowledge commitments that can stand in need of justification and serve as a reasons within a discursive social practice of deontic scorekeeping. 'Ought' is therefore a meaningful term because it can be used in inferential practice this way. It has the particular meaning that it does because of the particular inferential role that it plays, namely its being used to make explicit the propriety of particular kinds of material inference (moral practical reasoning).[35]

Thus, by adopting a deflationary attitude to semantic content, we shift the explanatory burden of one's theory of moral language to the level of metasemantics. So could semantic deflationism save Chrisman's account? It certainly doesn't save his more general claim that metaethics generally is better thought of in metasemantics terms. Semantic deflationism is a substantive and controversial commitment, so it should not have to underlie any general debate. However, this does not mean that it is not an option for any particular theory. So we might still hope to save metasemantic nondescriptivism by embracing semantic deflationism. Most contemporary nondescriptivists accept some sort of deflationism about truth anyway, and so it is not a huge step to the more radical sort of deflationism sketched above.

---

[33] Thus the kind of deflationism presented here is not helpfully presented as some kind of anti-realism about content—c.f. Boghossian (1990).

[34] §1.3.

[35] Having a suitable metasemantic would also preclude the introduction of deviant expressions such as Dreier's 'is hiyo'-predicate (1996).

Obviously, however, the approach is theoretically more costly, as a fully worked out theory would need to defend not only metaethical nondescriptivism but also semantic deflationism. However, nondescriptivism just is philosophy of language applied to moral discourse, and so we should expect commitments from each to entitle and stand in need of justification to commitments from the other. Moreover, semantic deflationism, while raising distinctive problems of its own, is attractive for many of the same reasons that nondescriptivism is. These two approaches are to a large degree complementary. In the remainder of this chapter, I raise some problems that arise from holding these views jointly.

3.5.2 Troubles in Paradise

Nondescriptivism often starts with the thought that moral discourse is in some sense importantly different to descriptive discourse, paradigmatically perceptual discourse. One might worry, however, that semantic deflationism threatens to undermine this distinction. For it might be thought that by deflating what are traditionally considered to be 'representational' semantic notions, such as truth and reference, we thereby rob the descriptivist of the resources to formulate her position. This in turn would rob the nondescriptivist of the resources to formulate the distinction between descriptive and nondescriptive discourse that motivated her theory in the first place.

Thus stated, I hope it is clear that the worry is mistaken. It is true that semantic deflationism denies that there is any distinction to be drawn between descriptive and nondescriptive discourse at the level of semantics. However, it does not follow that there distinction cannot be made in other ways—particularly, in one's metasemantics. So we can still use representational concepts in our metasemantic account of descriptive discourse. They will just be nonsemantic. In fact, this has to be the case anyway, regardless of semantic deflationism. For if one provides a representationalist metasemantics, and metasemantics is in the business of explaining semantics in terms of more fundamental *non*semantic facts, then we need some nonsemantic conception of representation anyway. The general idea should be familiar from 'theories of reference' that aim to reduce reference to some naturalistic relation.[36] So perhaps these theories could provide the conceptual resources for formulating descriptivism.

This only goes so far, however, and certain tensions remain. At this level of abstraction, the move might look somewhat ad hoc. For it might seem that we maintained the descriptive/nondescriptive distinction by fiat, stipulating two distinct senses in which a term can 'refer', for example. I think that the tension here is real, and I

[36] Chrisman suggests in passing that representation might be cashed out in terms of causal regulation (2015: 192). See, for example, Field (1972), Millikan (1984), and Fodor (1990) for three quite different conceptions of the kind of causal relation this might be. Chrisman's notion of descriptivism as representing ways that reality might also be thought to lend itself to Lewis' 'reference magnetism' (1984).

do not intend to settle the issue here. It certainly looks ad hoc when presented at this level of generality. However, perhaps it will not when we have a worked out theory of what each sense means. After all, substantive representational notions in philosophy are terms of art, even if their life begins in the everyday. So it would be premature to try and settle the matter here.

