

2013-12-13

Lightweight Moral Realism: Objectivity and Reasoning Without Heavyweight Facts

Mark Douglas Warren

University of Miami, MarkDouglasWarren@gmail.com

Follow this and additional works at: https://scholarlyrepository.miami.edu/oa_dissertations

Recommended Citation

Warren, Mark Douglas, "Lightweight Moral Realism: Objectivity and Reasoning Without Heavyweight Facts" (2013). *Open Access Dissertations*. 1138.

https://scholarlyrepository.miami.edu/oa_dissertations/1138

This Open access is brought to you for free and open access by the Electronic Theses and Dissertations at Scholarly Repository. It has been accepted for inclusion in Open Access Dissertations by an authorized administrator of Scholarly Repository. For more information, please contact repository.library@miami.edu.

UNIVERSITY OF MIAMI

LIGHTWEIGHT MORAL REALISM: OBJECTIVITY AND REASONING WITHOUT
HEAVYWEIGHT FACTS

By

Mark Douglas Warren

A DISSERTATION

Submitted to the Faculty
of the University of Miami
in partial fulfillment of the requirements for
the degree of Doctor of Philosophy

Coral Gables, Florida

December 2013

©2013
Mark Douglas Warren
All Rights Reserved

UNIVERSITY OF MIAMI

A dissertation submitted in partial fulfillment of
the requirements for the degree of
Doctor of Philosophy

LIGHTWEIGHT MORAL REALISM: OBJECTIVITY AND REASONING WITHOUT
HEAVYWEIGHT FACTS

Mark Douglas Warren

Approved:

Amie Thomasson, Ph.D.
Professor of Philosophy and
Parodi Senior Scholar in Aesthetics

Simon Evnine, Ph.D.
Associate Professor of Philosophy

Bradford Cokelet, Ph.D.
Assistant Professor of Philosophy

M. Brian Blake, Ph.D.
Dean of the Graduate School

Jamie Dreier, Ph.D.
Professor of Philosophy
Brown University

WARREN, MARK DOUGLAS
Lightweight Moral Realism: Objectivity and
Reasoning Without Heavyweight Facts

(Ph.D., Philosophy)
(December 2013)

Abstract of a dissertation at the University of Miami.

Dissertation supervised by Professor Amie Thomasson.
No. of pages in test. (173)

In this dissertation, I develop a non-representational approach to metaethics that avoids the ontological and epistemic problems that traditional realist views face: accounting for the nature of moral facts or properties and our knowledge of these facts in a way that comports with naturalism. According to expressivism—the most popular non-representational account in contemporary metaethics—moral claims don't aim primarily to represent moral facts or properties, but instead function to express our non-cognitive attitudes like approval or disapproval. This is a promising approach, but the view has substantial difficulties accounting for important features of moral discourse, like objectivity and complicated moral reasoning. I develop an alternative to expressivism: metaethical inferentialism, which I argue gives us a new way to deal with the problems expressivism faces. If we understand the meaning of moral concepts by way of the inferential roles they have instead of the attitudes they express, we can see why these concepts have meanings that stay constant even in contexts where they aren't used in the typical way, and in doing so solve the infamous Frege-Geach problem. This approach offers simple answers to at least two problems that have vexed contemporary expressivists—the “problem of permissions” and the commitment to “mentalism”, both of which are problems readily solved by an inferentialist approach. I also argue that moral objectivity isn't best explained by reference to some realm of moral facts waiting to be

accurately described by a successful discourse, but that instead the objectivity of morality is a function of the inferential roles that we should expect moral concepts to have, given their purpose—the coordination of social behavior.

For Stephanie, my very favorite person.

Acknowledgments

I'd first like to thank Amie Thomasson for offering countless hours of conversation, reading, revision, and support as I wrote this dissertation—I couldn't have hoped for a better advisor, and I can't imagine having done this without her help. H!(Amie). I also want to thank my excellent committee members for their help and encouragement—Brad Cokelet, Simon Evnine, and my outside committee member Jamie Dreier, one of my favorite metaethicists in the world. I want to thank my wonderful and supportive parents, Chris and Barbara Warren, who came around to the idea of having a philosopher for a son. I should also thank my brother, Dr. Adam Warren, who inspired me to pursue graduate school by getting a Ph.D. in a less challenging field. I'd like to thank all the wonderful friends I made at the University of Miami: Noel Alphonse, for agreeing to various virtue-inducing wagers over the years; Ryan Lake, for being such a marvelous delight; Micah Dugas, for his childlike wonder; all the Hufflepites—Ryan, Micah, Ben, Robin, Fredrik, and Stephanie—for giving me a home in Miami; and Ben Burgis, for showing the honesty and strength of character to change his philosophical commitments about qualia and the analytic/synthetic distinction in the face of overwhelming evidence. I want to thank my best friend Chad Jackson, who by my reckoning owes me upwards of 20,000 dollars. And I especially want to thank my sweetie, Stephanie Saline, for her patience and love.

Table of Contents

Chapter One: Non-Representationalism and Accommodation	1
Chapter Two: Objectivity without (Heavyweight) Realism	27
Chapter Three: The Frege-Geach Problem	50
Chapter Four: Moral Inferentialism and the Frege-Geach Problem	83
Chapter Five: Objectivity for Lightweights	128
References	169

CHAPTER ONE: NON-REPRESENTATIONALISM AND ACCOMMODATION

1. Overview

In this dissertation, I develop and defend *moral inferentialism*, a *non-representational* response to some old questions in metaethics: How do moral facts fit in the world? Given the restrictions of naturalism, how can we make sense of moral properties or states of affairs? If moral facts are just a species of physical facts, why do they seem to have motivational import; why is it that coming to understand that lying is wrong can move me to decide I won't lie?

Non-representationalism about metaethics doesn't attempt to answer these questions by inquiring into the nature of moral truthmakers, but instead seeks "to dismiss or demote such metaphysical puzzles in favour of more practical questions, about the roles and functions of the matters in question in human life." (Macarthur and Price, 94) This approach in metaethics traces its roots to the emotivism of A.J. Ayer (1952). On his account, propositions with moral predicates do not primarily function to express beliefs, but instead express evaluative attitudes like approval or disapproval. "Sentences which simply express moral judgements do not say anything. They are pure expressions of feeling and as such do not come under the category of truth and falsehood." (108)

Emotivism's contemporary successor is the expressivism advocated Simon Blackburn and Allan Gibbard. According to this tradition, moral claims do say something, but their meaning is a function of the essentially evaluative attitudes they express. For this reason, emotivism and expressivism are considered types of "non-cognitivism". From a

strictly non-cognitivist perspective, moral propositions do not express cognitive mental states—that is, beliefs—and so are not appropriately understood in terms of their truth values; like the sentence, “Hurrah for the Lakers!”, moral claims cannot be either true or false. This approach is characteristic of metaethical non-representationalism in two ways: first, it denies that moral assertions function to report on moral facts. Second, it doesn’t take this state of affairs to entail an error theory about morality; instead, the approach is to look at moral discourse itself in a new way, such that the absence of moral facts does not undermine the discourse. Again, according to Ayer, moral statements are essentially evaluative, and not truth-apt. Because the question of a moral statement’s truth shouldn’t come up, we have no reason to worry that moral statements do not correspond to any moral facts in the world.

The metaethical inferentialism I argue for differs from emotivism and expressivism in many ways, chief among them my account of the function of moral claims. I’ll argue that moral claims don’t function primarily to express non-cognitive mental states like approval and disapproval, but instead derive their meaning from the inferential import they have. This means that my approach is not best characterized as a type of non-cognitivism; as we will see, one of the advantages of an inferential account is that it makes sense of our normal talk about moral *beliefs*—and even, as we will also see, our talk about moral *truths*. I’ll argue in this dissertation that this is the primary advantage inferentialism has over traditional non-representational approaches: it offers us a more conciliatory account of moral discourse; it offers us a way to make sense of our intuitive understanding of moral thought and practice, especially as they regard the phenomena of moral objectivity and ethical reasoning.

Nevertheless, one might understandably think that my view, like Ayer's, must entail a kind of revisionism about moral discourse: if moral claims don't function primarily to report on some range of truthmakers, it might be thought that our commonsense ways of thinking and talking about moral topics is fundamentally misleading. But we shouldn't jump to that conclusion too quickly; to do so underestimates the tools available to non-representational accounts generally. The dialectic between representational and non-representational approaches isn't limited to metaethics; it's a pattern we find repeated in other domains of philosophy. Traditional *representational* approaches to philosophical topics like epistemology, truth, modality, etc. work under the common assumption that these discourses function primarily to describe particular kinds of objects, relationships, properties, or states of affairs, and as such the important philosophical work to be done is an investigation of the nature of these objects or properties.

This is a tempting starting point, but those who take it must wrestle with difficult questions—what Huw Price calls the placement problem: how do we place these objects in the world as it's described by science? How is it that we natural beings are able to come to know about such objects? Non-representationalists like Huw Price, Simon Blackburn, and Paul Horwich endeavor to avoid these questions by denying that the primary aim of the discourse in question is representation of a special sort of object. They instead explain the discourse in question by identifying the purpose it serves in human life.

It will be helpful here to first characterize *representationalism* about metaethics—to see how Price's placement problems manifest in the ethical domain. In what follows, I'll take you on a very brief tour of meta-ethical representationalism. I'll then explain the

non-representational approach in contrast to these views, first using Amie Thomasson's modal normativism as a clear exemplar of non-representationalism in another domain. Thomasson's view and my own share the same advantages, and the same motivations: offering an account of our respective discourses in ways that promise to avoid the placement problem. Before outlining the structure of this dissertation, I'll present the fundamental aim of my dissertation—offering a metaethical account that makes sense of our common intuitions about moral discourse.

2. Representationalism in Metaethics

The question that is characteristic of representational approaches to metaethics is this: "What makes moral claims true?" One way a representationalist might answer this question is by positing that the truthmakers for moral claims are a type of non-natural fact. This is the position of ethical *non-naturalism*, the view most closely associated with GE Moore. Non-naturalists like Moore hold that moral facts are fundamentally distinct from the kinds of facts that we discover with the natural sciences. Moore argues that moral facts can't be reduced to or otherwise analyzed in terms of natural facts, and so moral properties are *sui generis*, entirely unlike the empirical properties science investigates.

The great challenge for non-naturalism in modern philosophy comes from the widespread (but by no means universal) commitment to naturalism, the belief that nothing exists apart from the natural world—elements, properties, and events that are the proper study of science. How are we to understand the nature of these *sui generis* moral facts in a naturalistically respectable manner? Furthermore, how do creatures such as ourselves—

hairless language-using apes that evolved by natural selection—apprehend truths in ethics? Do we have some (non-natural, perhaps) faculty or intuition that connects us to these facts? How would such a faculty interact with our physical bodies to produce reliable moral judgments? The only answers available seem to violate the restriction that naturalism imposes.

Representational *ethical naturalists* begin the hunt for these moral truthmakers in the natural world. One way of doing this is by seeking a reduction of moral properties to some set of natural properties—properties that are appropriately the subject matter of some science, such as biology or psychology. We might say, for example, “War is morally bad” is true because war has some physical property—maybe the propensity to cause great suffering, or perhaps more specifically the tendency to activate the neural correlates of pain in human brains—that the term “bad” refers to. We can understand this sort of naturalism in terms of its metaphysical and semantic components: Metaphysically, the wrongness (or rightness, or permissibility) of an act can be identified with its natural or physical properties. Semantically, this is what we *mean* when we say it is wrong—we’re just saying that this act happens to have these properties.¹

¹ But as a semantic thesis, this naturalist project seems wanting. Moore (1903: 67-70) pointed out that if the meanings of moral claims pick out natural properties, then if *N* is the property that makes acts wrong, questions like

Act *X* has natural property *N*, but is it wrong?

Act *X* is wrong, but does it have natural property *N*?

should be closed. That is, any person who understands what it is for an act to be wrong will know, just from understanding that *X* has property *N*, whether it is wrong, and conversely will know, just from understanding that *X* is wrong, whether it has property *N*. But such questions aren’t closed—it will always be possible for a person who can competently use moral language to wonder whether or not an act that is wrong will always have a particular natural property, and to wonder whether or not with a particular natural property is wrong. That is, it won’t be analytically obvious to a competent speaker that act *X* is wrong because it has natural property *N*, just on the basis that it has that property. But if our moral terms picked out natural properties, Moore argued, we should expect that such questions like Q1 and Q2 would be closed to a competent speaker, in just the way that

He is an unmarried male, but is he a bachelor?

is closed, for anyone who understands what a bachelor is.

Both natural and non-natural metaethics have trouble accounting for the *prescriptivity* of moral claims. They seem to miss out on the “oomph” of moral talk, what Kant would call its categorical nature. If the meaning of moral terms is given by the physical or non-natural properties they refer to, why do we think these properties are so very important? Why do moral truths seem to give us reasons to act in certain ways? Usually, when we come to accept that something is the right thing to do, this understanding has a kind of motivational import: the right thing to do is, after all, the thing to be done; when we understand that something is right, we typically find ourselves moved to do it. And this is something these approaches struggle to account for.

According to the Humean picture of human motivation, there are two principle kinds of psychological state. There are beliefs—cognitive states—and there are desires—conative, or non-cognitive states. These are fundamentally different mental states, according to Hume: beliefs are apt to be true or false, but desires do not purport to describe the world. “For any belief and desire pair that we imagine, we can always imagine someone having the desire but lacking the belief, and vice versa.” (Smith 1994:7)

Modern day naturalists recognize the power of Moore’s “Open Question” argument, but try to defend the metaphysical thesis of naturalism—that moral claims have their truthmakers in the natural world—by altering its semantic thesis. We don’t know *a priori* that moral properties are identical to natural properties in the way that we know bachelors are identical to unmarried males, they argue. Instead we discover these relationships *a posteriori*, in the same way that we discover the identity between water and H₂O. There is nothing in the meaning of water that makes it obvious that it is identical to H₂O; we have to engage in empirical research to discover this relationship. Moral philosophers, then, are not unlike scientists, investigating the natural world for an explanation of the properties that make our claims true or false.

This is a promising response, but it’s still left to the naturalist to explain how it is that moral claims have the significance that they do: in virtue of what does a moral claim refer to its moral truthmaker? Philosophers like David Brink, Richard Boyd, and Nicholas Sturgeon posit a kind of causal theory of reference much like that Kripke (1980) uses to explain how the term “water” comes to rigidly designate H₂O, wherein our uses of moral terms rigidly designate the natural property in the world that causally regulates our common use of these terms¹. So whereas Kripke posited that our use of the term ‘water’ is causally regulated by a chemical natural kind—H₂O—these philosophers posit that our use of an ethical term like “good” is causally regulated by some other natural property—perhaps the tendency of some acts to increase human happiness. This is an ingenuous response, though one might worry whether there is indeed some unique physical property that plays this role in our use of moral terms.

What's important is that we normally cite desires, not beliefs, to account for our motivations. We cannot look to our beliefs on their own to explain our reasons for acting:

To cite a cognitive propositional attitude [i.e., a belief] is to give at most a partial specification of a reason for acting; to be fully explicit, one would need to add a mention of something non-cognitive, a state of the will or a volitional event. (McDowell 1981: 155)

So for example, if you wanted to know what my reason was for going to the refrigerator, it wouldn't suffice to know that I believed there was celery in the vegetable drawer. You'd also need to know that I wanted to eat some celery. It wouldn't make sense to say that I went to the refrigerator because of my belief that celery was there, if I didn't want celery.

But our moral beliefs do seem to give us reasons to act, just on their own. If I come to believe that eating meat is wrong, this means that I have a reason to not eat meat, and in contrast to the standard Humean model, it doesn't seem like I need to cite any particular conative state on my part to explain this reason. I have a reason to not eat meat, whatever my desires are. This seems to be built into our use of moral claims. If we went to a foreign land where the people had terms that referred to just the same physical properties our own moral terms putatively refer to, but the community didn't take these terms to have any practical significance for how they should behave, would we consider their terms to be moral? Many think not: an essential feature of the meaning of moral terms is the impact they have on how we live².

JL Mackie held that the prescriptivity of morality argues against the very idea of there being moral properties:

If there were objective values, then they would be entities or qualities or relations of a very strange sort, utterly different from anything else in the universe. Correspondingly, if we were aware of them, it would have to be by some special

² See (Hare 1952: 146-149)

faculty of moral perception or intuition, utterly different from our ordinary ways of knowing everything else. (Mackie 1977, 38)

Some take this to indicate that there are moral facts, but they are subjective—the motivational import they have depends on some contingent desire on our part. Our objective way of speaking about moral truths is misleading, and so some revision is called for. Mackie argues there can be no facts that on their own prescribe what is to be done, and so he cuts to the chase and proclaims all moral talk to be in fundamental error, like talk committed to the existence of gnomes or phlogiston. Moral claims do attempt to represent some special sort of objective facts, but there are no such facts, so moral claims fail. Our hunt for truthmakers, he argues, is doomed to come up short.

Where the representationalist understands moral claims as expressions of cognitive states that attempt to represent some aspect of the world, non-representationalists see such claims as expressions with very different functions. If the primary goal of ethical claims isn't the representation of moral facts, non-representationalists needn't provide an analysis of the nature of such facts, and so needn't worry about violating the restrictions imposed by naturalism. And because the function associated with moral claims is characteristically motivating—one implication of my belief that lying is wrong is that I feel pressure to not lie—the non-representationalist has no problem explaining the motivational 'oomph' of moral assertions.

This is a characteristic benefit of taking a non-representational approach: by understanding the discourse in question according to its pragmatic function, as opposed to its truthmakers, the non-representationalist gains an important advantage over the representationalist. The placement problem needn't be faced, because it doesn't come up in the first place. And the motivational import is a natural consequence of the non-

representational understanding of moral discourse. In the next section, I want to illustrate how this advantage plays out for such an approach in another domain.

3. Modal Non-representationalism

Amie Thomasson argues that if we characterize the role that we should expect *modal* concepts and terms to play given the purpose of metaphysical modal talk, we can explain features of that discourse (e.g., how we come to know modal truths, how there can be modal facts, how they can be objective) without having to explain the underlying nature of modal properties. The advantages of this approach over traditional representational options are important to my project. Like metaethical inferentialism, we will see that Thomasson's modal normativism avoids difficult placement questions about modal truths. Her view will play a key role in both the development and exposition of my metaethical position; it will be helpful to explain the view and its motivations.

Ayer's non-representational treatment of necessity gives us a good model for the origin of modal normativism. His positivistic convictions led him to deny that modal statements were factual at all; instead he held that they were analytic statements—empty tautologies, none of which “provide any information about any matter of fact.” (Ayer, 79). This left him in the awkward position of explaining how it is that a proposition without factual content could be informative, as many necessary claims seem to be. He explains this phenomenon by arguing that such claims “do enlighten us by illustrating the way in which we use certain symbols.” (ibid.) In making such an illustration, one is not

depending on any matters of fact, but is instead “simply calling attention to the implications of a certain linguistic usage.” (ibid.) On Ayer’s approach, claims of necessity are at bottom analytic claims, and depend in some way on the linguistic conventions that govern our use of terms.

The way in which we understand this dependence is crucial, however. The most common (mis)construal of “modal conventionalism” is one in which the conventions of language themselves act as the truth-makers for our claims of necessity or possibility. Such a reading of the position is disastrous, because it collapses the view into a kind of idealism: if a sentence like “Necessarily, all bachelors are male” is true just because the term “bachelor” is only used to refer to unmarried males, we could presumably render the proposition false, just by adopting a new linguistic convention. The view ends up looking something like magic.

But it is a misconstrual, and modal normativism highlights this by stressing that modal claims do not function to describe our linguistic conventions—they don’t function to *describe* at all. Instead we see the function as non-representational: we employ modal language to reflect (and endorse) the kind of semantic commitments we’ve made with the terms in use. The role of modal vocabulary is not to report on modal truth-makers, but is instead to “make explicit the semantic, conceptual connections that are already implicit in the use of ordinary empirical vocabulary. (Brandom, 2007: 2)” The modal normativist can dispel the charge of idealism by pointing out that while we can make a decision to change the way we use the term “bachelor”, doing so won’t mean that, given the way we use the term now, it could turn out as possible that some bachelors aren’t male. The choices we

make about how to use terms may be (in part) up to us, but the truth values of claims made with those terms is up to the world.³

Like ethical expressivists, modal normativists deny that claims in their considered domain are meant to be representations of metaphysical facts, but insist instead that modal claims serve a non-representational function—we use modal claims to convey the semantic rules that govern our use of terms. The idea is that these semantic rules play a central constitutive role in our linguistic use, much in the same way that, e.g., the rules of chess play a central role in one’s use of chess pieces. A knight in chess is in a sense constituted by the rules of the game—if it weren’t governed by the rules that say how it can move, it wouldn’t count as a knight. In the same way, expressivists argue, the semantic rules governing a term are constitutive of the meaning of that term. We use modal vocabulary to illuminate those rules we engage in with our use of normal language.

So the modal normativist contends that the sentence:

(MS) Necessarily, bachelors are men.

³ Of course these old objections are not the only ones with which modal normativists must wrestle; the position is a relatively old one, and some modern developments in analytic philosophy need to be addressed. For example, Saul Kripke’s work on synthetic necessities—modally charged claims that we must go into the world to empirically discover—indicates a challenge. How can we explain all necessities by way of analyticity, as Ayer hoped to do, if there are some necessities, such as “Water is necessarily H₂O”, which are not analytic in nature? This objection has been treated as conclusively damning of the conventionalist project, but it needn’t be. Alan Sidelle forwards plausible response: To understand how a statement like “water is H₂O” is necessary, we have to look at the individuating principles which govern our use of natural kind terms like “water”. (Sidelle 1989) For such terms, we will find a framework analytic principle in effect, like:

Natural kinds have their deep explanatory structure—whatever it is—necessarily.

Such an analytic principle is supplemented by non-modal empirical statements like:

Water has H₂O as its deep explanatory structure.

These taken together will yield a necessary a posteriori statement. Notice that the necessity is not out there in the world, waiting to be discovered; only the non-modal truth about water’s deep explanatory structure is a posteriori and depends on the world for its truth. The modal import of

□(Water = H₂O)

is tucked safely away in the analytic principles governing our use of natural kind terms, where it does no harm to the normativist’s basic thesis.

is best understood in terms of a conveyance of a linguistic rule: “Use the term ‘bachelor’ only if you may also use the term ‘man’.” Such a rule is a requirement of correctly using the term ‘bachelor’.

Notice also that such individuating principles aren’t limited to linguistic requirements. Rules of use also cover permissible linguistic moves. We probably originally learn these simply by ostensive definition as children (“That is a dog, and so is that, even though it only has three legs.”), but we can more directly express these rules by use of the modal term, ‘possible’, as in:

It is possible for a dog to only have three legs.

Such a modal claim reflects that a particular linguistic rule is in effect regarding the term “dog”—namely that we can refer to an animal with that term even if it doesn’t have the normal number of legs. (Thomasson, 2007a)

So much for the basics. What is the motivation for this position? This normative approach contrasts with traditional modal accounts, which are largely concerned with finding the truth-makers for modal claims, or the conditions or states of affairs that would make modal statements true. David Lewis (1986), for example, accounts for the truth of modal claims in terms of possible worlds—innumerable causally isolated worlds that are as real as our own. According to Lewis, the proposition, “Necessarily, bachelors are men” is true, because in all possible worlds that have bachelors those bachelors are men.

I sketch Lewis’ take on modality only to gesture towards the motivation of the modal normativist. Because it places truth conditions for modal statements in representational terms of correspondence with some realm of modal facts—be they possible worlds, combinations of states of affair, etc—realism about modality must

answer difficult metaphysical and epistemological questions. If, for example, necessity is a matter of truth across all possible worlds, we should ask just what sort of things these possible worlds could be. Lewis provides a notoriously fantastical ontology in response, but even more conservative modal realists must wrestle with Hume's observation on the topic: The possible (and the necessary) outstrip the actual, and so it seems to explain the possible, we must go beyond what is found in the actual world. (Hume, 1999)

Even if we can find a way to make such a position sit comfortably within the constraints of naturalism, we must then give an adequate account of how we come to know that modal statements are in fact true. No amount of knowledge about the way things are, the expressivist argues, is sufficient to justify any claim about the way things must be, or could be. Reflection on how we do in fact justify our modal knowledge will be of no help to the modal realists. For example, how do we know that water is necessarily H₂O? We can't come to this conclusion purely on the basis of empirical investigation; we must also undertake the kind of counterfactual imaginings that Kripke advises. We think of a possible world in which some substance has all the features of water here on earth—a drinkable liquid you can swim in, etc.—but which is not itself H₂O. Could such a thing be water? If we answer “no”, we might be committed to the necessity of the identity between water and H₂O—we can't conceive of it being otherwise. But why should we imagine that conceivability can give us justified beliefs about modal truth? From the realist perspective, there's no ready story for why the limits of our imagination should be expected to track the non-modal facts that give truth-conditions for modal propositions; what's conceivable could just as well outstrip what is possible, or vice-versa.

The modal normativist, on the other hand,

doesn't think of acquiring modal knowledge as a matter of coming to see new, different features of the world, still less of perceiving some other possible worlds or a platonic world of essences. Instead, [the normativist] demystifies modal knowledge by considering the move from using language to knowing basic modal facts to be a matter of moving from mastering the rules for properly applying and refusing expressions (as a competent speaker), to being able to explicitly convey these constitutive rules. (Thomasson 2008: 150)

We can understand the sort of imaginative tests Kripke applies as part of that process—considering counterfactual possibilities brings us to a more explicit awareness of how we use a term in question, and can even be a part of coming to decide on new uses for that term.

Notice that treatments of issues in metaethics and modality deal with parallel concerns. What sorts of facts or features make modal or ethical claims true? Where are we to place moral and modal facts within a natural framework? David Hume suggests we cannot: the possible and the necessary outstrip the actual. This is what led Hume to his quandary about the 'necessary connexions' found in nature. No amount of knowledge about the way things are is sufficient to justify any claim about the way things must be or could be. Hume noted a similar problem for ethics: what is the case cannot on its own determine what ought to be the case. (Hume, 1968)

The modal normativist sidesteps these problems by refusing to see the fundamental question about modality as one about modal truthmakers; she instead looks to the purpose that modal discourse plays in our lives. So too the metaethical non-representationalist avoids the placement problem for ethics: we needn't posit any of Mackie's "queer properties" to make ethical sentences true, because (roughly speaking) the primary goal of ethical discourse is not to assert any matter of fact—though, as we'll see, we can make sense of their truth without adverting to truthmakers.

Both forms of non-representationalism share advantages over their representational counterparts, but they also share a similar burden: making sense of our normal moral or modal discourse without presupposing a robust range of moral or modal facts or properties. A great part of the evolution of non-representationalism about ethics—including the work I do in this dissertation—can be understood as a task of *accommodation*.

The idea is to interpret the discourse in a manner that comports with so-called commonsense assumptions of that discourse; in other words, one wants to accommodate those commonplace assumptions. So there are features of ordinary moral discourse—for example, it is (or appears to be) fact-stating; we take there to be right answers to many moral questions; and so forth—that should (if possible) be accommodated by a plausible story about such discourse. (Timmons 1999: 11)⁴

The task of accommodation is the one undertaken by Simon Blackburn’s quasi-realist, who seeks to show that from a non-representational perspective

there is nothing illegitimate in our ordinary practice and thought. The respects in which we talk as if there are, for instance, moral facts, are legitimate. (Blackburn 1993, 216)

This is the task Amie Thomasson undertakes when she explains what it is to gain modal knowledge; it is the task Alan Sidelle undertakes when he shows how we can explain the truth of modal claims that don’t obviously redound to analytic truths⁵. It is this task that I will turn to for non-representationalism about metaethics.

⁴ Timmons calls this task “internal accommodation”, in contrast to “external accommodation”, the task of construing the metaphysical and epistemological commitments of moral discourse in a way that is consistent with one’s other philosophical commitments. I ignore the distinction here.

⁵ See fn. 2 above.

4. Accommodating Moral Discourse

As we've seen, emotivism is a radical position—again, according to Ayer, moral claims have no more content than exclamations of “Boo!” or “Hooray!”; it follows that common sense is wrong when it tells us that some moral sentence like, “Murder is wrong” must actually be true. This also means that genuine moral disagreements don't really exist—if you think abortion is wrong and I think it's permissible, there may be some implicit disagreement about empirical matters (like, say, the neural development of a fetus), but barring this, there can be no authentic “dispute about questions of value”. (Ayer 1952: 110)

Modern day ethical expressivists attempt to correct these extreme tendencies and protect our ordinary ways of thinking and talking about moral issues. In both moral and modal non-representationalism, minimalist theories of truth play an important role in the project of accommodation. Like other non-representational approaches, a minimalist or deflationary view of truth doesn't rely on correspondence with truth-makers, or successful representation of facts. According to the minimalist, there is no general way to characterize truth at all beyond the equivalence schema:

(ES) ‘p’ is true if and only if p.

Truth, the minimalist says, has no nature, and so no attempts to systematically theorize about it will succeed. All we have to understand the truth of sentences is a deflationary or disquotational sense: “Murder is wrong” is true iff murder is wrong.

The implication is that we can talk of metaphysical, mathematical, modal truth because that is just to repeat our commitments in these areas. But what is going on when we have such a commitment is to be understood in other terms. (Blackburn 1993: 78)

The idea is that I can move from the sentence, “Murder is wrong”, to the sentence, “It is true that murder is wrong”, without adding any real content to the original, and therefore do so without taking on any metaphysical burden of explanation. Blackburn calls such a maneuver “moving along Ramsey’s ladder”. And as he argues,

Because of the minimalism we can have for free what looks like a ladder of philosophical assent: ‘p’, ‘it is true that p’, ‘it is really and truly a fact that p’..., for none of these terms, in Ramsey’s view, marks an addition to the original judgement. You can as easily make the last judgement as the first—Ramsey’s ladder is lying on the ground, horizontal” (Blackburn 1993: 78-79)

So we can speak of murder being wrong, or it being a fact that murder is wrong, or murder having the property of wrongness; from the first, all the rest come for free. This introduces an important distinction: anti-realists like Mackie deny the existence of moral truths or properties altogether; non-representationalists, however, need not deny moral truths or properties, so long as they are understood to take a metaphysically deflated place on Ramsey’s ladder.

Amie Thomasson makes a similar trip along Ramsey’s ladder in the spirit of quasi-realism, to justify our right to claims about modal properties; following Stephen Schiffer (2003) she calls the metaphysically inconsequential steps along the ladder “pleonastic transformations”. They

enable us to restate modal truths redundantly in claims that entail the existence of modal facts redundantly in claims that entail the existence of modal facts or properties. All modal truths T may be restated redundantly as “It is a fact that T”, so from “Necessarily, all bachelors are men”, we can get “It is a fact that it is necessary that all bachelors are men. (Thomasson 2007: 149)

The important point here—moral and modal—is that we can speak of such facts or properties without committing ourselves to using them as explanations for what makes moral or modal claims true.

We can’t explain why poppies make us sleepy by appealing to the fact that they possess the dormitive virtue: true though it may be that they possess a dormitive

virtue, this is just a hypostatized way of expressing the fact that they make us sleepy and cannot be used to explain it. So similarly, if talk of modal facts and properties is simply derived from hypostatizations out of modal truths, any attempt to ‘explain’ modal truths by appeal to the existence of modal facts or properties would just yield a dormitive virtue explanation. This result, however, is not a liability for the expressivist, since she does not require modal truthmakers to explain modal truth. (Thomasson, *ibid.*)

I will argue in Chapter Five of my dissertation that while minimalism about truth gives the non-representationalist an essential tool for accommodating our commonsense intuitions about moral discourse, we must be careful that we’re not using it for an ill-suited purpose. I’ll argue specifically that neither minimalism about truth nor pleonastic transformations will allow us to move from “murder is wrong” to “it is *objectively true* that murder is wrong”.

The same quasi-realist motivations that lead us to a minimalist account of truth and facts may also bring us to reconsider the cognitivist/non-cognitivist distinction that originally characterized the difference between moral representationalism and non-representationalism. After all, if we want to account for the features of moral talk that make sense of our normal use of the truth predicate in connection with moral claims, and our talk of moral facts and properties, then surely we also need to make a place for moral *beliefs*. We can do this by making a distinction between “the idea that moral judgments are not primarily representational of moral properties and facts (non-representationalism), and the idea that moral judgments do not express beliefs (non-cognitivism).” (Timmons and Horgan 2006, 256) This is possible if we have a (what some might call ‘minimal’) notion of the conceptual connection between moral assertions and beliefs: “if someone makes an assertion, and is supposed sincere, it follows that she has a belief whose content can be captured by means of the sentence used.” (Wright, 1992: 14) The result of this line of thought is that moral judgments do express beliefs, at least when they are made

sincerely, because beliefs themselves can be non-representational.⁶ “Thus... one can be a cognitivist in ethics—someone who thinks that moral judgments are beliefs—but also be a non-representationalist in ethics, claiming that the content of such beliefs is not (or not primarily) representational, how-the-world-might-be-content.” (Timmons and Horgan 2006: 76)

The upshot of all of this is that by shedding some of the hereditary baggage of its predecessors, the modern day expressivist can secure a more robust, “realist-seeming” account of moral discourse. If we use minimalism to deny Ayer’s contention that moral propositions have no truth-value, we find a way to deny moral anti-realism and justify truth claims in moral discourse, with our non-representationalism securely in tow. As we will see in **Part One**, this move is important for expressivists wishing to answer charges of relativism and secure a robust objectivity. And by denying that beliefs must be representational, the expressivist opens up a space for rejecting non-cognitivism. These terminological distinctions—between anti-realism and realism, non-cognitivism and non-representationalism—reflect a more ambitiously quasi-realist strain of expressivism that has become popular recently. I will argue in **Part Two** of my dissertation that a final distinction needs to be drawn to round out this problem: Contemporary critiques of expressivist semantics (e.g., van Roojen 1996, Schueler 1988, Schroeder 2008) all proceed with the expressivist premise that moral language has a primarily expressive function. But this need not be the case; one may be a non-representationalist without being an expressivist. As we’ll see, I think there are good quasi-realist reasons to be the former, but not the latter.

⁶ see Timmons (1999: Ch. 12)

5. The structure of this dissertation

This dissertation is divided into two parts. In **Part One**, I discuss the methods and limits of the expressivist approach, considering in particular two common objections to the quasi-realist project in traditional expressivism: our commitment to moral objectivity, and our use of moral claims in complex reasoning. I argue that the most promising response for the expressivist to both of these objections comes from reflecting on the pragmatic import of moral discourse in human life. We see in each case that this approach goes some way towards answering the challenges, but that doing so with the tools of traditional expressivism leads to further difficulties.

In **Part Two** of the dissertation, I consider a non-representational account of morality that doesn't depend on the mental states expressivists must posit. I argue that this view, moral inferentialism, can pick up where traditional expressivism left off. Understanding the pragmatic import of moral discourse will still play a key role in the project of accommodation, but because inferentialism locates the meaning of moral claims in terms of their inferential import—rather than in terms of the mental states they function to express—it can offer responses to the challenges from objectivity and moral reasoning, responses that are both simpler and more robust than those offered by its expressivist counterpart.

5.1 Part One

One of the most common objections to metaethical views like emotivism and expressivism is the charge of relativism. Indeed, non-representational theories generally

face similar charges in their respective domains—whether the domain is moral, modal, or epistemological. The underlying worry is apparently that because non-representational theories do not appeal to a range of facts or properties as truthmakers, the expressivist has no way to non-arbitrarily adjudicate between contradictory claims. In **Chapter Two**, I look at traditional responses to this charge, show how they parallel responses available for the modal normativist, and argue that both forms of non-representationalism can claim a more convincing justification for objectivity.

My central claim in that chapter is that to justify objectivity within a non-representational account like expressivism, we must look to the function of that discourse. Moral discourse is objective, not because there is a unified world of ethical truth waiting to be described, but because the characteristic function of ethical discourse is one that cannot be fulfilled if contradictions are accepted. If moral discourse aims to coordinate social behavior, this aim is bound to be frustrated by accepting contradictions.

In this chapter, I also apply the same functional response to charges of ontological relativism against modal normativism, and show how it also has a pragmatic role—the orchestration of communication—that grants this objectivity-justifying characteristic. These considerations point to a proper criterion for objectivity in non-representational accounts generally: a non-representational discourse is objective iff the function of that discourse is one that cannot be fulfilled if contradictory claims are accepted.

Finally I turn to the implications of this criterion. I consider possible problems that might be raised for my account, for example cases in which it seems coordinated action is possible even when we accept contradictions—so long as the contradictions are not relevant to the behavior at hand.

In order to give an adequate account of a given area of discourse, a non-representationalist theory must not only show why it gives a better naturalistic account; it must also explain how the various propositions in that area of discourse have the inferential relationships they actually have in normal discourse. It must explain how we can use these claims—non-representational though they may be—in ways that make sense of the logically binding force of argumentative forms like Modus Ponens. In **Chapter Three**, I will begin by explaining the moral expressivist’s infamous Frege-Geach problem, and show that it is equally a problem for modal non-representationalists. My account will make use of a “higher-order theory” like the one Simon Blackburn proposes (1984); to clarify the problem and explore the implications of Blackburn’s formalization, I will use it as a framework for the modal solution to the same problem, substituting in operators that reflect commitments to linguistic rules of use for Blackburn’s operators that express commitments to approval and disapproval. Using this formalization shows us how we can answer part of the Frege-Geach problem—the higher-order account avoids the charge of equivocation by holding the potential expressive force of these clauses as constant.

But once this task is accomplished, we can’t just stipulate that the formalization is offers an adequate model of our actual use of moral terms. We have to explain why it is that the non-representational way we use moral or modal language is such that it can produce valid and invalid arguments just as well as normal representational language does. I call this task “justifying validity”, and argue that the way to justify validity is—echoing the lesson of the previous section—to look to the use and purpose to which we put the language in the relevant domain. Because the language-derived coordination I posited as the use of moral discourse in the Chapter One is best undertaken in terms of giving and

asking for reasons, we are compelled to be coherent in our moral outlook; I propose a parallel explanation for modal discourse.

But because the semantics of expressivism is necessarily tied to the mental states that moral claims express, this kind response to the Frege-Geach problem will involve the expressivist in some unlikely commitments and complications. I will end this chapter by exploring some of these, and argue that they present a strong case for considering a non-expressive form of non-representationalism.

5.2 Part Two

Like expressivism, moral inferentialism rejects a representational account of meaning, but unlike expressivism it explains meaning in terms of the *inferential role moral concepts play*, instead of the *attitudes moral claims express*. According to the expressivist tradition, the meaning of sentences like “It is wrong to lie” are best understood in terms of the expression of some mental state, for example a disapproval of lying. We use such moral claims to express these evaluative states, and this use gives the claims their meaning. But those familiar with the Frege-Geach problem know that this alone won’t suffice: the use associated with a concept like “wrong” doesn’t explain its meaning because that meaning remains constant even in contexts when the use is absent. Drawing on Michael Williams’ recent work on inferential theories of meaning, I argue in **Chapter Four** that an appropriate understanding of the use of moral discourse—the facilitation of coordinated social behavior—explains the kind of inferences that are licensed by moral concepts. I offer a sketch of the inferential roles the moral ‘ought’ plays,

and argue that once we recognize the relevant inferential roles are meaning-constitutive, we will be in a position to solve the Frege-Geach problem in a new way, one that doesn't fall prey to the objections of the previous chapter.