Notice that the same issue arises in relation to the corresponding ontological categories. We saw that one way in which Chrisman seems to distinguish between descriptive and nondescriptive uses of language is between those that are ontologically committing (e.g. property denoting) and those that are not. Again, ontological naturalism—the thesis that all that exists are the objects, properties, and states of affairs recognised by the natural sciences—is often cited as a motivation for nondescriptivism. Demarcation problems aside, semantic deflationism about semantic categories invites a corresponding deflationism about ontological categories. For example, insofar as we accept a predicate '*F*', and we accept the claim that '*x* is *F*', then it seems a harmless step to infer that *x* has the property of being *F*. This is an example of what Stephen Schiffer (2003) calls a 'something-from-nothing-transformation': a conceptually valid inference from a statement involving a certain kind of semantic category to a statement that refers to a 'pleonastic' entity corresponding to that category.[37]

Again, perhaps a working distinction could be made.[38] However, it must be principled and explanatory. It is no good simply postulating a bifurcation between those expressions that are "*really*" representational and those that are not. While some writers have urged that we should give up on representation entirely in theorising about language, this seems premature if drawn from very general considerations about semantic deflationism.[39] One might think that Chrisman has a way out here. For recall Chrisman's emphasis on how 'ought' is an intensional operator and not a predicate. It seems possible that such a semantic distinction might reflect the fact our use of 'ought' fundamentally differs from representational terms (paradigmatically predicates) even if the semantic distinction does not consist in this difference. So perhaps the fact that 'ought' is not a predicate provides evidence for its being used nondescriptively.

However, if one were to take this approach, then one would have to eschew talk of semantic properties, if semantics were not in the domain of the real. Again, however, there is some room for manoeuvre here. Consider, for example, Brandom's prosentential theory of truth, which is tied up with his general anaphoric account of semantic vocabulary (1994). Brandom argues the correct analysis of 'is true' reveals that it is not a predicate but a pro-sentence forming operator. On this picture, it is argued that

[37] Thomasson (2014) goes so far as to argue that semantic deflationism is only plausible with ontological deflationism.

[38] Compare Wright's (1992) pluralist reconstruction of the realism/antirealism debate.

[39] See, for example, Macarthur and Price (2011). In later work, Price (2013) develops a more nuanced position in which he distinguishes i-representation, which is semantic representation, from e-representation, which is covariation with an environment.

inflationism about truth is based on a misunderstanding of the correct logical form of 'is true'. However, this approach in many ways goes against the general spirit of deflationism by drawing metaphysical conclusions from semantic premises. Even Chrisman himself uses the sentence 'the joke is hilarious' as a plausible counterexample to the claim that all predicates refer to properties (2015: 10). Perhaps one could argue that the correct logical form of 'is hilarious' is not predicative, but this seems implausible.

Finally, the kind of position being suggested here might look hard to distinguish from certain forms non-natural descriptivism. For it seems conceivable that there might be a descriptivist theory which denies that there is any answer to the grounding *or* constitutive question about semantics, yet nonetheless maintains that the semantics are robustly descriptive.[40] I have already noted this possibility when I claimed that semantic deflationism does not strictly rule out a substantive metaphysics for semantic notions. Rather, it tells a story that aims to undermine the motivations for having one in the first place, and show the theoretical possibilities and payoffs of taking an anti-metaphysical stance to semantics. Moreover, it might be argued that the naturalistic credentials of deflationism make it the preferable approach.

Again, I cannot address the question adequately here. These matters are best settled by examining each particular theory in detail and seeing what commitments it incurs and how it differs from other possible positions. Such matters raise a whole host of issues in the philosophy of language, metaphysics, and philosophical methodology generally. We cannot always expect to taxonomise philosophical theories in terms of yes and no answers to some fundamental question—do moral sentences express propositions? are moral judgments cognitive? are there any moral facts?—though no doubt these serve as helpful, perhaps indispensible, starting points in our inquiry.