In **Chapter Five** of the dissertation, I revisit the problem of objectivity, this time in an inferentialist rather than an expressivist key. The task for an inferential justification of moral objectivity is to make sense of our commitment to a particular inferential norm: For any genuine moral disagreement, it is necessarily the case that at least one disputant is in error. To make sense of this, I will consider Huw Price's arguments about the role that disagreement plays in general, and apply it to the case of moral disputes. Echoing the lessons of the previous three chapters, I argue that we can understand moral objectivity by showing that an inferential rule tying error to disagreement expedites the pragmatic aim of moral discourse.

The inferential approach has advantages over its expressivist counterpart: because it doesn't identify the source of moral disagreement in terms of the pragmatic clash between evaluative attitudes, it needn't account for those attitudes in a way that guarantees they'll always clash. I argue that it gives us a way to reconceptualize the problem of objectivity: I urge that instead of seeing the issue in terms of questions about the objectivity of truthmakers for moral claims, we should understand it in terms of questions about the sensibility of certain inferential norms. This approach also sheds light on the role disagreement plays in aesthetics and modal discourse, and gives us reasons to expect that the latter will have the resources to justify a robust objectivity in a way that the former will not.

If the arguments in this dissertation succeed, I will have carved out a space for a different approach to the old challenges metaethicists face. We do not need to understand morality in terms of the types of properties or states of affairs that make moral claims true; indeed refusing to do so will give us a metaethical perspective that sits comfortably with naturalism. But we also don't need to understand moral claims in terms of the evaluative attitudes they express. If we instead understand moral claims by looking at their inferential role, we'll find a way to make sense of moral reasoning and objectivity that secures the advantages of non-representationalism without the traditional baggage of expressivism.

PART ONE

CHAPTER TWO: OBJECTIVITY WITHOUT (HEAVYWEIGHT) REALISM

1. Overview

Moral expressivists have faced charges of relativism for a long time now. This makes sense; because expressivists do not explain the claims of moral discourse in terms of a realm of truthmakers—no range of moral properties or states of affairs to which true ethical claims must correspond—it's easy to see where the charge of relativism comes from. Such an account ties the meaning of these claims to the moral perspectives of those who hold them; since different people have different perspectives, and because the expressivist cannot appeal to a range of moral facts to adjudicate between perspectives, relativism looms.

A specific kind of response has become popular for the modern expressivist: to see the charge of relativism as a proposal for a first-order, engaged normative theory—a relativistic theory—and to dismiss it on these grounds as utterly inadequate. In this chapter, I will explain the charge of relativism the expressivist faces, and then review this response. It's not entirely satisfactory, I argue: it gives us the grounds for rejecting a certain kind of overly inclusive relativism, but doesn't give us the tools to justify a more robust kind of objectivism—one that gives us reasons to reject contradictions whenever we find them.

In this chapter, I'll argue that such objectivity can only start to be attained when we reflect on the *purpose* of moral discourse—the coordination of social behavior. I

recommend the same approach to one of expressivism's non-representational cousins, Amie Thomasson's modal normativism. After showing the consequences of such an approach, I suggest a general criterion for objectivity in non-representational accounts: A non-representational discourse X is objective if and only if the function of X is one that cannot be fulfilled where contradictions are accepted. I consider the consequences of such a criterion for moral expressivism, and suggest solutions to problems that may come up. Finally, I consider other potential worries for this approach, laying ground for the inferential account of objectivity I'll advance in Chapter 4.

2. The Charge of Relativism

The charge of relativism that expressivists must face generally falls out of two characteristic features of the expressivist thesis: the *semantic* and the *metaphysical*. Expressivism's semantic thesis is that ethical claims have a different function than traditional moral realists think they do—ethical claims do not aim to report on or describe some mind-independent moral domain; they instead function to express some *non-representational* mental state. There is a difference of opinion about what kind of mental state that is: in early work Simon Blackburn argues that moral claims express attitudes of approval and disapproval (1984); in later work, he speaks of being “tied to a tree” of patterns of inference governing these attitudes (1988); Allan Gibbard describes the state of mind expressed as an acceptance of norms which govern those emotional responses we typically consider ethical—anger, guilt, resentment, and so on (1990); Mark Timmons argues that we shouldn't hold out for the necessary and sufficient conditions that

characterize these states (1999). For the sake of brevity, I will use *approval* and *disapproval* as shorthand to stand in for any of these non-representational mental states.

A clear contemporary template for the semantic argument can be found with Jussi Suikkanen (2009), who begins a similar criticism of expressivism with a premise about the *correctness conditions* for various claims: “a necessary and sufficient condition for a statement... to have meaning is the existence of norms which determine when the sentence can be correctly used to make statements.” (368) Typical representational claims, for example, have meanings that are informed both by their truth- and sincerity-conditions; a sentence like “grass is green” is related to its truth-condition by the conventional connection between its utterance and the speaker’s belief that grass is green. Absent this belief and truth-condition, such a claim cannot receive a “positive semantic evaluation”; making the claim without the attendant belief amounts to violating a semantic norm.

The condition of semantic correctness for a moral claim is to be found in the biconditional:

Φ -ing is wrong if and only if I can now use the sentence ‘ Φ -ing is wrong’ to make a statement that receives a positive semantic evaluation. (Suikkanen, 369)

Claims in a representational language have a further semantic norm they must answer to: they are semantically appropriate “only and always when the relevant statements would be *true*, i.e. when they would correspond to the facts.” (Suikkanen, 369) The utterance “grass is purple” violates a semantic norm even when it expresses a sincerely held belief, since it’s false. But Suikkanen argues that because expressivists don’t make sense of moral claims by adverting to truth-conditions, the only standard for positive semantic evaluation of moral claims available to the expressivist is sincere assertion. This entails:

Φ -ing is wrong if and only if I can now sincerely assert the sentence ‘ Φ -ing is wrong’.

This is the conclusion reached by Frank Jackson and Philip Pettit:

Not only will the sentence ‘That is good’ express the attitude alleged [i.e., approval], it will be true just in case the attitude is present and false otherwise: it will in some sense, however broad, report the presence of that attitude. (1998: 242)

If this indeed is the correct account of moral truth from the expressivist perspective, then relativism obviously follows, inasmuch as two people can be sincere in their assertion of contradictory moral claims.

We might wonder why, on the expressivist view, moral statements don’t answer to further semantic appropriateness conditions in the way of normal representational statements; why we can’t also insist that

The sentence ‘ Φ -ing is wrong’ can be used to make a statement that receives a positive semantic evaluation if and only if it is both the case that ‘‘ Φ -ing is wrong’ is uttered sincerely *and if Φ -ing is wrong*.

It seems that this option is not available to the expressivist; because of her unwillingness to rely on any robust notion of moral facts or properties, the charge goes, she can make no sense of Φ -ing being wrong apart from the conditions in which it would be semantically appropriate to claim Φ -ing is wrong. This is an apparent consequence of the *metaphysical* stance of expressivism, a metaphysical quietism that refuses to posit a set of moral facts that function as *truthmakers* for moral claims.

Because this chapter is concerned with how non-representational accounts of various discourses can establish objectivity, and given that such accounts share this metaphysical quietism, this merits an important aside here about what exactly such quietism amounts to. It is crucial that we not conflate it with *anti-realism*, the view that there are no properties or facts of a given sort (e.g., no moral properties or facts; no modal properties or facts; etc.). Non-representationalism about ethics has traditionally (and

understandably) been associated with the latter. But this is unnecessarily radical. As we saw last chapter, armed with a deflationary notion of truth and a pleonastic account of properties, sophisticated non-representationalists attempt to lay claim to the full range of properties and truths that are traditionally understood as being found only in the jurisdiction of conventional realism about the subject.⁷

The nomenclature is a bit muddled at this point in the debate: advocates of this sort of position have used the label “irrealism” (Timmons 1999) or “quasi-realism” (Blackburn), but such terms seem to imply that the view being considered is an alternative to realism. I will argue that this is wrong; my aim in this dissertation is to show that non-representationalism can offer us a realist meta-ethics—“realist” in the only sense that matters. What is crucial here is the rejection of what might be called “heavyweight” realism, the view that the claims in question can only be true (or for that matter, objective) in virtue of truthmakers to which they must correspond. As we saw last chapter, the heavyweight approach must explain the nature of these truthmakers—how they are to be placed in a naturalistic worldview. The non-representational view under consideration here can avoid this requirement, so I will characterize the metaphysical thesis of expressivism as a variety of “lightweight realism”. The metaphysical and semantic theses of expressivism are deeply linked: because the expressivist denies that moral claims are in the game of corresponding to truthmakers altogether, she needn’t posit moral properties or facts or their truths to make sense of such claims.

Nonetheless (and returning to the contemporary charge of relativism), Suikkanen insists that the expressivist cannot rely on a notion of *robust*, non-deflationary truth to set

⁷ We might be tempted to take as the upshot of this that the debate has been resolved: since moral claims can be true, and since (barring dialetheism) the law of non-contradiction holds, objectivity comes for cheap. I’ll argue against this easy path in Chapter 5.

an additional standard for positive semantic evaluation. Doing so would amount to giving up the game, to admitting:

that there are moral facts and that our moral statements aim at matching them. This would mean that our moral statements are apt for being robustly true. If our moral statements could be true in this way, then it would seem to follow that we should be able to also believe them to be true. This would vindicate cognitivism. (373)

Eschewing “robust truth-conditions”, Suikkanen argues, the expressivist must make due with the “mere assertibility conditions” of sincerity as the truthmaker for her moral claims.

The charge of relativism comes in two easy steps, then. First we note that the semantic thesis of expressivism ties moral claims to mental states. And because expressivists are lightweight about moral properties, they cannot appeal to any further facts beyond mental states to non-arbitrarily adjudicate disputes. Given that such disputes take place, and that they are a consequence of the contrary moral attitudes different people hold sincerely, it would seem that moral relativism follows.

Obviously if this charge is correct, it undermines the expressivist’s claim to objectivity. It’s important to notice a troubling implication of this argument. It’s not simply that, because the expressivist does not rest her claim to objectivity in terms of the types of objects, facts, or properties to which moral claims correspond, she can have no hope of objectivity whatsoever. Instead, an objective expressivism is threatened by this charge because it undercuts any sort of *pressure towards resolution*: If the charge of relativism leveled at expressivists is correct, there is no reason for moral disputants locked in disagreement to persist. Since morality is relative, the argument implies, no resolution should be expected.

The charge of relativism shouldn't just be troubling for metaethical non-representationalists like expressivists. Non-representational accounts of other topics have their own versions of expressivism's semantic and metaphysical commitments, so one should expect that similar arguments would lead to relativistic conclusions within those domains. In the next brief section, I want to show how a parallel argument applies to expressivism's non-metaethical cousin, non-representationalism about modality. Later I'll argue that the responses to these charges run in parallel.

3. Modal Relativism?

Modal normativism is also vulnerable to charges of relativism—in this case, ontological relativism. Both non-representational views seem to make the truth of particular claims contingent on the commitments of those who utter them; the objection is that where moral expressivists tie the truth of claims in that domain to the moral commitments of the speaker, modal expressivists tie the truth of claims to the linguistic rules the speaker is committed to.

“Conventionalism,” Crawford Elder writes, “ultimately founders on its refusal to allow that any objects in the world possess mind-independent existences” (2004, p.20). Modal conventionalism is forefather to modal normativism—both views tie the truth of modal claims to the linguistic rules or conventions that govern our use of terms—and so the objection Elder presents is often taken to properly apply to normativism. The essence of an object is delineated by its modal properties—by what changes it can and cannot undergo, by what characteristics it must have. And so, according to critics, on the

expressivist account modal properties aren't really mind-independently "out there" in the world, so neither then are any objects. Consider the unraveling of a sweater:

The conventionalist suggestion is that the same physical processes of unraveling, with the same qualitative results, can result either in the persistence of a single object [i.e., the string of yarn which makes the sweater], or in the destruction of an object, *depending on what our conventions are*: it is because we have the concept 'sweater' that there is something which does not survive unraveling. But this is to say that we can, simply by agreement, and from a distance, make it the case that the material object I see before me today either will or will not exist tomorrow. (Elder 2004: 148, emphasis mine)

We can frame the argument in terms of semantic appropriateness. We can approach Elder's case with another biconditional:

It is true that the object survived unraveling if and only if I can now use the sentence 'the object survived unraveling' to make a statement that receives a positive semantic evaluation.

Since modal normativism ties the modal profile of an object to the linguistic conventions we have, it seems like the only conditions of appropriateness for such a claim will be conventional:

The sentence 'the object survived unraveling' can be used to make a statement that receives a positive semantic evaluation if and only if the conventions we have for the term 'object' entail that it can survive unraveling.

The argument is that if the modal claim 'an object cannot survive unraveling' is true, it will only be so in virtue of the fact that we have adopted the linguistic convention,

"Wherever the term 'object' applies, do not use this term where unraveling has occurred."

If such linguistic conventions we adopt are mind-dependent—as of course they are—and if modal properties in turn depend on the linguistic conventions we use, then ultimately, the existence of objects is mind-dependent. The upshot is that any claim regarding the existence of an object—or, more directly, any claim regarding a particular modal truth—can only be assessed with reference to the semantic conventions or rules that are being followed. And because modal normativism is lightweight in the same way expressivism

is, the normativist cannot point to any particular modal truthmaker that could non-arbitrarily decide which linguistic convention is the most appropriate. As with expressivism, the charge of relativism falls out of the semantic and metaphysical commitments of a lightweight non-representationalism.

And again, the non-representationalist must deal with a deeper issue: given a lightweight realism which rejects explanations by way of robust, non-minimal modal facts and properties, and a non-representationalism which holds that modal language does not function to report such facts, what is it that guarantees a pressure towards coherence? Why, when two people disagree about the modal profile of some object, must one of them be wrong?

4. The Standard Response

The modern day expressivist can avail herself of a response to these charges that has become the industry standard over the past couple of decades (see Blackburn 1998, 1999; Timmons and Horgan 2006); this offers insight into the modal charge. The charge of relativism, this response holds, misunderstands the semantic and metaphysical implications of the expressivist account. I will outline this response here, show how it applies to modal normativism, and in the next section explore a problem that comes with it.

Terry Horgan and Mark Timmons argue that charges of relativism misunderstand the primary role of ethical language. Ethical judgments might be understood in terms of the semantic correctness conditions of sincerity that Suikkanen suggests, but we shouldn't

assume that such an understanding commits us to a reductive equivalence of the following type:

(RE) I say, “X is good” truly iff I approve of X

While this is perhaps a good account of the semantic norms that govern my use of the term “good”, Timmons and Horgan stress that we shouldn’t conflate semantic appropriateness with truth. If I say, “X is good” *without* approving of X, then of course I am being deceptive, and whether the statement is true or not, my utterance doesn’t follow the relevant semantic norms. But because ethical claims do not function to descriptively report one’s ethical framework, we cannot exhaustively understand their truth conditions in terms of this framework.

To clarify the distinction between expressivism and subjectivism, Horgan and Timmons urge us to consider two possible perspectives from which to view predications of ethical truth: an ethically engaged and an ethically disengaged perspective. (RE) above is an account of correctness from an ethically disengaged perspective. Why “ethically disengaged”? Consider the following:

According to the norms of the Taliban, it is true that a woman should not talk back to her husband.

This is an explicitly relativized use of the term ‘true’, which can be glossed as:

It’s true that the ethical outlook of the Taliban includes ethical norms that require women to be obedient.

And while this gives us an understanding of the semantic norms that govern a committed Talibanista’s use of ethical terms, it certainly doesn’t affirm his ethical judgment, because our understanding of this state of affairs doesn’t motivate us to accept that it is true that women should be obedient.

Timmons and Horgan argue that our use of ethical language from an ethically engaged perspective is the more fundamental one. As Timmons argues,

Moral judgment... is *categorical* in nature. In other words, in the typical case, one has a certain kind of commitment—a categorical commitment—that serves to distinguish moral stance taking from having a preference, a taste, or a personal ideal. It is the sort of commitment that typically manifests itself in being disposed, for example, to assert and uphold one's stance over and against the conflicting stances of others, and so to categorically deny conflicting moral views rather than regard them as true or correct for those individuals who sincerely hold them... (Timmons 1999, 143)

It is because this first-order use of ethical language is categorical that we should consider it the primary use of ethical language.

Charges of relativism, they contend, come from a conflation of the engaged with the disengaged use of ethical language. The disengaged use of ethical language gives a descriptive account of how psychological states map onto ethical statements, but in doing so, separates us from the essentially evaluative and motivational role of first-order ethical language. So in presenting first-order, engaged ethical discourse as primary, Timmons and Horgan argue that they are showing a way out of relativism. The argument for relativism as an entailment of expressivism can only proceed when we conflate the conditions of semantic aptness with the *truthmakers* for moral claims. Because ethical expressivism is committed to a semantic thesis of non-representationalism, it's not in the business of hunting for truthmakers to which moral claims must correspond. (Timmons and Horgan 2006: 89)

But without any sort of heavyweight realist foundation, how can we hold out for a robust moral objectivity? Our opposition to the treatment of women by the Taliban is an engaged ethical opposition, but if there are no ethical truthmakers, how can we understand the claim that a woman should not talk back to her husband as being *objectively* wrong?

Obviously, the expressivist cannot adjudicate the dispute between us (denizens of Western society) and them (members of the Taliban) in terms of whether or not oppressing women has the in-the-world property of wrongness, so what sense is to be made of the dispute? To use the Paul Bloomfield's phrase (2003), what could give us the "moral high ground" in this situation?

Simon Blackburn (1999) offers simple answers to these questions: what's wrong with the moral commitments of the Talibanista is that they are unfair, patriarchal, mean-spirited and non-egalitarian. It's true that we can only make this judgment from within our own moral perspective, but this is to be expected. We have to start out from where we are, with the norms we presently use to justify our claims. Since the non-representationalist doesn't understand ethical discourse in terms of truthmakers to be tracked, they do not need to hold out for an algorithm to track them. We should not hold out for some principled stance from which we can evaluate all of our moral commitments with absolute authority.

This is not to say that we can rest on our laurels, and blithely rely on whatever moral standards we have. We cannot unjustifiably assume that we ourselves aren't subject to the kinds of mistakes in moral reasoning we see others making. We must remain vigilantly fallibilist, and, as Blackburn points out, we can do so without relying on a heavyweight foundation.

How can I make sense of fears of my own fallibility? Well, there are a number of things I admire: for instance, information, sensitivity, maturity, imagination, coherence. I know that other people show defects in these respects, and that these defects lead to bad opinions. But can I exempt myself from the same possibility? Of course not (that would be unpardonably smug). (Blackburn 1993, 318)

So we make our moral judgments with the evaluative tools we have available, and hold those tools to further standards still, but they are still our standards—inescapably so. To

use a metaphor popular among expressivists: we're all of us locked onto Neurath's boat, and can only make adjustments and changes as we travel with the materials at hand (Gibbard 1990, 314-15; Blackburn 1993, 318).

This response can be framed in terms of semantic correctness. Recall that Suikkanen makes his charge of relativism by introducing the concept of semantic "correctness conditions" for assertions; the representational claim that grass is green can be assessed in terms of sincerity conditions—one has violated a semantic norm if one makes this claim absent the sincere belief that grass is green—and also conditions of truth—even if the claim has been made sincerely, one has violated a semantic norm if the claim is not true. But then it seems we should also consider *justificatory* norms for assertions; if I make the above claim unjustifiably, I have violated a condition of assertion, even if I believe the claim and it is true. Blackburn's argument, then, can be read as pointing out the important role of justificatory norms for moral claims.

This sort of move should not come as a surprise to those making the charge of relativism, though. Jackson and Pettit, for example, consider briefly "more complex versions" of expressivism that "talk of expressing higher-level attitudes, or attitudes formed after due deliberation or negotiation with others" (240), and Suikkanen echoes Blackburn on the notion of moral improvement:

Our moral sensibilities are better the more coherent, consistent, and informed factually they are. If this is right, then there is a best possible set of attitudes I could have. This would be the set I would have if I made all possible improvements to my current actual attitudes. Thus, there exists the most coherent, consistent, and informed set of attitudes I could end up with. (Suikkanen, 381)

A version of expressivism that correlated truth in ethical claims to consistency with such a set of attitudes would still have worrying relativistic implications; plausibly, two agents

can be in such a state of deep disagreement that even a process of idealized improvement could not bring them to consensus.

Timmons and Horgan argue that the objectivity in expressivism need not be held hostage to the possibility of consensus between all idealized sets of attitudes, though. To think otherwise would again be to ignore the import of the semantic and metaphysical theses of expressivism: it would be to conflate the aptness conditions for a moral claim with the truthmakers for that claim. In making moral claims, we *express* attitudes of approval or disapproval. A sincere expression of attitude qualifies a claim as semantically apt; this does not mean that the claim functions to *report* the attitude. The expressivist dispels charges of relativism by reminding us that the attitude is not the truthmaker.

The same point applies for modal normativism. Though the linguistic conventions that are reflected in modal claims set the norms for correct assertibility, this doesn't entail that those conventions act as truthmakers for modal claims. We should remember that

the claims of metaphysics are not descriptions or statements of rules of language (made in a metalanguage) but rather are made in the object language, using the terms in question. Since the terms are *used* in making the claims of metaphysics, these claims are in that sense about the things, if any, our terms and concepts refer to—not about our terms or concepts themselves. (Thomasson 2007a: 152)

But still a nagging worry remains. If moral or modal claims don't aim to correspond to moral or modal facts, then what supplies the pressure towards resolution in disputes? In the next section, I will clarify this worry—what I believe the best argument against objectivity in expressivism would amount to—and will then look for grounds to defuse this argument.

5. A Practical Approach to Lightweight Objectivity

To bring the problem into focus, consider the disanalogy between the above counterarguments and a *prima facie* case for objectivity in a representational discourse. If two people disagree about some empirical matter—whether or not grass is green, say—we generally rest assured that at least one of them has made an error. Why? Of course, it might be the case that one of the parties to the disagreement is speaking insincerely; she says grass is not green even though she does not believe this to be so. Or it might be the case that even though both parties are expressing sincerely held beliefs, one of them is not justified in her belief. But even if we find that both parties are sincere and justified in their contradictory assertions, we nevertheless assume someone is speaking incorrectly. This is because our representational discourse aims to represent the world as it is—to correspond, as it were, to the greenness that acts as a truthmaker for the claim. Contradictory claims about the world can only serve to undermine this function because the world is unified and coherent, and so we have a reliable reason to reject all contradictions.

A lightweight account can't make use of such an explanation of objectivity. Moral discourse, the expressivist holds, does not function to represent a mind-independent moral domain, so she cannot bank on the unified nature of such a domain to explain why one party must be in error in cases of disagreement. And this seems to be at the heart of the worry: a convincing objectivity—the kind we might plausibly hope for in empirical matters—can only be secured if we have a confidence that there is some justifiable pressure to move towards concurrence in cases of disagreement.

Unfortunately, two of the semantic norms we've discussed—sincerity and justification—are not plausible candidates for giving us this pressure. This is because, in

the first case, we cannot rest assured that all instances of disagreement will reduce to a matter of one party speaking insincerely. Nor can we (pace Ideal Observer theories of morality) assume that all moral disagreements will boil down to at least one party making unjustified assertions. If these norms are not up to the task of establishing objectivity, how can the expressivist guarantee that disagreement entails error? This is, after all, the operating assumption that drives the pressure towards resolution.

Timmons (2004) argues that moral discourse has what he calls ‘objective pretensions’—when engaged in ethical discussion, we bring to the table the concepts of genuine ethical disagreement, ethical reasoning, and non-relative ethical truth—and it is a condition of correct use of moral language that these pretensions be accounted for. It’s a conceptual truth about morality that it has categorical, motivational force, and for it to have that sort of force, we can’t believe that there are alternate ethical truths. So, when we use ethical language in its primary, ethically engaged manner, we do so in virtue of paying heed to these constraints.

This seems to put the cart before the horse, though. We aren’t assured of truths full stop merely by the presence of putative conceptual truths. Perhaps these “objective pretensions” are just that—pretensions. Imagine for example a population in which it is a conceptual truth about simple judgments of taste—claims about which flavor of ice cream is best, say—that they are objective: if one person claims chocolate ice cream is the best, and another claims vanilla is the best, at least one of these people must be in error⁸. When members of such a population engage in such taste-talk, they bring to bear certain standards of use that do not allow for disagreement without error. Now, to give an

⁸ It’s not too hard to imagine such a population: just disagree with a seven-year-old about what ice cream is best.

adequate account of taste discourse, you'd need to explain this feature but note that this explanation on its own wouldn't give us any justification in believing taste claims are objective. The point is that the semantic features of a discourse are not on their own sufficient to justify these conclusions about objectivity, since these features might themselves be nonsensical. Timmons is right that the place to start looking is at the structure of moral discourse, but we must look deeper.

I want to argue that the place to look isn't the domain of facts to which moral claims aim to refer. If it were, the search for objectivity would be over before it began for *any* lightweight account. We should instead look to the *aim* of moral discourse. I argued above that the pressure towards concurrence that is the hallmark of a robust objectivity for a representational discourse makes sense, given that the aim of that discourse—accurate representation of a unified, coherent world—cannot be met if we have no pressure to resolve the tension between contradictory claims. But this is only one way to make sense of such a pressure within a discourse; there is more than one way to have objectivity.

What in the function of ethical discourse could give this guarantee? If there are no ethical truthmakers in the world, what would be so wrong about accepting that two contradictory ethical claims are both true? People might in fact speak as if there were objective moral truths, but is there any expressivist justification for this way of talking? Among modern ethical expressivists, there is a type of answer that one finds repeated.

Timmons argues:

[Ethical beliefs] are involved in reasoned action guidance. The sort of inconsistency involved with ethical beliefs can be understood in terms of possibility of action: *inconsistent ethical beliefs cannot coherently guide our behavior*. And, in failing to guide our behavior, they frustrate one of our main practical goals: effective action in the world. (Timmons, 173, emphasis mine)

We can see then, why the need might arise to treat our normative commitments—especially our ethical commitment—as categorical and non-relative. When we have ethical disagreements, or otherwise engage in ethical inquiry, it won't do to become relativists about the truth of the matter; such a move will leave unanswered that most important question, "What do we do now?" This rough sketch reflects a central argument of my dissertation: where heavyweight realists can rely on truthmakers to explain the features of moral discourse, non-representationalists must rely on the function of the discourse itself to make sense of these features. (But the fight for moral objectivity isn't over yet—see section 5 below for complications.)

Can we generalize this approach to explain objectivity in non-representational accounts of other discourses? Only if the domain in question has the right sort of pragmatic features; as I argued above, lightweight accounts cannot justify claims of objectivity by reference to a realm of facts or properties to be objectively described, and so their only recourse is to be found in the function of the discourse in question. From this, we can infer a master criterion that any non-representational account must meet in order to justify objectivity:

(Crit) A non-representational discourse *X* can establish objectivity if and only if the practical goal of *X* is one that could not be fulfilled without pressure towards coherence.

For there to be pressure towards coherence within a discourse, there must be some pragmatic feature of that discourse that explains why we cannot have disagreement without error.

Consider for example the function of modal discourse. Of course, it doesn't function to guide social behavior in the way that moral discourse does, but it does guide a particular aspect of social behavior—specifically, linguistic behavior. Robert Brandom

(2007) gives a compelling account of the use of modal discourse: Our ability to describe anything in the world—simply talking about everyday objects, properties, and events—“already presupposes grasp of the kind of properties and relations made explicit by modal vocabulary.” (7) We use modal propositions to make explicit the kind of normative rules to which we must conform in any “autonomous discursive practice”—that is, in any substantial language game. These are rules that we must have in order to count as asserting anything at all. We use modal discourse to orchestrate communication; it is a tool for guiding us in the common meanings we give our terms.

So what comes to the rescue here isn't some fact of the matter, or some algorithm for deciding between disputants, but rather a practical commitment to using words in the same way. It is built into the act of engaging in language that we try to use language consistently—that our terms refer to same things whenever they're used, whoever uses them; modal language functions to make this use explicit. Because we are engaged in conversation there is pressure to mean the same things with our words (or have some way to signal that we are using a homophone with a different meaning). If we disagree, we must not simply let things remain unresolved—not if we are seriously committed to communicating—because inconsistent modal beliefs frustrate the shared attempt at communication. If we do not accept objectivity as a guiding principle, we cannot communicate, and conveying constitutive rules for the terms we use to communicate is what modal language is for.

6. Further Concerns

The approach I argue for here is one in which the inconsistency of those claims that fuel disagreements is best understood in terms of the practical consequences of such disagreements. For an expressivist, these involve claims that express sensibilities that would dispose the disputants to incompatible behavior. If we accept that both disputants speak truly—that is, if we embrace a relativistic account of their dispute—we undermine the coordinating function of ethical discourse. “Relativism thus threatens the social function of morality—which perhaps explains why it is such a disturbing thesis, at least in extreme forms. “ (Price 1988: 179)

This is an approach with precedent among philosophers friendly to the expressivist project. Huw Price argues:

What matters is a certain correlation between utterance and action in general: roughly, that within a speech community utterances of a given sentence tend to be correlated to mental states with similar causal or functional roles in the determination of behavior. (Price 1988:152)

If such a stable correlation is established, we can start to see why those committed to conflicting claims should try to resolve disputes with argument—such claims express mental states that bring us into practical conflict. For example, in his account of clashing beliefs, Paul Horwich points out:

An essential property of our faculty of belief—its *raison d’être*—is the role it plays in determining how we are inclined to act. And the conflict associated with contradictory beliefs consists in their potential, through inference, to engender conflicting desires and decisions. If I disagree with you about the truth of some empirical proposition, <T>, then that can easily result (via theoretical reasoning and given other premises) in our disagreeing about the truth of some more directly action-guiding belief, <If A is done then X will occur>. And if we both want X to occur then one of us will, on that account, be in favor of A being done, and the other won’t. We might even come to blows! So one can see how divergent empirical beliefs might correlate with a practical tension. (Horwich 2010: 183)

The success of this expressivist approach to objectivity, then, depends on establishing that moral claims express mental states that exhibit the same “correlation between utterance and action” as beliefs do—a correlation that produces some sort of practical tension when those utterances are incompatible. The expressivist must account for these mental states in such a way that they explain this practical tension in every case of genuine moral disagreement.

And this is where the problems start to creep in for the expressivist, because not just any mental state will be adequate to the task. Consider for example, the state of being in pain. Say we come up with a predicate, “throbbly”, that we use to express the pain that a headache is giving us. “It’s throbbly in here,” you say, and because I have no headache, I say, “It’s not.” (see Dreier 2009: 102-110) It doesn’t seem like this is a real disagreement:

[W]hen I say that it is throbbly in here and you reply that it isn’t the least bit throbbly, all that is happening is that I have a headache, and you do not, and we both know this... and there is no intelligible sense in which we disagree. (*ibid.*, 103)

Now the crucial question is this: are the mental states that moral claims function to express more like Horwich’s beliefs or more like pains? Because “one precondition, maybe the only one”, for moral disagreements to actually count as disagreements and not faux-disagreements, “is that the state of mind expressed... is the kind of thing that we can disagree with.” (*ibid.*)

Headaches aren’t the only sorts of mental states that fail to produce genuine conflict. Consider desires, for example. Say we attach the predicate “yummy” to certain objects in order to express our desire for them. “Those cookies are yummy,” you say, and I demur because I’m very full—“Those cookies aren’t yummy at all.” In what sense can this be considered a disagreement? There’s no practical conflict here to buttress the notion

of genuine inconsistency between our claims. It's not even clear that you're contradicting yourself if, given these standards for using the predicate, you say, "Those cookies are yummy but they're not", since there may be something unfortunate—but not illogical!—in having desires that pull in different directions.

Allan Gibbard offers a partial solution to this problem by suggesting the notion of a commitment to plans as the right sort of candidate for the mental states that moral claims express. If, for example, a moral sentence like, "Murder is wrong" expresses a commitment to a kind of plan to refrain from murder and to keeping others from murdering, then we can see how someone who makes such a claim is in practical conflict with someone who thinks that murder isn't wrong. If I desire a cookie, but also don't desire it, I can have the cookie or not without being vulnerable to charges of being incoherent. But if I am both committed to a plan to refrain from murder, and at the same time committed to a plan to murder, the goal of one of these commitments must be frustrated. They have the right sort of practical clash.

That's the easy case. But not all ethical commitments can be cashed out in terms of commitments that have such obvious behavioral consequences:

There is only one dose of painkiller left and we each have a headache. You judge that people ought to forswear pharmaceutical relief from pain and tough it out, while I think you are mistaken—people ought to embrace the pain relief offered by medically tested drugs. Now my aims are met to the extent that you would manage to act on yours. There is no practical clash. And indeed the clash would come precisely if we *agreed* that each of us ought to try and grab the Ibuprofen. Why, then, do we disagree in the first case and agree in the second?" (Dreier 209: 105)

The challenge, it seems, is to supply a mental state that will explain why this apparent moral clash about what's to be done with the Ibuprofen actually qualifies as a disagreement—a mental state that correlates to behavior in such a way that there is some practical clash to be resolved.

Notice, though, that solving the Ibuprofen challenge alone won't do. There are many ways in which we have moral disagreements that don't obviously boil down to practical clashes. We disagree over ethical questions about past events ("Were the bombings of Hiroshima and Nagasaki justified by the quick end they put to the war?"), about practically impossible counterfactuals ("Should Superman really be spending his time stopping bank robberies instead of, say, supplying the world with cheap energy using his heat vision?"), and so on, none of which bears in any obvious way on mental states that should produce clashing behavior. The challenge for the expressivist is to supply an account of mental states that will do the trick.

In Part Two of this dissertation, I will argue that an inferentialist approach to metaethical non-representationalism has advantages over expressivism in accommodating moral discourse. I will return in Chapter 5 to the above challenges, and argue that an inferential approach holds special promise for making sense of moral objectivity. The simplest way to make sense of moral disagreements in contexts like the Ibuprofen dispute or disputes about counterfactual or past behavior doesn't require an identification of mental states that produces the right sort of practical clash. Instead, the inferentialist must identify and explain the relevant inferential norm that informs our moral disagreements—a norm whose function doesn't require a particular sensibility in order to operate. As we will see, this is a norm we would expect moral discourse to have, if it is to play its characteristic role of coordinating social behavior. The advantage of the inferentialist approaches to accommodation—both of objectivity, and as we will see in Chapters 3 and 4, of meaning—is that their plausibility doesn't depend on an identification of specific mental states.

CHAPTER THREE: THE FREGE-GEACH PROBLEM: THE METHODS AND LIMITS OF AN EXPRESSIVE APPROACH

1. Overview

Using Frege's distinction between content and use, Peter Geach raised a problem for non-representational accounts of morality by bringing up questions about the meaning of non-representational claims when they're embedded within conditional propositions, and about how such claims can be validly used in logical arguments. In this chapter, I will explain the *Frege-Geach problem* for moral expressivism, show that like the problem of objectivity, it generalizes to non-representational accounts in other domains, and indicate two tasks an non-representational theory of any ilk must accomplish in order to adequately respond to it. Simon Blackburn proposes a solution to the Frege-Geach problem for moral expressivism; I'll use it as a model of an attempt to accomplish these tasks and consider a parallel approach for an explicitly expressivist version of modal non-representationalism.

In the last half of this chapter, I'll point out the complications such expressivist approaches face, if we take seriously the challenge of explaining meaning in terms of use. I'll conclude that given these problems, exploring a non-expressive approach—one that doesn't rely on the inferential characteristics of the mental states moral claims putatively function to express—is an attractive route for the non-representationalist grappling with the Frege-Geach problem.

2. The Frege-Geach Problem

Expressivists give an account of the meaning of particular claims in terms of the commitments they are used to express. Moral expressivists argue that sentences like:

It is wrong to lie.

are best understood as expressions of some mental state, like a disapproval of lying. The idea is that we can understand the *meaning* of sentences like the above by understanding their *use* when they are asserted.

Peter Geach (1965) argues that those who hold that the use of term or phrase gives its meaning have things backward—“a thought may occur now asserted, now unasserted, without change of content”(Geach, 457)—and so it seems we must first have an understanding of meaning before we can give an account of use. This, he argues, is the only way to make sense of conditionals with embedded moral clauses, such as:

If it is wrong to lie, it is wrong to get your brother to lie.

In asserting this conditional, I’m neither expressing my disapproval of lying or of getting your brother to lie; what I am asserting is that a relationship holds between these two putative facts. The Frege-Geach problem was originally presented by way of conditionals, but of course we find examples of such “force-stripping” clauses not only in conditionals, but also in moral expressions like:

Is lying wrong?
It’s not the case that lying is wrong.
Lying is either harmless or it’s wrong.

If the meaning of the word “wrong” in these can’t be understood in terms of the speaker’s disapproval of lying, the expressivist must give some other account of meaning.

Notice that the problem is just as acute for other non-representationalists as it is for the moral⁹. For example, a modal expressivist would hold that claims of necessity and possibility express a commitment or intention to use the constitutive semantic rules for the terms used. An important note: The modal non-representationalism we considered in the last chapter is Amie Thomasson’s modal normativism. In this chapter I want to treat the Frege-Geach problem for an expressivist version of modal non-representationalism—that is, a version that, like moral expressivism, explains the meaning of claims in terms of the mental states they express. As we’ll see, this will help us keep track of the kind of semantic commitments an expressivist treatment makes.

So a modal expressivist would analyze the proposition “Necessarily, all bachelors are men” as an expression of our commitment to use the semantic rule “Apply ‘bachelor’ only where ‘man’ applies.” We can see then that the Frege-Geach problem also arises for modal expressivist. For example, the proposition:

If dogs are necessarily reptiles, then Rascal is a reptile.

is obviously not an expression of commitment to use a semantic rule that would have us use the term “reptile” wherever the term “dog” applies; we can sincerely assert this claim without being committed to such a rule.

If we try to sidestep this argument by claiming that the meaning of modal or moral propositions in such force-stripping conditionals is in fact different from the meaning they have when presented bare, an obvious trap awaits. This sort of response leaves us guilty of *equivocation* in modus ponens arguments like the following:

1. If he did it, he should be punished.
2. He did it.

⁹ In fact, Geach’s target in (1965) wasn’t limited to moral non-representationalism specifically, but to any “anti-descriptive”(463) account of a discourse that grounds meaning in terms of use instead of reference.

3. He should be punished.

For this argument to be valid, the meaning of “he should be punished” must be the same in (1) and (3). If it isn’t, the argument is invalid because of the equivocation involved, as in the following argument¹⁰:

4. If he works at the bank, his feet will get wet.

5. He works at the bank.

6. His feet will get wet.

This argument looks valid, but is not, if the meaning of “bank” in (4) is “side of a river”, but the meaning in (5) is “financial institution”. If the expressivist contends that moral or modal propositions have different meanings in force-stripping contexts, the Frege-Geach problem charges expressivism with the same sort of equivocation.

To answer these charges, the non-representationalist must accomplish two tasks:

Refuting equivocation: an account must be given of the content of the claims in question that doesn’t leave them vulnerable to the above charge of equivocation. It must be shown that the meaning of a claim remains constant even in when its typical use isn’t in force.