In this thesis I have tried to illuminate some of the starting points that one might take to developing a propositional nondescriptivism. However, while the overall aim has been general, the conclusions were reached from an examination of detailed attempts by other thinkers to forge the beginnings of such a position, even if perhaps this was not their primary goal. I hope to have shown that working out even the basic shape of a propositional nondescriptivism is a task far more complex than the initial characterisation of the approach might suggest. Nonetheless, I believe that our inquiry has shown that there are a wide array of conceptual resources available for challenging ingrained assumptions (perhaps dogmas) of the contemporary metaethical landscape. This might not be much, but one might take from this a cautious optimism for further inquiry.

[40] c.f. Boghossian (1989).

3.6 Conclusion

In this chapter, we examined the suggestion that propositional nondescriptivism might be achieved by formulating its central claims at the level of metasemantics, which provides general grounding explanations of semantics. It was argued that such a position does not imply nondescriptivism. We also saw how inferentialism might provide a more plausible nondescriptivist alternative to expressivism. It was suggested that metasemantic inferentialism might be combined with semantic deflationism to achieve propositional nondescriptivism. Some problems were raised both inferentialism and its combination with deflationism that would need to be answered by any such approach.

Bibliography

Alwood, Andrew. 2016. 'Should Expressivism Be a Theory at the Level of Metasemantics?'. *Thought* 5: 13-22.

Ayer, A. J. 1946. *Language, Truth and Logic* (2nd ed). London: Gollancz.

Ayer, A. J. 1984. *More of My Life*. London: Collins.

Bar-On, Dorit and Chrisman, Matthew. 2009. 'Ethical Neo-Expressivism' in Russ Shafer-Landau (ed.) *Oxford Studies in Metaethics* 4. New York: Oxford University Press.

Bar-On, Dorit; Chrisman, Matthew; Sias, James. 2014. '(How) Is Ethical Neo-Expressivism a Hybrid View?' in Guy Fletcher and Michael Ridge (eds.) *Having It Both Ways: Hybrid Theories and Modern Metaethics*. New York: Oxford University Press.

Blackburn, Simon. 1984. *Spreading the Word: Groundings in the Philosophy of Language*. Oxford: Clarendon Press.

Blackburn, Simon. 1993a. 'Errors and the Phenomenology of Value' in Blackburn 1993c.

Blackburn, Simon. 1993b. 'Supervenience Revisited' in Blackburn 1993c.

Blackburn, Simon. 1993c. *Essays in Quasi-Realism*. New York: Oxford University Press.

Blackburn, Simon. 1998. *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Clarendon Press.

Boghossian, Paul. 1989. 'The Rule-Following Considerations'. *Mind* 98: 507-549.

Boghossian, Paul. 1990. 'The Status of Content'. *The Philosophical Review* 99: 157-184.

Brandom, Robert. 1994. *Making It Explicit*. Cambridge, MA: Harvard University Press.

Brandom, Robert. 2000. *Articulating Reasons: An Introduction to Inferentialism*. Cambridge, MA: Harvard University Press.

Brandom, Robert. 2008. *Between Saying and Doing: Towards an Analytic Pragmatism*. New York: Oxford University Press.

Brandom, Robert. 2010a. 'Reply to Michael Dummett's "Should Semantics Be Deflated?"' in Bernhard Weiss & Jeremy Wanderer (eds.) *Reading Brandom: On Making It Explicit*. Abingdon: Routledge.

Brandom, Robert. 2010b. 'Reply to Jerry Fodor and Ernest Lepore's "Brandom Beleaguered"' in Bernhard Weiss & Jeremy Wanderer (eds.) *Reading Brandom: On Making It Explicit*. Abingdon: Routledge.

Burgess, Alexis. 2011. 'Mainstream Semantics + Deflationary Truth'. *Linguist and Philos* 34: 397-410.

Burgess, Alexis, and Sherman, Brett (eds). 2014a. 'A Plea for the Metaphysics of Meaning' in Burgess and Sherman (2014b).

Burgess, Alexis, and Sherman, Brett (eds). 2014b. *Metasemantics: New Essays on the Foundations of Meaning*. Oxford: Oxford University Press.

Castaneda, Hector-Neri. 1975. *Thinking and Doing: The Philosophical Foundations of Institutions*. Berlin: Springer.

Chrisman, Matthew. 2015. *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics*. Oxford: Oxford University Press.