One way for the moral expressivist to refute equivocation is to give a formalization of moral language that tracks the way we use it in reasoning to conclusions, and keep the content of moral propositions the same whether or not they’re embedded within a conditional. This is the approach Simon Blackburn (1984) and Alan Gibbard (2000) take; it’s accomplished by positing mental states that moral claims express whether embedded or not.

But once this task is accomplished, we can’t just stipulate that the formalization is an adequate account of the meaning of moral terms. We must also account for how we use moral claims in complicated reasoning.

¹⁰ Taken from Blackburn (1984, 190).

Justifying validity: We have to explain why it is that the non-representational way we use the language in question is such that it can produce valid and invalid arguments just as well as normal descriptive language does.

To justify validity, we need to see why a commitment to (1) and (2) above compel us to accept (3) also. As we will see, the way to justify validity for an expressivist account is similar to the path I took last chapter, where I considered how an expressivist can justify claims to objectivity: we must identify the impact that the pragmatic aim of the language in the relevant domain has.

3. Refuting Equivocation

In response to the Frege-Geach problem, Simon Blackburn offers an expressivist formalization of a language for translating moral claims. I will treat Blackburn's formalization as a proof of concept. I'll also offer a parallel solution to the modal version of the problem; examining what they accomplish and fail to accomplish will show us just what will be involved in a satisfying solution to the Frege-Geach problem.

Blackburn begins by suggesting an explicitly expressivist language, one that contains no ethical predicates but instead "wears its expressive commitments on its sleeve" (Blackburn 1984, 193). He suggests B! and H! operators, which are used to express disapproval or approval, respectively. So an assertion like:

Stealing is wrong.

can be formalized without recourse to moral predicates as:

B!(stealing)

This formalization indicates the disapproval evinced by the moral predicate “wrong”; it gives us a way to model the mental state that use of such a predicate expresses. Blackburn argues that we can understand this in terms of the moral *sensibility* the speaker expresses, where a moral sensibility is a function from some factual input—the act of stealing, for example—to an evaluative output—here, disapproval.

This notion of a sensibility is crucial for Blackburn’s account of the content of moral propositions, because we don’t only approve and disapprove of acts, but also of sensibilities themselves. To represent this sort of higher-order attitude, Blackburn introduces vertical lines into his formalization, so that:

H!|B!(lying)|

indicates an attitude of approval towards disapproving of lying.

Sensibilities function to produce approval or disapproval as output, then, but they aren’t limited to acts or other sensibilities as inputs. They can also take as input *certain relations* between different sensibilities, or sensibilities and other non-moral facts, and Blackburn argues this is an important function:

For one of the main features affecting the desirability of the world we live in is the way other people behave, and the way other people behave is largely a function of their sensibility... And amongst the features of sensibilities which matter are... the interactions between them. For instance, a sensibility which *pairs* an attitude of disapproval towards telling lies, and an attitude of calm or approval towards getting your little brother to tell lies, would not meet my endorsement. (ibid, 192)

Blackburn suggests the use of a semi-colon to indicate the presence of a relation. This gives us a way to understand the conditional:

If it is wrong to lie, it is wrong to get your brother to lie.

with the formalization:

H!|B!(lying);B!(getting your brother to lie)|

Note that the content of “it is wrong to lie” is given as an expression of the same sensibility *whether it’s embedded or not*. The content therefore stays same even in force-stripping contexts; this is how Blackburn refutes the charge of equivocation. The key move here is to give an account of the meaning of simple moral claims that comes from the mental states they express. The meaning of complex claims with moral clauses is a function of the meaning of its parts; we get a compositional semantics in which complex moral claims express mental states that are a function of the mental states that would normally be expressed by their clauses.

A similar approach to refuting equivocation is available for the modal expressivist. Remember that modal non-representationalists hold that statements like “Water is necessarily H₂O” and “Bachelors are necessarily male” do not purport to *describe* some non-modal facts. Modal talk doesn’t stand in need of modal truth-makers, they claim, because the purpose of modal talk isn’t to *describe*, but is instead to *express* our commitment to analytic principles of individuation (cf. Sidelle 1989) or rules concerning the application and co-application conditions for our terms (cf. Thomasson 2007). Robert Brandom (2008) argues that the purpose of modal discourse is to make explicit the sort of commitments and entitlements we need when using non-modal discourse; this idea of constraints and permissions¹¹ is important in my treatment of the modal Frege-Geach problem.

Modal propositions involving necessity and possibility function to make explicit the kinds of linguistic constraints and permissions we accept. For example:

Necessarily, cats are animals.

¹¹ Brandom speaks of “commitments” and “entitlements”. I’ll eschew this terminology, as I wish to continue to use the term “commitment” in the special sense above.

is an object-language statement that expresses our commitment to the following meta-linguistic constraint:

Wherever “cat” is used to refer, one can use “animal” to refer to the same thing.
“Cat” cannot be used where “animal” doesn’t refer.

Furthermore:

It’s possible for water to be ice.

Expresses our acceptance of the following permission, which tells us:

It is sometimes permissible to use “ice” where “water” refers.

I propose we mirror Blackburn’s formalization of moral language, and consider a modal language that wears its constraints and permissions on its sleeves. Let the operator C! indicate a commitment to a constraint, and P! a commitment to an permission. A colon can be used between terms or phrases to indicate direction of fit, so that:

Necessarily, cats are animals.

is formalized as:

C!(‘cat’:‘animal’)

and:

It’s possible for water to be ice.

is formalized as:

P!(‘water’:‘ice’)

We can make sense of modal claims embedded in conditionals by stealing the notion of a sensibility from Blackburn. Of course, a moral and modal sensibility will have quite different functions: one functions to indicates approval or disapproval of acts and other sensibilities; the other functions to give us license to make such inferences from one

modal proposition to another, or from a modal proposition to a non-modal one¹². On this account, a modal claim functions to express a state of commitment to use terms in a certain way, so I'll call the modal version of a sensibility a *linguistic sensibility*.

Conditionals such as:

If it's necessary that cats are mammals, then its possible Zoe is a cat.

can be formalized as:

C!|C!(‘cat’: ‘mammal’); P!(‘Zoe’: ‘cat’)|

This indicates a linguistic sensibility—a commitment to a constraint that takes us from another constraint to a permission; the modal claim is read as an expression of commitment to a linguistic rule (the one with the widest scope) constraining the speaker to being permitted (to defeasibly use ‘cat’ wherever ‘Zoe’ refers) on the occasion of accepting the original constraint (to use ‘mammal’ wherever ‘cat’ refers).

Again, we're in a position to refute the charge of equivocation, because the meaning of a modal sentences like “Cats are necessarily mammals” is given by the same mental state, a linguistic sensibility that remains the same whether it is embedded or not. As in the moral treatment, complex modal claims like the conditional above express a mental state that is a function of the mental states that would be expressed by its parts.

5. Justifying Validity

We're now in a position to think seriously about how such an account of the meaning of moral or modal claims can explain the validity of arguments that involve such claims. And this is where the problems will start, because “it is not enough just to do

¹² Compare Ryle (1950) on the notion of inference licenses.

compositional semantics; expressivists need a compositional semantics which correctly predicts why complex sentences have the properties they do.” (Schroeder, 22)

Representational approaches to morality or modality will have an easy time accounting for the validity of moral and modal arguments. Because such approaches start with the truth-conditions of moral and modal claims, there’s no problem explaining why the following is valid:

7. If it is wrong to lie, it is wrong to get your brother to lie.
8. It is wrong to lie.
9. It is wrong to get your brother to lie.

The meaning of (7) is a function of the truth-conditions of its parts: it can only be true if its antecedent is false, or if both the antecedent and consequent are true. Since (8) holds that the antecedent is true, the conclusion of the argument follows. Validity is just a matter of the truth of the premises guaranteeing the truth of the conclusion, so there’s nothing mysterious going on here. But because their metaphysical quietism bars them from using the truth-conditions of moral claims as *explanans*, expressivists can’t make use of the same tools that the representationalist enjoys. As with the problem of objectivity, so too the problem of meaning: the non-representationalist has to supply pragmatic considerations to do the job that truth-conditions do for the representationalist.

Blackburn proposes that we can understand the content of the above argument with the formalization:

- 7'. $H!|B!(lying);B!(getting\ your\ brother\ to\ lie)|$
- 8'. $B!(lying)$
- 9'. $B!(getting\ your\ brother\ to\ lie)$

Why do the premises lead to the conclusion here? Consider the attitudes expressed in the premises:

Disapproval of lying, and approval of making (disapproval of getting [your] brother to lie) follow upon (disapproval of lying). Anyone holding this pair must hold the consequential disapproval: he is committed to disapproving of getting [your] brother to lie, for if he does not his attitudes clash. (Blackburn, 195)

But the ‘clash’ between attitudes that Blackburn proposes can’t be understood in the manner the representationalist would have it: it can’t simply be a clash between beliefs that can’t both be true. Instead, Blackburn explains the clash pragmatically: A person holding both attitudes would have

a fractured sensibility which cannot itself be an object of approval. The ‘cannot’ here follows not (as a realist explanation would have it) because such a sensibility must be out of line with the moral facts it is trying to describe, but because such a sensibility cannot fulfill the *practical purposes for which we evaluate things*. (Blackburn 1984, 195, emphasis mine)

Remember that, according to the analysis of last chapter, moral discourse aims to guide us into complex coordinated behavior. If our moral claims expressed sensibilities that didn’t answer to a pressure towards consistency, they’d be of no use in guiding our behavior.

Now, one might accept the notion that having some moral sensibilities commits one to having others, but is this a *logical* commitment? If I approve of disapproval of lying, but don’t myself actually disapprove of lying, it seems that having such a fractured sensibility would be problematic—maybe hypocritical—but no more *illogical* than wanting a piece of cookie, but wishing you didn’t want it. (Schueler 1988, 500) Such conflicts are familiar to us all, and we don’t normally think that they involve any sort of logical *contradiction*. We can understand why we must honor the demands of consistency in our representations of facts—the goal of belief is to accurately map the way the world is, and if we have inconsistent beliefs, that goal cannot be met. This isn’t clearly the case with the sensibilities involving approval or disapproval, though. The problem is that not all mental states produce contradictions in the way that beliefs do: believing p and $\sim p$

entails some sort of inconsistency, but wondering p and $\sim p$, for example, does not. And disapproving of p and $\sim p$ seems more like the latter than the former.

This doesn't mean we have to abandon Blackburn altogether, but it does suggest an improvement: moral sensibilities aren't best understood in terms of approval and disapproval, but in terms of *planning* or *being committed*.¹³ Language-derived coordination is best undertaken in terms of giving and asking for reasons for our commitments, and so we are compelled to be coherent in our moral outlook. Coherence itself is part of the moral language game—we simply aren't engaging in moral discourse inasmuch as we're not making ourselves available to good reasons for our commitments—and so coherence is a goal in the moral domain just as much as it is in the doxastic domain. If I am both committed to getting my brother to lie and committed to not being committed getting my brother to lie, the goal of one of these commitments must be frustrated. Beliefs have as their goal accurate representation of the world; moral evaluations have as their goal a commitment to attitudes that effect coordination. Both require consistency for these goals to have any hope of being met. This account of coherence in our commitments is just what Blackburn needs to justify validity: suffering

¹³ There's an ambiguity in the term "commitment" we should clarify. Here the word is not meant to indicate something like "promise"—promising, like approval and disapproval, isn't something that obviously produces *logical contradictions*. I can approve of something and disapprove of myself for feeling that way, and I can promise something after promising I wouldn't, and while both of these indicate an inconsistency on my part, the inconsistency doesn't have the same logical implication in, say, believing it's raining and believing it's not raining. See Schueler (1988), for more on this point. Instead, I'm using the term "commitment" in the same way it's used when we speak of committing to a *decision*.

I treat the issue at greater length elsewhere in my dissertation, but we should for notice that this marks a break with Blackburn's schema in *Spreading the Word*. If his formalization only codes for approval and disapproval, it will not give us an adequate account of validity, because it won't be able to make sense of the *logical contradictions* we find in statements like "Killing is bad but it's not bad". To get around this problem, we could speak instead of being committed or deciding to disapprove of killing, and also being committed or deciding to not disapprove of killing. See Gibbard's *Wise Choices, Apt Feelings*, Ch. 7 for an interesting attempt to account for the moral emotions, which are governed by norms we express when making moral judgments. See also his *Thinking How to Live*, Ch. 3 and Ch. 4, for more on the role of contradiction and negation. Notice that he uses *planning* where I use *being committed*.

from a fractured sensibility is more than just a matter of having conflicting objects of approval and disapproval; it's a matter of having conflicting *commitments*.

The takeaway from the above considerations suggests how we should expect to solve the Frege-Geach problem for expressivism of any flavor:

1. Identify the mental states that the claims in question function to express.
2. Account for the meaning of those claims in terms of those mental states, so that even in embedded contexts, moral clauses express the same states and so have the same meanings.
3. Make sure you have a pragmatic account of the discourse that makes sense of the need for coherence, so that we have a way of understanding of why arguments with particular structures can be valid.

This last step is crucial, and it sheds light on the requirements for a good solution to the problem.

Consider again the Frege-Geach problem for modal expressivism. The proposal is that we understand *modus ponens* arguments like the following:

- 7'. If it's necessary that ice is water, then it's necessary that this swan sculpture is water.
- 8'. It's necessary that ice is water.
- 9'. It's necessary that this swan sculpture is water.

with a formalization like this:

- 10'. C!|C!(‘ice’: ‘water’); C!(‘this swan sculpture’: ‘water’)|
- 11'. C!(‘ice’: ‘water’)|
- 12'. C!(‘this swan sculpture’: ‘water’)

The idea is that if (10') above expresses a linguistic sensibility to a higher-order constraint, namely to the constraint of using ‘water’ wherever ‘this swan sculpture’ refers

on the occasion of accepting the constraint of using ‘ice’ wherever ‘water’ refers. If I combine this higher order constraint along with the constraint that I use ‘male’ wherever ‘bachelor’ refers (expressed by (11’) above), I must also accept the constraint expressed in the conclusion of the argument, on pain of inconsistency. And we can once again explain the logic of the argument with pragmatic considerations.

Why do we need to commit to these linguistic constraints and entailments?

Brandom argues for what he calls the “Modal Kant-Sellars” thesis, which holds that when speaking non-modally—when simply talking about everyday objects, properties, and events—we’re already implicitly committed to the rules which modal discourse makes explicit. To engage in meaningful discourse, we must understand one another, and this involves holding others and being held responsible for the implications and entailments of our claims. Notice that the use of modal language shares this feature with the use of moral language—the goal of both types of language demands coherence among language users. In the moral domain, this demand must be met in order to ensure the goal of coordinated social activity; in the modal domain, this demand must be met in order to ensure the goal of shared meaning of terms.

6. Problems Creep In

Non-representationalism about morality has been focused on the nature of the attitudes moral claims express since Ayer’s emotivism, a preoccupation that was reinvigorated by Blackburn’s quasi-realist treatment of the issue. And though Gibbard’s norm expressivism focused on issues of rationality (the rationality of blaming, of guilt, etc), it never broke the basic mold set by Ayer: issues like the Frege-Geach problem were to be solved by

reflecting on the roles played by the mental states expressed. In the remaining chapters of my dissertation, I want to stress that it doesn't have to be this way, and that in fact there are advantages to solving the traditional semantic problems that plague expressivism without worrying ourselves with mental states. To motivate this, though, we should first look at the burdens invoked by this classical expressivist tactic.

Let's start with problems facing my own work: I've defended Simon Blackburn's higher order account of expressive semantics, in which the logical properties of complex sentences with normative parts are understood in terms of a normative attitude—an evaluative sensibility—towards potential combinations of the attitudes that would be expressed by those parts. We saw that such a defense is possible, and therefore that Blackburn's higher order account is an adequate response to the embedding problem, only given a couple of conditions. The first is that we abandon the attitudes of disapproval and approval as those which sustain the sort of semantic relationships we find among complex normative sentences: the problem with disapproval and approval, as Schueler pointed out and as Blackburn later conceded (2007), is that they don't seem to produce logical inconsistencies; it may be unfortunate if I both approve of cheating on one's taxes, but also disapprove of that approval, but such capriciousness doesn't amount to a genuine logical contradiction.

In response to this pressure, I instead suggested that we understand the mental states expressed as sensibilities involving intentions, or commitments, in the hopes that such states would work where approval and disapproval didn't. To illustrate, consider once again the argument:

If lying is wrong, getting your little brother to lie is also wrong.
Lying is wrong.

Therefore, getting your little brother to lie is wrong. I argued that the validity of this *modus ponens* argument might be justified if we understand that to be committed to not being committed to a getting my brother to lie, and to at the same time to be committed to getting my brother to lie, is to be in a state of contradiction—logical contradiction. This is not because such commitments would by necessity fail to mirror some mind-independent moral state of affairs, but instead because the goal of one of these commitments is bound to be frustrated.

Now, this modification of Blackburn's semantics is only plausible if we agree on a few more things. The first is that such higher-order embeddings of commitment or intention (or some attitude like them) really do logically bind us in the way I argued; not everyone is convinced that if one intends to intend to do a thing, for example, one is thereby compelled on pain of logical contradiction to also intend to do that thing. The second and more consequential point on which we must agree is that there really are such higher-order mental states, and that people actually do express these when they sincerely utter complex moral claims. Even if the argument is up to this point convincing, we cannot assume that this is the case—the burden of proof (or evidence, at least) is on the higher-order theorist. That moral claims express approval and disapproval, *à la* Ayer, is *prima facie* plausible—though perhaps ultimately untenable. But the assertion that these claims express complex commitment states needs independent justification, and this justification cannot simply be that “It makes the theory work.” None of this is to say that the situation is hopeless, of course; I'm just pointing out the explanatory burden such an approach shoulders.

Of course, higher order semantics aren't the only way that expressivists have gone about answering the Frege-Geach problem. But all of the approaches tried to date have

had this feature in common: they rely on some account of the mental states being expressed to justify the inferential relationship between normative claims. Such an account must involve a plausible explanation of the inferential properties of mental states, such that these make sense of how different attitudes can be related to one another in ways that make sense of validity and contradiction. And as with the case of higher-order semantics, this reliance puts the expressivist in a tough spot. I want to look at the recent work Mark Schroeder, who has specialized in pointing out those spots, to see what we can learn from his arguments, and what options they leave us.

In *Being For*, Schroeder argues that in order to give a robust account of our everyday use of moral terms, the moral expressivist must undertake a host of increasingly problematic semantic commitments. I will look at the arguments he offers for two conclusions:

First Commitment: In order to give a consistent treatment of sentential connectives, the expressivist must treat all such connectives—whether they stand between representational or non-representational atomic sentences—as expressions of mental states.

Second Commitment: The expressivist must hold that both representational and non-representational claims function to express the same kinds of mental state.

Schroeder's arguments for the second commitment are not conclusive; as we will see, the pragmatic principles I have considered already in Chapter Two and in the first half of this chapter might give us the tools to reject it. Nevertheless, I think his argument for these commitments are compelling and disturbing for the moral expressivist. It will lead us to other considerations of the expressivist project, which will motivate my work in the Part Two of the dissertation, where I outline a non-expressive approach to non-representational semantics for ethical claims.

7. Argument for the First Commitment

What's particularly worrisome is that Blackburn's account of conditionals

has the disadvantage that it does not give us a single unified account of conditional utterances, from which the required account of conditionals with embedded moral clauses falls out as a special case—conditions in general are not expressions of moral sensibility. I suspect that Blackburn has confused two notions of endorsement that we give to any proposition when we assent to it, the first the semantic endorsement that we give any propositions when we assent to it, and the second the peculiarly moral endorsement that we give to an act or state of affairs of which we approve. (Price 2011, 64)

Consider for example, if we accept both the moral and modal responses to the Frege-Geach problem considered here, what are we to make of claims that combine moral *and* modal clauses, like

13. If it's possible that mammals are conscious, then eating meat is wrong.

We face a dilemma. Does this sentence express an evaluative sensibility or a linguistic sensibility?

H!|P!(‘mammal’: ‘conscious’);B!(eating meat)|

or

C!|P!(‘mammal’: ‘conscious’);B!(eating meat)|

If it expresses a linguistic sensibility, we will have to have a new way to justify the validity of an argument that starts that combines (13) with the expression of a moral attitude:

Eating meat is not wrong.

To the conclusion:

Therefore, it's possible that mammals are conscious.

Likewise, if (13) expresses an evaluative sensibility, how does coupling it with the expression of a linguistic sensibility:

It's possible that mammals are conscious.

entail the conclusion:

Therefore, eating meat is wrong.

Here are the expressivist's options: First, we could explain how mental states of completely different types enter into inferential relationships with one another. This is the option Timmons and Horgan pursue; I'll explore its implications in the section below. Second, we can attempt some sort of univocal treatment of conditionals and other complex sentences that makes sense of the validity of this sort of argument. But this would amount to treating complex moral and modal claims as if they are expressions of the same fundamental sort of mental state, which is at least on its face rather dubious. It also would imply that if we're tempted by any other sort of expressive treatment—of epistemology or causality, for example—we'll have to give a univocal account of the attitudes claims in all of these domains express.

To avoid the dilemma altogether we might be tempted to abandon our expressive treatment of one of these discourses—and so presumably opt for a representational account of either modality or morality, with all of the attendant problems such an account brings. But there's reason to think that once we adopt an expressive account of the meaning of any one sort of claims, the dilemma is going to rear up again when we combine those claims with standard representational claims. Consider the following complex sentences.

14. If stealing is wrong, so is murder.
15. If stealing is wrong, then Todd's in jail.
16. If he stole the diamond, then we should tell the police.
17. If he stole the diamond, then Todd's in jail.

Expressivists account for the meaning of (14) by holding that the indicative conditional expresses a mental state whose content is a function of the attitudes that would be expressed by its constituent parts—disapproval of stealing and disapproval of murder. What are we to make of sentences like (15) and (16), though, which have one atomic part that expresses a normative state and one that simply reports on a state of affairs? One option is to hold that in each of the sentences above, the function of the conditional is different, and to introduce new accounts for each of the logical connectives: the meaning of the conditional in (14) would be given by mental states that would be expressed by its antecedent and consequent; the meaning of the conditional in (17) would be given in the traditional representational manner, presumably, by a truth function between the propositions referred to in its antecedent and consequent; the meaning of the conditional in (15) and (16) would be given as a function between the mental state expressed in its antecedent or its consequent, respectively, and the proposition referred to by the complement.

Such an approach is not just messy but rather implausible, first because our use of all of the sentential connectives—conditionals, conjunctions, disjunctions, and negations—in natural language is apparently uniform; we don't indicate our conditionals are functioning differently when they govern normative versus descriptive sentences. Also, an account of the sentential connectives that is not univocal will have difficulty explaining a plausible feature of logic:

Take any inference of the form: 'If A then B; but A: hence B.' This inference is valid in exactly the same way regardless of whether or not 'A' or 'B' contains normative language. Expressivists must not only explain how normative terms can occur, without shifts of meaning, both in embedded contexts and elsewhere; they must also explain how the meaning of normative terms and of the logical operators allows all the customary forms of inference to be valid, in exactly the same way as

they usually are, regardless of whether or not they contain normative language. (Wedgwood 2001: 44)

So there is considerable pressure to admit that the conditional plays the same fundamentally expressive role in these. “If any of the conditionals [in (15)-(17)] are the same as the conditional [in (13)], then it must be *the same function from two mental states to a third mental state*.” (Schroeder, 2008a: 95, emphasis mine) This entails what Schroeder calls ‘mentalism’, the view that complex representational sentences like (17) above get their meaning (and semantic properties) from the mental states they express. This holds for complex representational sentences like (17), but it is also implied for its atomic parts¹⁴.

Mentalism is, as Schroeder puts it, “non-trivial” (23). The expressivist approach to the semantics of complex normative claims is forced into a substantive position about the content of complex representational claims. On the “standard” semantic picture, the content of a sentence—and the belief that it expresses—can be read off of the truth conditions for the proposition to which it corresponds. But on the expressive semantic picture argued for here, this is turned on its head: the meaning of a representational sentence doesn’t come directly from the proposition to which it corresponds, but instead from the mental state that it expresses.¹⁵

¹⁴ Schroeder goes on to argue that it is *entailed* for the atomic parts—that the expressivist isn’t only committed to mentalism about complex representational sentences, but also mentalism about atomic sentences. As we’ll see, I don’t think this argument is dispositive.

¹⁵ See Schroeder (2008a: 31-34)

8. Arguments for the Second Commitment

None of this is conclusive. Mentalism may, after all, be the right account of descriptive content. It does show us, however, that the expressive approach to semantics—in which we understand the semantic properties of normative sentences in terms of the mental states they express—cannot be made in a vacuum. It has implications for fundamental semantic issues.

For one thing, it returns us to our dilemma. Given mentalism, it would be simplest to assume that the mental states expressed by representational claims are simple beliefs, but Schroeder argues that similar considerations will lead us to the conclusion that on any adequate expressive account, both normative and representational sentences must function to express the *same kind of attitude*. And we reach the same conclusion for modal claims, whether we understand them representationally or expressively. Consider an expressivist treatment of a complex normative sentence like “Either murder is wrong or stealing is wrong”. On the combinatorial semantic account the expressivist needs, such a statement is best understood as an expression of a mental state whose content is a function of the sensibilities that would be expressed by its atomic parts, perhaps a mental state “in which if one side is closed off to me, I am to switch to the other—or withdraw the commitment.” (Blackburn 2002: 171) The same would hold for the complex modal sentence, such as “Either water is possibly H₂O, or it’s necessarily H₂O.” Likewise, the sentence “Either he did it or Todd is alive” expresses a mental state—one might presume a belief—the content of which is in similar fashion given as a function of the mental states that would be expressed by its parts.

Now compare this to a sentence that mixes the normative and the representational: “Either murder is wrong or Todd is alive.” The first atomic part of this expresses a normative attitude towards murder, and the second expresses a particular belief about Todd. Now we have the dilemma again: does this disjunction itself express a belief, or a normative attitude? Schroeder argues that here

we can try to generalize the semantics for descriptive sentences to apply to normative sentences. This would be to go with the view that every sentence expresses a belief. But that is just to give up on expressivism altogether. (Schroeder 2008a: 91)

We can’t consistently maintain an expressive account of normative sentences if we’re accounting for the meaning of a mixed sentences like the above in terms of representation; after all, it is characteristic of expressivists that they insist that normative attitudes don’t function primarily to represent. So Schroeder argues we must conclude that this disjunction itself is an expression of a normative attitude—“there is nothing for it but to conclude that all sentences must really express the same kind of attitude.” (ibid)

From this, Schroeder argues for a commitment-like state he calls “being for”, which is the fundamental mental state underlying belief and moral attitude, so that “Murder is wrong” expresses a state of *being for* blaming for murder, and “Grass is green” expresses a state of *being for* proceeding as if grass is green (ch.9). And perhaps we can extend this to account for modality: modal claims express a state of being for adopting the relevant linguistic rules¹⁶. But the task of justifying validity remains: even if we postulate

¹⁶ Schroeder himself thinks that his “being for” account of representational beliefs and evaluative attitudes leads to bizarre consequences when it’s used to account for modal claims, claims involving tense (p.170-171), and claims with restricted quantifiers like “most” (p.175-176). The lesson he argues we should take from this is that “if the most promising way of addressing the simplest possible questions leads so directly to an account that cannot accommodate the semantics of “most” [or “necessarily”], then expressivism should not inspire our confidence.” (176) Because I will turn my attention in the remainder of this dissertation to a treatment of non-representationalism that doesn’t rely on the expressive assumption, I won’t address these arguments, though see Yalcin (2012) and Marbitto (2009) for critical assessments of Schroeder’s project.

a type of mental state that all claims ultimately express, we'd still have to explain why we can make valid inferences between different kinds of claims—moral, modal, or straightforwardly representational. That is, we'd have to give an account of the inferential relationships that hold between the mental states in question, such that those relationships make sense of our logical commitment to the conclusion of e.g. *modus ponens* arguments—and to make sense of the inconsistency inherent in accepting the premises of such arguments, but denying the conclusion. I've shown above what's involved in justifying validity for arguments with complex and simple moral and modal claims: the trick is to identify the purpose of the discourse in question and show that the pressure to maintain coherent mental states of the type the claims in the discourse typically function to express is a precondition for the discourse to serve its characteristic function. What would be needed, if we agree with Schroeder that all assertions express the same basic mental state (i.e., being for), is an identification of the purpose assertive discourse plays that would be undermined without a systematic commitment to coherence. That is, to say the least, a tall order.

Recall that another option is available, though. What's to stop the expressivist from assuming complex sentences like:

If stealing is wrong, then Todd's in jail.

do express a sensibility, but a complicated one that is a function of the mental states—one a normative attitude, and the other either a linguistic sensibility or a belief—of its two constitutive atomic parts? Perhaps the attitude expressed is a sort of commitment to either believe that Todd is alive, or be committed to a disapproval of murder. Likewise, perhaps the sensibility expressed in the modal/moral sentence:

Either it's possible that animals are insensate or eating meat is wrong. is a similar sort of commitment, to either have the linguistic sensibility that permits one to defeasibly use 'insensate' wherever 'animal' refers, or the evaluative sensibility that functions from the input an act of eating meat to the output of, say, commitment to disapproval. So long as we're able to understand the sentential connectives as expressing such attitudes, as the arguments above suggest, the validity of any particular argument would remain untouched. Consider:

- 18. Either murder is wrong or Todd is alive.
- 19. It's not the case that Todd is alive.
- 20. Therefore, murder is wrong.

(18) and (20) above would express evaluative sensibilities—where the sensibility expressed by (18) is a function of the two different kinds of sensibilities that would be expressed by its parts—and (19) would express a belief, and not a normative attitude, as Schroeder insists. Given that the normative attitude expressed in (18) would reflect the ways in which those distinct states could be combined, and that (19) would shut off one possible such combination, it doesn't seem particularly surprising that (20) would follow. This would be a *prima facie* reasonable way for the expressivist to account for the meaning and semantic properties of a disjunction that mixes normative and descriptive atomic claims without accepting Schroeder's argument that all claims must ultimately express the same fundamental kind of mental state.

But Schroeder has another argument leading to this commitment—that the mental states expressed by both descriptive and normative claims are ultimately the same kinds of normative attitudes—and it's one that speaks to the above case. Consider the alternate conclusion:

- 20a. It's not the case that murder is wrong.

Obviously {18,19,20a} is an inconsistent set of propositions, and in the expressivist framework, that inconsistency must boil down to inconsistency between mental states. But as we've seen, not all mental states produce inconsistency. For example, wondering or supposing are both attitudes that do not "transmit" inconsistency; one can wonder p and wonder $\sim p$ at the same time without contradiction.

So we must look to mental states that do transmit inconsistencies. Intention and belief offer good paradigm cases of inconsistency-transmitting attitudes: to believe p and at the same time believe $\sim p$ is to be in contradiction; the same holds for intention. These are both attitudes that create what Schroeder calls "A-type inconsistency":

An attitude A is inconsistency-transmitting just in case two instances of A are inconsistent just in case their contents are inconsistent. (43)

The problem is that beliefs can be A-type inconsistent with other beliefs, and intentions can be A-type inconsistent with other intentions, but there is no such inconsistency to be found between beliefs and intentions. On the explanation limned above, the inconsistency found in the attitudes expressed by {18,19,20a} is a B-type inconsistency—an inconsistency between mental states that involves bearing different attitudes towards the same content.

Schroeder argues that there are two central problems with explanations of semantics that rely on B-type inconsistency. The first is such an account would be radically different

from the familiar kinds of inconsistency for which expressivists have good models elsewhere. All the other good models of inconsistency between mental states arose in the case of inconsistency-transmitting attitudes. These were all cases of the *same attitude toward inconsistent contents*. (48, emphasis mine)]

We can, of course, stipulate that there is an inconsistency between the beliefs and commitments in {18,19,10a}, but simply assuming such an inconsistency is “taking for granted everything that expressivists need to explain. (ibid)”

Furthermore, once we go down this path of stipulation, Schroeder argues we’ll be committed to infinitely many such different attitudes. This argument gains its force through a critique of the expressivist account of semantics argued for by Horgan and Timmons’ (2006) in their paper “Cognitivist Expressivism”, in which they work from the assumption that B-type inconsistency accounts for the inconsistencies we find in sets like {18,19,20a} above.

The central argument of Horgan and Timmons’ paper is a rejection of the cognitivist/non-cognitivist distinction that is traditionally used to characterize expressivism. They argue that our moral judgments are in fact a species of belief that can be expressed with assertions, though unlike other typical assertions, moral claims express a non-representational kind of belief. Representational beliefs are classified as mental states of “is-commitment”; non-representational beliefs are mental states of “ought-commitment”. Is- and ought-commitments attach to some specifiable content, with different resultant meaning, such that an is-commitment toward the content that Obama is president yields the belief that Obama is president, whereas an ought-commitment toward the same content yields the belief that it ought to be the case that Obama is president.

These two species of belief, involving two distinct ways of mentally affirming a core descriptive content, are both *sui generis*: neither type is reducible to some kind of nonbelief state such as an attitude of approval. On their account, each combination of sentences—and each of their negations—expresses a new and distinct kind of mental

state, a “logically complex commitment state”, which is constituted by its inferential role with other states. “Such a state is one that is poised to interact with other potential beliefs to inferentially generate yet further beliefs that are inferentially ‘in the offing.’” (278)

Complex sentences like “Either he did it or Todd is alive” or “If stealing is wrong, so is murder” can themselves be understood as expressions of is- or ought-commitments, respectively (see Appendix Note 13 and 14: 295), but complex sentences involving both normative and descriptive elements, like “Either murder is wrong or Todd is alive” are neither is- nor ought-commitments, *but are instead themselves sui generis belief states of logically complex inferential commitment*. These states can combine with is-commitments, ought-commitments, or other logically complex inferential commitments to yield entailments or contradictions. This would explain why a logically complex sentence like (18), when combined with (19), entails (20); it also explains why the {18,19,20a} is an inconsistent set.

The problem with all of this, Schroeder points out, is that this view

amounts to the hypothesis that there is an unfathomably huge hierarchy of distinct kinds of mental state, together with the unsupported confidence that these mental states will have the right inconsistency relations with one another. (2008a: 51)

Horgan and Timmons provide a model theoretic account of expressive semantics, but don’t give any justification of why we should think there are such mental states with these inferential properties. They give a *desideratum* for their semantics—that we actually have these complex mental states. What they don’t do, Schroeder complains, is supply any independently plausible reason we should expect such a desideratum to be fulfilled.

Now, it’s not obvious to me that such an explanation won’t be available. My tentative solution to the Frege-Geach problem in the first half of this chapter shows how this is to be done: supply an independently feasible account of the function of the attitudes

expressed in moral discourse, and then show how such a function is undermined by committing to mental states that have the right kind of (A- or B-type) inconsistency. This seems like something we could plausibly do in the case of Horgan and Timmons' logically complex commitment states: perhaps these states function to bring about social coordination of acts, emotions, and dispositions. In order to do so, it is crucial that representational doxastic commitments are brought into accord with other commitments; when they are not (as with the set {18,19,20a} above), the function of those mental states is bound to be frustrated.¹⁷

So I believe Schroeder's arguments for the second conclusion are inconclusive. Still, though, the kinds of explanations available for expressive semantics are troubling, and not simply on the grounds that they lead to Schroeder's first commitment—mentalism about the content of descriptive sentences. Contemporary expressive accounts—Allan Gibbard's hyperplans, Simon Blackburn's notion of being tied to a tree of commitments, and Horgan & Timmons' logically complex commitment states—all share in common that they

try to account for inconsistency of sentences in terms of the states of mind expressed, entailment among sentences in terms of commitment between states of mind expressed, and so on. These are the materials with which expressivists have to work. (Schroeder 2008a, 39)

This is the operating assumption of expressivist semantics—that the inferential properties of a claim are given by the mental state that claim expresses. This means that any semantic account given by an expressivist, no matter how elegant, will bear a justificatory burden: not only must such an account explain all the realist-seeming features of moral discourse, but the complex mental states on which such a semantic account is

¹⁷ In a review of Schroeder's book, Robert Mabrito offers a strikingly similar approach with his proposed assumption: "A set of attitudes is inconsistent just in case the aims associated with those attitudes are not mutually realizable." (2009)

predicated must be the kinds of states we can reasonably believe people are actually in, when they sincerely make sincere moral utterances. This is not damning—such an account may be forthcoming. But I wonder if it’s really a burden that the non-representationalist must bear.

9. Conclusion

Up to this point in the dissertation I have worked from this operating assumption of expressivism—that the meaning of moral claims is a function of the mental states the express—to show how far it could take us. In Chapter Two I argued that expressivists could justify claims to objectivity in moral discourse if the function of the mental states expressed by that discourse would be undermined without pressure toward systematic coherence; I arrived at a related conclusion above, where I argued that a solution to the Frege-Geach problem can only come for the expressivist when she has given an account of the function of the mental states moral claims express such that those states can stand in the appropriate inferential relationships to one another. Given the difficulty of such an approach, and the commitments it must assume, there seems to be good reason to drop this assumption, though, and to proceed with a non-representational account of moral discourse that doesn’t rely on the expression relation, but instead looks directly to the use to which we put moral terms.

Obviously, one reason to attempt such a revision of moral non-descriptivism is to avoid the briar patch of a semantics based on the inferential relationships between mental states—and my objective in the remaining chapters of the dissertation is to see just how

well a non-expressive form of moral non-representationalism can handle the typical issues that plague expressive accounts. Another reason for the undertaking, though, is that we should have a *prima facie* preference for such an emendation, purely on grounds of simplicity: an explanation that relies on the functional role of moral discourse itself is more straightforward than an explanation that relies on the functional role of the mental states that claims in the discourse characteristically express. The semantic order of explanation for the expressivist starts with mental states that are expressed by moral claims and then spells out the inferential roles the states themselves have; the content of moral claims comes from the attitudes expressed. It would be more direct if we could look directly to the inferential role the moral terms themselves have—without worrying about the mental states expressed—to explain the meaning of moral sentences and the kinds of inferences we can make with them.

And such a move would not be *ad hoc*; as I've already pointed out, other forms of non-representationalism explain the inferential relationships among the sentences within a given domain without reference to mental states—it is, in fact, the norm for non-representational accounts in other domains. Amie Thomasson has also recently wrestled with the Frege-Geach problem as it applies to modal normativism. I have in this chapter developed a parallel, explicitly expressivist version of her non-representational account of modality. Like my own, her actual account explores the inferential role played by the necessity and possibility operators, but at no point does she stop to consider the importance of the mental states being expressed by the use of modal terms. And at no point does she need to.

In the following chapter, I will start establishing some of the specifics of a non-representationalism about morality that doesn't rely on the idea that the meanings of moral terms comes from the mental states they express. This will begin with an inquiry: What is the function of particular moral concepts, and what inferences do these license? After outlining this project, I will move back and reconsider the problems that expressivist semantics faces, surveying the prospects for responses without the basic expressive assumption about the meaning of moral claims. I'll argue for a non-expressivist version of moral non-representationalism—moral *inferentialism*—and show how it solves the Frege-Geach problem without letting the expressivist's problems creep in. Then in Chapter Five, I'll turn my attention back to the question of moral objectivity, and argue that an inferential approach will also shed light on some of the worries about objectivity that we were left with in Chapter Two.