Darwall, Stephen. 1997. 'Reasons, Motives, and the Demands of Morality: In Introduction' in Stephen Darwall, Allan Gibbard and Peter Railton (eds.) *Moral Discourse and Practice: Some Philosophical Approaches*. New York: Oxford University Press.

Darwall, Stephen; Gibbard, Allan; and Railton, Peter. 1992. 'Toward *Fin de siècle* Ethics: Some Trends'. *Philosophical Review* 101: 115-189.

Davidson, Donald. 1984a. 'Theories of Meaning and Learnable Languages' in Davidson (1984c).

Davidson, Donald. 1984b. 'Truth and Meaning' in Davidson (1984c).

Davidson, Donald. 1984c. *Inquiries Into Truth and Interpretation*. New York: Oxford University Press.

Davis, Wayne. 2003. *Meaning, Expression, and Thought*. Cambridge: Cambridge University Press.
Devitt, Michael. 2010. *Putting Metaphysics First*. New York: Oxford University Press.
Dreier, James. 1996. 'Expressivist Embeddings and Minimalist Truth'. *Philosophical Studies* 83: 29-51.
Dworkin, Ronald. 1996. 'Objectivity and Truth: You'd Better Believe It'. *Philosophy & Public Affairs* 25: 87-139.
Field, Hartry. 'Tarski's Theory of Truth'. *The Journal of Philosophy* 69: 347-375.
Fodor, Jerry. *A Theory of Content and Other Essays*. Cambridge MA: MIT Press.
Fodor, Jerry and Lepore, Ernie. 2001a. 'Brandom's Burdens: Compositionality and Inferentialism'. *Philosophy and Phenomenological Research* 63: 465-481.
Fodor, Jerry and Lepore, Ernie. 2001b. 'Why Compositionality Won't Go Away: Reflections on Horwich's 'Deflationary' Theory'. *Ratio* 14: 350-368.
Fodor, Jerry and Lepore, Ernie. 2007. 'Brandom Beleagured'. *Philosophy and Phenomenological Research* 74: 677-691.
Frege, Gottlob. 1963. 'Compound Thoughts'. *Mind* 72: 1-17.
Geach, P. T. 1960. 'Ascriptivism'. *The Philosophical Review* 69: 221-225.
Geach, P. T. 1965. 'Assertion.' *The Philosophical Review* 74: 449-465.
Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Oxford: Clarendon Press.
Gibbard, Allan. 2003. *Thinking How To Live*. Cambridge, MAL Harvard University Press,
Harman, Gilbert. 1973. Review of *The Significance of Sense: Meaning, Modality, and Morality*. *Philosophical Review* 82: 235-239.
Horgan, Terence and Timmons, Mark. 1992. 'Troubles on Twin Earth: Moral Queerness Revived'. Synthese 92: 221-260.
Horwich, Paul. 1998a. *Meaning*. Oxford: Clarendon Press.
Horwich, Paul. 1999b. *Truth* (2nd ed). Oxford: Clarendon Press.
Horwich, Paul. 2005. *Reflections on Meaning*. Oxford: Clarendon Press.
Kalderon, Mark. 2005. *Moral Fictionalism*. New York: Oxford University Press.
Kaplan, David. 1989. 'Afterthoughts' in J. Almog, J. Perry, and H. Wettstein (eds) *Themes from Kaplan*. Oxford: Oxford University Press.
Krazter, Angelika. 1981. 'The Notional Category of Modality' in Hans Eikmeyer and Hannes Reiser (eds.) *Words, Worlds, and Contexts*. Berlin: de Gruyter.
Krazter, Angelika. 1991. 'Conditionals' in Armin von Stechow and Dieter Wunderlich (eds.) *Handbuch Semantik*. Berlin: de Gruyter.
Kukla, Rebecca and Winsburg, Eric. 'Deflationism, Pragmatism, and Metaphysics'. in Steven Gross, Nicholas Tebben and Michael Williams (eds.) *Meaning Without Representation: Truth, Expression, Normativity, and Naturalism*. New York: Oxford University Press.
Lewis, David. 1984. 'Putnam's Paradox'. *Australasian Journal of Philosophy* 62: 221-236.
Mackie, J. L. 1977. *Ethics: Inventing Right and Wrong*. London: Pelican.
Macarther and Price. 2011. 'Pragmatism, Quasi-Realism, and the Global Challenge' in Price (2011b).
Miller, Alex. 2013. *Contemporary Metaethics: An Introduction* (2nd ed). Cambridge; Malden MA: Polity.
Millikan, Ruth. 1984. *Language, Thought, and Other Biological Categories*. Cambridge MA: MIT Press.
Ogden, C. K., and Richards, I. A. 1923. *The Meaning of Meaning: A Study of the Influence of Language Upon Thought and the Science of Symbolism*. New York: Harcourt, Brace & World, Inc.
Pérez Carballo, Alejandro. 2014. 'Semantic Hermeneutics' in Burgess & Sherman (2014a).
Price, Huw. 2011a. 'One Cheer for Representationalism' in Price (2011b).
Price, Huw. 2011b. *Naturalism Without Mirrors*. New York: Oxford University Press.