PART TWO

CHAPTER FOUR: MORAL INFERENCEALISM AND THE FREGE-GEACH PROBLEM

1. Overview

The basic semantic approach that expressivists use makes them vulnerable to the Frege-Geach problem: how can expressivists account for the uniform content of moral terms even when they are uttered in “force-stripping” contexts, where their typical use is not in effect? Though there are several attempted responses to this issue (in the last chapter we saw one such attempt) recent literature on the topic suggests that these are deeply problematic (Unwin 1999, 2001; Schroeder 2006; van Roojen 1996; Schueler 1988), and leaves the expressivist with a heavy explanatory burden, most of which involves explaining exactly how mental states are supposed to inform both the meaning of moral claims and explain the validity of arguments with moral components. In this chapter, I will argue that there is an alternative, moral inferentialism, which can sidestep this burden but that nevertheless retains expressivism’s core advantages.

After giving a brief sketch of metaethical inferentialism, I’ll explain the Frege-Geach problem and how expressivists deal with it. I’ll then offer a sketch of how a moral inferentialist can do the same, with an account of the inferential role played by the moral term ‘ought’. I’ll argue that such inferentialism enjoys distinct advantages over expressivism: not only is it a simpler position, it can solve the Frege-Geach problem without recourse to the mental states on which its expressivist cousin relies. This means

that inferentialism can avoid many of the semantic complications that come in the wake of an expressivist solution to the Frege-Geach problem. I'll consider two well-known complications from the literature, and show how an inferentialist account avoids them. This chapter is intended as an argument for inferentialism as a metaethical position that preserves the advantages of expressivism while avoiding some of its problems.

2. Inferentialism: Expressivism without Expressing

The representational approach assumes that once we've accounted for the pragmatic role of terms on the 'near end' of the word-world semantic relationship inherit in moral claims, we must then go on to characterize the truthmakers for such claims—those properties or relations that lie on the 'far end' of this semantic relationship.¹⁸ The non-representational approach rejects this assumption; if we can explain the features of moral discourse *pragmatically*, without involving the referents of terms in any explanatory role, we won't need to engage in this sort of truthmaker-hunting project; all the explanatory work will have been done. Metaethical inferentialism is a very close cousin to expressivism; both can rightly be understood as grandchildren of AJ Ayer's emotivism (1936). As we've seen, on Ayer's conception, moral claims lack "factual meaning" (p.107); moral claims do not correspond to anything that can be considered true or false. Instead, they simply evince one's moral approval or disapproval. This approach is non-representational in that it denies that moral assertions function to represent moral facts, properties, or states of affairs.

¹⁸ Cf. (Price and Macarthur 2007, p.94-97)

But as we've seen in Chapter One, such an emotivist approach does entail a radical reformation of ethical discourse. We think we can speak non-problematically about moral beliefs and moral errors, but if moral claims have no content, the structure of ethical discourse is grammatically and logically misleading, and we'd do well to look past its surface features. This is a high price to pay to be an emotivist: if Ayer is right, it's hard to make sense of some very basic features of ethical discourse in a way that doesn't deeply undermine the discourse. The denial that ethical claims have meaning seems to undermine important ethical notions like moral error, disagreement, and improvement.

Expressivism avoids this conclusion by arguing that ethical claims do in fact have content, while denying that such content is given in terms of *representation*. It retains Ayer's basic contention—that moral claims function to express rather than describe—but argues that it is this expressive function that actually gives moral claims their meaning. Expressivists explain the semantic features of moral claims by reference to this essentially expressive function; the meaning of moral claims is explained in terms of the mental states they are used to express.

If we understand the progression from emotivism to contemporary expressivism in this way, metaethical inferentialism can be characterized according to another distinction from its relatives. It retains the non-representational aspect of expressivism's semantic account—the rejection of the idea that moral language functions to represent moral facts—but denies that the fundamental non-representational function of moral language is best understood in terms of the mental states moral claims serve to express. The inferentialist understands “the commitments which constitute a claim's meaning not in ontological or psychological terms but in inferential terms”. (Chrisman 2010: 118) The

metaethical inferentialist argues that we can identify the meaning of moral claims by articulating their inferential connections to other claims, and that these connections are best understood by looking to the pragmatic function moral discourse fulfills.

According to the general inferentialist approach to meaning advocated by Robert Brandom (1994), the content of a statement is not given by the conditions which would make it true, but is instead best understood in terms of the inferences that license the statement, and those that it licenses. The notion of inference here is quite broad; it is not limited to the formal inferential connections that hold between atomic sentences regardless of their content (for example, the inference from $A \& B$ to A) but also includes material inferences—those inferences that one is justified in making in virtue of the content of a claim (for example, the inference from “It is green” to “It is colored”). On this approach, the meaning of a statement is found in its role in the social practice of asking for and giving reasons to justify the statement, and making inferences that follow from the statement.

According to Brandom, these inferential or conceptual roles associated with a claim jointly constitute its meaning.

The conceptual role played by [a] claim in the social practices of the community is determined by what further assertions that community would accept as appropriate justifications of it, and what assertions they would take it to license or justify. (Brandom 1983: 642)

According to Brandom, it is this social practice of keeping track of both what inferences a claim entitles one to make and what sort of inferential obligations one undertakes in making a claim that is essential to an account of its meaning. These inferences will include language-language transitions, which denote the inferential relationships between propositions, but also “language exit” and “language entry” moves that “take in the

circumstances and consequences of application” (Brandom 2010: 159) of claims, where the circumstances of application may be perceptual rather than propositional (so that we are justified in saying, “That is red” when we see something red), and the consequences may include specific behavior (so that one normal consequence of accepting, “That tastes horrible” is to avoid eating it). Metaethical inferentialism advocates an inferential treatment of moral claims.

The development away from an expressive account should not be such a surprising step in the evolution away from emotivism, if for a moment we look away from debates in metaethics and cast our eye on non-representational accounts of other types of discourse. For example, Gilbert Ryle explains the import of law-like scientific claims in terms of the inferences such claims serve to license; a sentence such as “Salt dissolves in water” licenses the inference from “X is salt” to “X will dissolve in water” (Ryle 1949). And as we’ve seen, Paul Horwich argues that use of the truth predicate in assertions is important only as a generalizing device; the meaning of “P is true” is given by the equivalence schema: the proposition expressed by ‘P’ is true iff P. (Horwich 2010). Amie Thomasson argues that modal sentences like “Necessarily Q” make explicit the semantic rules that govern our use of the term ‘Q’. (Thomasson 2006) Notice that like expressivism, all of these forms of non-representationalism account for the discourses at hand in ways that that don’t rely on some peculiar moral/alethic/causal/modal range of facts waiting to be described. Each instead explains the discourse in terms of how we *use* the concepts in question—and our use in each case is not primarily representational. But it is only expressivism that ties this use down to mental states. Viewed this way, this seems a peculiar fact. I suspect that this feature of expressivism can be understood as a result of

the contingent fact of its emotivist lineage; metaethical inferentialism offers us a new way to understand how the uses of moral terms inform their meaning.

As we will see, one reason to prefer the inferentialist revision of expressivism is to avoid the briar patch of a semantics based on the inferential relationships between mental states—some of which we glimpsed in the last chapter. Another reason for the undertaking, though, is that we should have a *prima facie* preference for such an emendation, purely on grounds of simplicity: an explanation that relies on the functional role of moral terms themselves is more straightforward than an explanation that relies on the functional role of the mental states that moral terms characteristically express. The semantic order of explanation for the expressivist starts with the mental states that are expressed by moral terms and then spells out the inferential roles the states themselves have¹⁹; the content of moral claims comes from the attitudes expressed. It would be more direct if we could look straight to the inferential role the moral terms themselves have—without worrying about the mental states expressed—to explain the semantics of moral discourse. We'll have a simpler explanation because we don't need to assume anything about the expression relationship between moral claims and mental states²⁰, and there won't be any explanatory pressure to explain just how it is that mental states enter into inferential relationships with one another.

¹⁹ See for example Horgan and Timmons (2006) and Gibbard (2003, esp. Ch. 3).

²⁰ See for example Schroeder (2008b), and Charlow (2013, p.4)

3. The Frege-Geach Problem and the Standard Response

Recall that the challenge of the Frege-Geach problem has two parts. Consider again a simple example of a *modus ponens* argument with a moral component:

- 1) If he cheated, he ought to be punished.
- 2) He cheated.
- 3) He ought to be punished.

First, for this argument to be valid, the meaning of “he ought to be punished” must be the same in (1) and (3). If it isn’t, the argument is invalid because of the equivocation involved.

A few expressivists have tried their hand at responding to this challenge (Blackburn 1984, 1988, 2006; Gibbard 1990; Horgan and Timmons 2006); I want to point to general features these responses have in common. What’s needed, as we’ve seen in Chapter Three, is a compositional semantics that understands the meaning of a complex sentence like (1) given as a function of the meaning of its parts. So the meaning of (1) is a function of the *mental state* expressed by (3). The implication is that the content of (1) is also given by a mental state being expressed—a complex state, the meaning of which is determined in part by the attitude that *would* be expressed in a force-stripped context by its consequent clause.²¹

“He ought to be punished” has the same content in (1) as it does in (3), even though (1) does not directly express approval of punishment, so this sort of response clearly refutes the charge of equivocation. But why does the content of (1) and (2) taken together license the inference to (3)? Why is the argument valid? This is the second part of the challenge. Expressivists can’t simply argue that such an inference is valid on the

²¹ See Schroeder (2008a: 20-21)

grounds that if (1) and (2) are true, (3) necessarily must be so too. Such an explanation of validity would fundamentally rely on *truth-conditions* for normative claims, and one of the chief advantages of expressivism is that it needn't posit any particular states of affair or properties that would stand as truth-makers. Instead, the validity of the inference from (1) and (2) to (3) must be explained, once again, in terms of the mental states these sentences express. The idea is proposed that if one has the mental states expressed by both (1) and (2), one is compelled by *force of logic* to also accept the mental state expressed by (3)—to accept (1) and (2) but reject (3) would entail a state of contradiction. The notion of contradiction at play in this explanation is fundamentally non-cognitive; expressivists don't cash it out in terms of logically inconsistent representations of the world, but instead in terms of jointly unsatisfiable plans, intentions, or evaluations.

This is a tidy sort of response, but it seems to carry with it a host of complications and unintuitive commitments, many of which can be found in the literature on expressivism in the past 20 years (Unwin 1999; van Roojen 1996, Schroeder 2008b; Dreier 2006; Schueler 1988), and some of which we saw in the last chapter. For example: It's not clear that the logical notions of consistency and inconsistency are appropriately ascribed to the relationships between non-cognitive mental states (Schueler 1988); making normative content dependent on non-cognitive mental states in the way expressivists argue for commits us to also explaining descriptive content in the same way (Schroeder 2008b), and to explaining the meaning of logical connectives in the same way—both of which would involve the expressivist in a radical modification of compositional semantics for non-normative language. Also, an account that explains the semantics of normative claims in terms of evaluative states also seems to involve conceptions of consistency and

inconsistency that entail validity in arguments that common sense would tell us aren't in fact valid (van Roojen 1996); such an account has trouble explaining the logical relations between obligations and permissions (Unwin 1999). Obviously, there are arguments and attendant counter-arguments to address these issues (see Dreier 2006, Mabruto 2009, Alwood 2010), but in this chapter I want to suggest that one can avoid at least some of these problems by adopting a semantic account that answers the Frege-Geach problem *without* depending on mental states as the fundamental driving force for a non-representational semantic account of moral language.

To accomplish this, I'll focus on the last two from the list of complications above, on the grounds that these have proven to be particularly tricky for the expressivist. I'll explain them in more depth in Sections 8 and 9 below, but to help keep our eyes on the target, let's name them both:

The Problem of Permissions. Because the expressivist relies on the notion of mental states being in contradiction with one another to explain logical contradiction between moral claims, she must account for the apparent contradiction between statements of obligation and permission with reference to a *sui generis* category of mental state—tolerance—that has its own unique semantic properties.

The Commitment to 'Mentalism'. In order to consistently account for the meaning of logically complex sentences that involve both normative and descriptive atomic parts, the expressivist seems to be committed to the unusual semantic view that the content of the logical connectives—and possibly the

content of both normative *and* descriptive claims—must be explained in terms of the mental states each typically serve to express.²²

Before I show how inferentialism can meet (or avoid) both of these challenges, though, I'll start by explaining how the view can tackle the basic semantic challenge the Frege-Geach problem poses.

4. Meaning from Use

To solve the Frege-Geach problem, the task for the moral inferentialist is to give an account of the meaning of moral terms by way of specifying the inferential relationships those terms have, relations that are necessitated by the use to which we put those terms. As with expressivism, this account is constrained in the following ways:

1. Our account should be lightweight, in the sense that I explained in Chapter Two: It cannot invoke as explanatory a range of moral facts, properties, or relationships that are being represented by moral claims.
2. This account must give the meaning of moral terms via the inferential roles that give these terms their meaning, and this meaning—and so too the inferential roles—must remain the same whether the term is embedded or not. This is the task of *refuting equivocation* that I explained in Chapter Three.

²² In Chapter Three, I agreed with Schroeder's contention that mentalism about sentential connectives was a commitment that expressivists needed to take on, but that his arguments for mentalism about the content of all atomic claims—whether normative or representational—were not so conclusive. As we'll see, I don't think the charge of 'mentalism' has the same bite for the inferentialist, so I won't treat the distinction between the two charges in great depth here.

3. It also has to explain why the meaning of these terms is such that we can use them to make valid inferences—how the conceptual roles played by moral terms informs the inferential relationships we take to be basic in moral reasoning, including the relationships between permissions and obligations, and the validity of argumentative forms like *modus ponens*. This is the task of *justifying validity*.

How can we accomplish 2. and 3. without relying on representational resources?

We can take a page here from Michael Williams (2010), who develops a meta-theoretical framework to characterize non-representational theories that explain the meaning of terms entirely according to their use. In order to identify the inferential roles that a term plays, he argues, we must first consider the *functional* roles the term plays in our discourse.

These will be the roles that such a term must fulfill in order to serve the *pragmatic* role that the discourse plays in human life²³. It is only once we have such a functional explanation that we can move on to explain the inferential roles the terms associated with the discourse have; these should be the inferential roles—the patterns of usage that competent speakers will be disposed or expected to employ—by which the term fulfills the function of the discourse. These inferential roles determine the content of the term in question. This means that such an account will follow an order of explanation: The *meaning of a term* is constituted by the *inferential roles* that term plays, which are explained by the *functions* that term must fulfill, given the *pragmatic significance* of moral discourse.

²³ Williams himself doesn't distinguish between the pragmatic role of the discourse itself and the functional role a term in that discourse plays; both are treated in what he calls the "functional component" of an explanation of meaning in terms of use. For the purposes of this chapter, it will be helpful to be explicit about the distinction.

Perhaps an illustration is in order. Consider for example a non-representational account from another field: Paul Horwich's minimalism about truth (2010). As we've seen, according to Horwich, the meaning of the truth predicate can be given by the very simple inferential role given in his *equivalence schema*:

(MT) The proposition *that P* is true if and only if P.

This inferential role is explained by the function of the truth-predicate—it is important as a generalizing device. The truth predicate enables us to simply say, “Everything John said is true”, instead of saying, “John said that he's happy, and he is happy; he said that he's coming to the party, and he is coming to the party; and he said that monkeys are funny, and monkeys are funny.” We can also use the truth predicate to endorse a proposition even when we're not sure of its actual content: “Chris said something about the transmission needing repair, and I don't know anything about cars. But Chris does, so I'm sure what Chris said is true.” Why do we need a term that fulfills such a function? Because of the pragmatic significance of truth-talk: “we need a device that enables us to overcome finite constraints in our effort to describe the world.”(Beall 2009, p.1)

We find in Horwich's minimalism a clear example of how to give the meaning of a term without invoking any sort of underlying property, fact, or relationship being referred to. The *meaning* of a term comes from its content-determining *inferential role*. “With respect to giving the meaning of ‘true’, the rule of use implicit in our acceptance of the instances of MT is explanatorily fundamental.” (Williams, p.322) The inferential role governing our use of the truth predicate is one that enables it to fulfill its *function* as a generalizing device. Having a term that fulfills this function makes sense given the broad *pragmatic import* of truth-talk. Locating the functional role of a term within a practical

context “gives the point of our having a word with those [inferential] characteristics.” (p. 324)

Recall that in Chapter One I argued that the semantic rules of use for a term play a constitutive role for that term, similar to the way that the rules of chess play a constitutive role for the chess pieces. So for example what it takes to be a pawn in a game of chess is that it is a piece that follows certain rules:

1. It starts in a certain position—on the second rank of the board.
2. It can move two spaces the first time you touch it, and only one space otherwise.
3. It captures diagonally.
4. It promotes to become another piece when you advance it to the eighth rank.
5. It can capture *en passant*.

To fully understand a game of chess, you have to understand all of these things. But importantly, that doesn't mean that the pawn must display all of these capacities in any given move. Usually a pawn just follows (1)-(3), but we wouldn't get why players were making the decisions they make if we didn't get that a pawn always has (4) and (5) as two of its constitutive roles.

In what follows, I will outline a similar explanation of the meaning of an ethical term—the moral ‘ought’. This necessitates a digression on the pragmatic import of moral talk. In the next section, I will look at some of the typical uses of moral discourse—various functional roles it fulfills—and argue that all of these indicate a unified pragmatic role for moral language: the achievement of coordinated social behavior. I will attempt to be as anodyne as possible: the points I make will be recognizable to anyone familiar with expressivist literature.

5. A Pragmatic Account of Moral Discourse

Judgments involving moral terms have “indirect and variable influences on action” (Blackburn 1984: 189)—moral claims typically imply certain motivations on the part of the speaker, and sway the motivations of others. This practical aspect of moral claims is traditionally understood by expressivists in terms of moral judgment internalism, the view that sincere moral judgments necessarily involve some sort of motivation to act in accordance with those judgments. In the interest of skirting controversy, I’ll avoid this commitment²⁴. For the purpose at hand, there’s no need to narrowly construe this feature of moral claims in terms of some necessary connection between judging that ‘ ϕ -ing is wrong’ and being motivated to not ϕ -ing. Obviously it’s possible to act against one’s moral judgments because of fear, lust, greed, and so on—and perhaps one can make moral judgments without any attendant motivational import whatsoever. I’ll limit my point to this: there is some sort of conceptual connection between moral claims and influence on actions, however defeasible that connection may be. If we came upon a term in a foreign language that had no such practical import, I suspect we should not translate it as a moral term in our own language. This indicates that an essential feature of moral discourse is its behavioral impact (see Hare 1952: 148-150).

Such impact often comes from the connection between judgments involving moral terms and dispositions to have certain influential emotions: guilt or shame at having done something one judges wrong, pride at doing right, anger, outrage, contempt, or disgust at

²⁴ This hands-off approach also makes sense, given the distinction I’m drawing between inferentialism and expressivism about morality. The expressivist is committed to explaining the meaning of moral claims in terms of the (typically motivating) attitudes they express; inferentialism doesn’t carry this commitment, and so can remain silent on whether the characteristic action-guiding feature of moral claims is accomplished in the way that moral internalists argue or by some other means.

those who commit wrongs, and so on (see Gibbard 1990). Notice that these are all emotions that affect social behavior: an aversion to guilt (among other things) keeps one on the straight and narrow; fear of reproof and the punishment or ostracism that is often motivated by moral outrage or disgust can keep even those bereft of moral sentiments from transgressing.

In this limited sense, moral claims are similar to imperatives, in that they carry a kind of force; like commands, they have a defeasible tendency to influence our behavior²⁵. Some have taken this common feature to be fundamental, and have analyzed moral claims as a type of prescription: “Murder is wrong” expresses a command to not murder. (Hare 1952) No doubt the practical force of imperatives and moral claims is essential to their character, but there are important and instructive differences between the two types of speech acts.