Price, Huw. 2013. *Expressivism, Pragmatism and Representationalism*. Cambridge: Cambridge University Press.

Price, Huw. 2015. 'Idling and Sidling Toward Philosophical Peace' in Steven Gross, Nicholas Tebben and Michael Williams (eds.) *Meaning Without Representation: Truth, Expression, Normativity, and Naturalism*. New York: Oxford University Press.

Railton, Peter. 1989. 'Naturalism and Prescriptivity'. *Social Philosophy & Policy* 7: 151-174.

Ridge, Michael. 2014. *Impassioned Belief*. Oxford: Oxford University Press.

Schiffer, Stephen. 2003. *The Things We Mean*. New York: Oxford University Press.

Schroeder, Mark. 2007. *Slaves of the Passions*. New York: Oxford University Press.

Schroeder, Mark. 2008a. 'What is the Frege-Geach Problem?'. *Philosophical Compass* 3: 703-720.

Schroeder, Mark. 2008b. *Being For*. Oxford: Clarendon Press.

Schroeder, Mark. 2013. 'Two Roles for Propositions: Cause for Divorce?'. *Nous* 47: 409-430.

Schroeder, Mark. 2015. 'Higher-Order Attitudes, Frege's Abyss, and the Truth in Propositions' in Johnson, R. N. & Smith, M. (eds.) *Passions and Projections: Themes from the Philosophy of Simon Blackburn*. Oxford: Oxford University Press.

Silk. Alex. 'How to Be an Ethical Expressivist'. *Philosophy and Phenomenological Research* 91: 47-81.

Soames, Scott. 2010. *What is Meaning?* Princeton: Princeton University Press.

Soames, Scott. 2013. 'Cognitive Propositions'. *Philosophical Perspectives* 27: 479-501.

Soames, Scott. 2014. 'Cognitive Propositions' in Jeffrey King, Scott Soames and Jeff Speaks (eds.) *New Thinking about Propositions*. Oxford: Oxford University Press.

Stalnaker, Robert. 1997. 'Reference and Necessity' in B. Hale and C. Wright (eds) *A Companion to the Philosophy of Language*. Oxford: Blackwell.

Unwin, Nicholas. 1999. 'Quasi-Realism, Negation and the Frege-Geach Problem'. *The Philosophical Quarterly*, 49: 337-352.

Unwin, Nicholas. 2001. 'Norms and Negation: A Problem for Gibbard's Logic'. *Philosophical Quarterly*, 51: 60-75.

Williams, Michael. 1999. 'Meaning and Deflationary Truth'. *The Journal of Philosophy* 96: 545-564.

Wedgwood, Ralph. 2007. *The Nature of Normativity*. New York: Oxford University Press.

Wright, Crispin. 1992. *Truth and Objectivity*. Cambridge, MA: Harvard University Press.

Yalcin, Seth. 'Semantics and Metasemantics in the Context of Generative Grammar' in Burgess and Sherman (2014a).

Zalabardo, José. manuscript. 'Empiricist Pragmatism'.