First, the practical import of moral claims is of a special kind; it reflects what Kant would call a “categorical” rather than “hypothetical” imperative. Moral judgments are not mere pieces of advice or commands, the force of which only applies if one happens to have the right goals or desires, or if they are being issued by a person or institution of appropriate authority. The practical import of moral judgments is not taken to be contingent on particular desires to, say, avoid punishment, but rather to give us reason to act in certain ways regardless of our desires; their authority over us does not depend (*pace* divine command theory) on the individual delivering them. (Joyce 2006, Ch. 2) For moral language to have a pervasive effect on human behavior it cannot be entirely dependent on our expectations of punishment and reward (though such expectations often play an

²⁵ For a detailed exploration of the semantic and metatheoretical relationship between imperatives and normative language, see Charlow (2013).

important part in influencing)—often the most critical moral decisions are those that are made outside the normal reach of promised punishment and reward.

Second, moral claims made in the indicative are semantically more robust than imperatives. We can embed commands in the consequent but not the antecedent of a conditional. “If the window is open, then close it” is grammatical, but, “If close the window, then it’s open” is nonsensical. Making moral claims in the indicative gives us the ability to reason from moral claims to their consequences, and to reason among moral claims. It also allows us to make use of negation in an important way. The imperative form can only be used to express commands, but not permissions. We can introduce a negation in “Shut the window” to make the command, “Don’t shut the window”, but neither of these permit (shutting the window or not). There is room in indicative moral claims to make another kind of negation: “You ought to do your taxes” can be negated in two importantly distinct ways: “You ought to not do your taxes,” or “It’s not the case that you ought to do your taxes”, which is equivalent to “It’s permissible to not do your taxes.” These semantic features are important, since moral judgments involve and are involved in our reasoning towards practical (i.e., behavior-related) conclusions; they aid in, justify, and make explicit our reasoning about what to do.

In short, moral discourse carries with it a kind of social-behavior-influencing *force*: by making us liable (to others and to ourselves), it influences our behavior with a toolbox of emotional and social incentives and disincentives that outstrip mere hypothetical imperatives in both power and semantic subtlety. The discourse acts as a kind of chaperone for human behavior, constraining selfish or anti-social impulses. It has a kind of practical clout that we can deliberate about in complicated ways; we offer our

moral reasoning as justification for our behavior and our expectations for the behavior of others (and our expectations regarding the expectations of others), and as challenges to those who disagree. This capacity to think together on moral issues gives rise to the possibility of reaching consensus, both in moral opinion and the behavior that flows from it.

Why would creatures like ourselves adopt a discourse with these features? For a highly social, language-using species there is an obvious advantage to employing language that regulates behavior in the ways limned above: it promotes our chances for evolutionary success. Moral discourse offers us a potential counterpoint to the selfish impulses that would drive individuals to uncooperative behavior, even in situations when cooperation would be mutually beneficial; “the evolutionary function of moral judgment [and its attendant discourse] is to provide added motivation in favor of adaptive social behaviors (Joyce 2006: 117).” As Mark Timmons argues, we use moral language ultimately to provide

norms that serve to guide behavior in ways that are crucial to the stability and cohesion of groups of human beings. Because of the importance of certain patterns of behavior (and emotional response), one would expect that humans, capable of complex coordinative social behavior, would evolve more or less shared normative systems that have the sorts of features characteristic of a typical moral outlook. In short, given what human beings are like, the setting in which they find themselves, and that they are capable of quite sophisticated coordinative behavior, one would expect them to develop strategies that would most effectively solve problems of coordination... (Timmons 1999: 158)

The purpose of moral discourse, it seems, is to guide us into coordinated social behavior.

This is a rough sketch of features that I believe are essential to understanding what is distinctive about moral discourse, but it is not meant to be exhaustive. There may be other facets that are essential: moral discourse doesn't only goad us into behavior in a

particular way; it also seems to goad us towards particular kinds of behavior. Moral questions are typically concerned with issues of kindness, fairness, and so on.

I will not consider all of these features in depth here, or try to decide which are necessary characteristics of moral discourse. For the purposes at hand, it will suffice to echo Allan Gibbard, and point out that these make a particular kind of sense, in light of the broad coordinating function of moral discourse I'm arguing for. Feelings of benevolence prime us to act on behalf of others' interests, and this can lead to coordinated behavior that is to our obvious evolutionary advantage—when it benefits close kin or those in a position to reciprocate, but also because it is to our long term advantage to be a member of a species with altruistic impulses (Gibbard 1990: 258). Our concern with fairness lays the grounds for cooperation. “Judgments of fairness stabilize bargaining” (262) by giving us a common framework for assessing when gratitude (or retaliation) are appropriate. With some imagination, we might see why other issues are routinely within the purview of moral discourse—questions of respect, of disgust and communal purity, of guilt and shame—these are all concerns about our behavior that weigh heavily on our ability to cooperate.

This completes the first stage of our explanation of the meaning of moral terms according to their use given the pragmatic function of the discourse. The next step, according to Williams' framework for an explanation of meaning in terms of use, is to provide an inferential account of their meaning that falls out of the above pragmatic considerations. Understanding the inferential role moral terms play is a matter of understanding how those terms bring the characteristic behavior-influencing motivational force of moral discourse to bear for the purpose of coordinating social behavior.

6. An Inferentialist Solution to the Frege-Geach Problem

To solve the Frege-Geach problem, the inferentialist needs an account of the meaning of moral terms—an account that enables the above use of moral language. A moral term like ‘ought’ has a typical use like the one I outlined above, but this use is not always in effect.

Consider again the following argument:

- 1) If he cheated, he ought to be punished.
- 2) He cheated.
- 3) He ought to be punished.

The ‘ought’ in the conclusion of this argument has a behavior influencing force that is not present in its embedded form in the premise—judging (1) to be true does not directly influence me in favor of punishing my student; judging (3) to be true does. To escape charges of equivocation, the moral inferentialist has to characterize the meaning of the term such that its content remains constant across these different claims.

It’s crucial here to maintain a distinction between the pragmatic import of moral discourse and the inferential roles that enable them to fulfill that pragmatic function; we must avoid slipping “into thinking that use is at bottom *only* pragmatic significance, forgetting the use-patterns that fix conceptual content.” (Williams, 325) In making moral claims, we are not merely *doing* something but *saying* something. What we are saying—the content of moral claims—will be given by the inferential roles that fix the meaning of the terms involved. It is the *meaning* of the term and not what we do with it in moral utterances that remains constant even in embedded contexts.

Readers may notice that my proposal has a more than superficial resemblance to Ralph Wedgwood's conceptual role semantics for normative terms (2001). Like Wedgwood, I will argue that the meaning of 'ought' is "given by its essential conceptual role in practical reasoning." (Wedgwood 2006, p.142) We also share the view that "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" (Wedgwood 2001, p. 80), But Wedgwood's view also differs importantly from my own; his approach is explicitly representational.²⁶ He requires a metaphysical component to his analysis—insisting that in addition to an inferential account, "the fundamental explanation of the meaning of normative statements must refer to a particular property or relation in giving truth conditions for these normative statements." (p.40) I deny that we need to refer to such properties or relations in order to account for meaning; once we've shown how the inferential roles associated with the moral 'ought' arise from the pragmatic role it plays, no metaphysical accounting is necessary.

So the inferentialist needs an account of the inferential roles a moral term like 'ought' plays. For this, she should look to the functional roles played by the term. These are those roles that enable moral discourse to fulfill the pragmatic role I sketched above.

²⁶ Some differences are also relatively pedestrian: Wedgwood offers an analysis that is both more general and more meticulously technical than my own. Whereas he considers the inferential import of 'ought' along many of its normative dimensions (including the broadly deliberative, the instrumental and the epistemic), I limit my analysis to the import of the moral 'ought'. Wedgwood also takes great care in enumerating the technical details of his account: he argues that 'ought' should be understood as a propositional operator rather than a relational predicate, one which operates according to deontic logic and that can be modeled using possible-world semantics. I ignore such complications here; I'm primarily concerned with providing just enough detail to show that a non-representational inferentialism has the tools to address the Frege-Geach problem without falling prey to some of the criticisms that an expressivist solution must face. I draw some comparisons between Wedgwood's and my own conclusions about the inferential role of 'ought' in footnotes 10 and 13 below.

Drawing on our considerations from the last section, we should expect the term ‘ought’ to have two important functions, each of which can be associated with inferential rules for the term: helping guide coordinated behavior, and facilitating deliberation about how or why we should coordinate. In the following sections, I will briefly comment on the inferential roles that the moral ‘ought’ has to fulfill these functions; after that we will be in a position to see that the Frege-Geach problem can be solved once we recognize that the relevant inferential roles are meaning-constitutive.

(1) The action-guiding function of ‘ought’: The moral ‘ought’ is one term (among many) that brings to bear the typical action-influencing forces associated with the use of moral language, the toolkit of motivations, emotions, and behavior that (directly or indirectly) influences the actions of both the speaker and the audience of moral claims involving the term. The term carries this practical import categorically—the behavior-influencing force accompanying the moral ‘ought’ is not contingent on particular goals or desires, and it doesn’t derive its power from any particular authority.

How tight the connection between ethical ‘ought’ claims and motivations, moral emotions, and dispositions to other sort of ethically directed behavior (such as punishment) is not clear, but this shouldn’t be a problem for this inferential account. The import of such claims can be direct—as when one is motivated to φ simply by accepting one ought to φ —or indirect: it can be mediated by receptiveness to behavior-influencing moral emotions like guilt and disgust, or it can come from the behavioral influences warranted by threat of ostracism, disapproval, or other forms of punishment, or even from the pressures produced by dispositions towards second-order punishment—punishing those who fail to punish others. The practical inferences licensed by the moral ‘ought’ will

be defeasible, as when someone is being disingenuous, or suffers from *akrasia*, or has other overriding concerns, but what's crucial to the argument here is that moral claims involving 'ought' have a conceptual connection to those behavioral influences that are particular to moral discourse.

According to the inferentialist approach, the meaning of a term is given by the inferential role it plays in a claim—and the meaning of a claim can likewise be understood inferentially.

This comes, broadly speaking, to two things—(i) what circumstances license making the claim and (ii) what further statements and actions are licensed by the statement. (Chrisman 2008: 350)

This means that an inferential account of the meaning of the moral “ought” will need to establish both *introduction* and *elimination* rules for the use of the term—those rules that tell us what conditions warrant the use of the term, and what further inferences are licensed by its use, respectively. Let's consider elimination rules first.

U's assertion “one ought to φ ” licenses the inferences²⁷:

(1.1) That *U* has attitudes in favor of φ -ing, including perhaps motivation to φ , a disposition to approve of one who φ 's, and the second-order belief “one ought to be motivated to φ ”. These attitudes and dispositions reinforce the categorical nature of the moral 'ought'; their practical significance is not contingent on the desires or goals of particular agents.²⁸

²⁷ There are two schools of inferential explanation available. Inferentialists like Horwich (1998: see pgs.45-47) advocate a “naturalistic” account, in which the meaning of a term comes from our disposition to accept particular instances of it in given sentences. Inferentialists like Brandom, on the other hand, explain the meaning of terms normatively, by adverting to the rules that govern its use—those socially-embodied commitments and entitlements that accompany our use of a term. With a little imagination, one could derive a naturalist version of the inferential rules for 'ought' from the ones that follow below. Here, though, I side with the normativist approach.

²⁸ Cf. Wedgwood on the inferential import of the broadly deliberative 'ought', which he argues is always implicitly or explicitly indexed to some agent and time: “Acceptance of the first-person statement ' $O_{\langle me, t \rangle}(p)$ '—where ' t ' refers to some time in the present or near future—commits one to making p part of one's plan about what to do at t .” (2006, p.137) Ignoring questions of indexing time and agent, we can see the similarity of his account to my own—if we keep in mind that mine is narrowly concerned with moral 'ought's, and don't flesh out the notion of a commitments and plans in psychological terms (i.e. of commitments and plans as kinds of mental states) but instead in terms of the kind of pragmatic influences

(1.2) That U has similar attitudes against not φ -ing, including perhaps a disposition to feel guilt upon failing to φ , disapproval or disgust of not φ -ing, perhaps to the point of effecting punishment and even pressure on third parties to react punitively to those who fail to φ , and the second-order belief “one ought to disapprove of not φ -ing.”

These licenses can be understood, following Brandom, in terms of the social significance of making an assertion.

In asserting a claim one not only authorizes further assertions [which we will see below], but commits oneself to vindicate the original claim, showing that one is entitled to make it. Failure to defend one’s entitlement to an assertion voids its social significance as inferential warrant for further assertions. (Brandom 1983: 641)

So by asserting, “You ought to φ ”, I’m holding myself liable to challenges. If I make such an assertion, but don’t evidence any of the attitudes or commitments that should attend it, I’ve opened myself to criticism. A complete inferential account will likewise specify the introduction rules that indicate the appropriate circumstances for use of the moral ‘ought’:

(1.3) U ’s issuing a categorical command to ψ to a person S —that is, a command that is not intended to have its force contingently, and that disposes U to approve of S ’s ψ -ing, to disapprove or feel disgust at S not ψ -ing, etc—licenses U to claim that S ought to ψ .

This kind of explanation might raise a red flag. In arguing for an inferential understanding of the moral ‘ought’, I’m making use here of the kinds of pro- and con-attitudes that are the hallmark of an expressivist account. One might suspect that I’ve slipped back into an expressivist mode of explanation. To understand the inferential import of moral claims, we must look past the varied mental states that are generally correlated with sincere instances of moral claims, to the practical impact that these claims typically have.

(i.e. towards inferential commitments to arrive at deliberative conclusions about how to act) enumerated in Section 5 above. Cf. Charlow (2013, fn. 20).

Perhaps sincere ethical claim-making is indeed... correlated with having the sorts of desires, intentions, and plans which explain motivation, but the exceptions to this psychological correlation indicate that it shouldn't dictate our explanation of the meaning of the relevant claims. We can recognize that certain sorts of ethical claims conventionally express desires, intentions, and/or plans without holding that they mean what they do in virtue of expressing those things. (Chrisman 2010, p.119)

The attitudes that attend these claims can shift: sometimes the payoff of a moral statement is a direct motivation; sometimes it's guilt at a past transgression; sometimes it's fear of punishment. This flux of attitudes serves an underlying function—to promote or discourage particular types of behavior. And it is this function—not the attitudes that serve it—that explains the inferences that are licensed by utterances involving the moral 'ought'.

(2) Using 'ought' to reason to and between practical conclusions: Like imperative language, the moral term 'ought' influences behavior, but it as I argued above, it differs from imperative language importantly: the moral 'ought' can be used in the indicative form, which enables complex reasoning about the appropriate application of the above behavior-influencing forces. For example, it enables us to express permissions—that is, it enables us to abrogate those typical motivational forces associated with the use of the moral 'ought'. This function of the term is enabled by the following inferential rule connecting obligations to permissions:

(2.1) $\sim O(\sim\phi)$ iff $P(\phi)$: 'Ought' can be used to license inferences to conclusions that involve the motivating features of moral discourse, but because we can use it with negations, we can also use 'ought' to *block* certain kinds of inferences. "One ought to steal" licenses the use of practical force in favor of stealing. It can be contradicted in two ways: "One ought to not steal" licenses practical force *against* stealing; "It's not the case that one ought to not steal", though, undermines such force—it gives *permission* to steal (more on this below).²⁹

²⁹ Compare again Wedgwood's treatment: "Acceptance of the first-person statement ' $P_{\langle me, t \rangle}(p)$ '—where ' t ' refers to some time in the present or near future—permits one to treat p part of one's plan about what to do

Now that we've established introduction and elimination rules for 'ought' and the language-language rules that relate 'ought' to 'is permissible that', and note that these can be placed outside of any well-formed sentence p , such that $O(p)$ and $P(p)$ qualify as assertions, we have a *prima facie* case that these terms act as sentential operators, and as such we can use them with standard sentential connectives to make complicated moral sentences. The validity of arguments involving those connectives would then follow from the meaning of the connectives themselves. Thus, we could explain the validity of any inference of the following forms

$$\begin{aligned} p \rightarrow O(q); p; \text{ therefore } O(q) \\ O(p) \ \& \ O(q); \text{ therefore } O(p) \\ O(p); \text{ therefore } O(p) \vee O(q) \end{aligned}$$

trivially, as consequences of the introduction and elimination rules that come with the logical connectives. This would get us a "fast-track" solution to the Frege-Geach problem—one that shows there is "sufficient similarity of logical role" between evaluative and representational claims "to make the temptation to exploit ordinary propositional logic quite irresistible—and that is what we naturally do".³⁰ (Blackburn 1998: 517) A solution to the Frege-Geach problem would come easy.

But to many, this will seem far too fast. We may wonder if it's been established that O and P do actually act as sentential operators, and not just act a lot like them. We may worry that this account overlooks something. In the next section, I want to give voice to this worry and consider some ways to respond to it. I'll then return to the task of explaining the inferences that are licensed by the moral 'ought'.

at *t*." (2006, p.137) Notice also that this gives us one of the inferential rules that is meaning-constitutive for the moral term 'permissible'.

³⁰ For a comparison of the merits of "fast-track" and "slow-track" solutions, see (Blackburn 1998).

6. The Commandos and the Task of Accommodation

There was once a society of people, called Commandos, who were much like you and me, except for one linguistic oddity. Instead of using grammar to indicate that they were giving one another commands, they just used the phrase "it is commanded that". So instead of saying, "Open the window", a Commando would say, "It is commanded that you open the window." Now, if someone said, "It is commanded that you open the window" in English, this would be true just in case such a command was given. But let us stipulate that Commando sentences don't have this function: they don't report on truthmakers, but instead function solely as devices for issuing imperatives. And this way of talking seemed to serve the Commandos fine³¹.

But the Commandos went astray. For once logicians realized that you could add "it is commanded that" to a well-formed sentence to make a new sentence—one wherein one commanded that the content of the original sentence be made to be the case—it became popular to employ "command" language in the same way we'd use a sentential operator

³¹ To those familiar with Jamie Dreier's treatment of embedding, it will be obvious that this borrows liberally from his 'hiyo' story. Dreier introduces a predicate, "hiyo" into our language. As a speech act, the predicate is used to accost someone; "Bob is hiyo" gets Bob's attention. What's more, "hiyo" is a predicate that can be used in claims with minimalist truth ascriptions, so that "It's true that Bob is hiyo" has the same meaning as "Bob is hiyo". Importantly, Dreier declines to tell you what property it denotes. In the robust sense of property, it does not pick one out at all. In the Minimalist sense, it picks out the property that an object has just in case it is hiyo. (Dreier 2006: 42)

By stipulation, "hiyo" is not a predicate that operates representationally.

But this account on its own will not make sense of assertions like "if Alex is hiyo, then Bob is hiyo". Just because our syntax allows for such constructions doesn't mean that such constructions mean anything. Our attempt to use hiyo claims minimally shows us that:

Not every combination of characterization of the speech act performed by an utterance and stipulation that the expression in question functions as a predicate is a consistent combination. We want some sort of guarantee that the stipulation makes sense. (44)

The same lesson holds true of normative predicates in ethical claims.

like, “It is necessary that”. Sometimes this was fine, like when Commandos reached certain kinds of conclusions:

If the window is open, it is commanded that you shut it.

The window is open.

Therefore: it is commanded that you shut it.

When the Commandos discovered deflationary theories of truth, trouble really started brewing. "Adding the truth predicate doesn't add any content to a sentence," they reasoned. "'Snow is white' is true just iff snow is white. So we can say the same about what is commanded: 'It is commanded that you shut the window' is true just iff it is commanded that you shut the window." This led to the Commandos engaging in all sorts of weird command-based reasoning, like

I think, 'It is commanded that you shut the window' is true, but you think it's not true. Given the law of non-contradiction, one of us must be wrong.

And strange instances of “validity” in arguments like:

If it is commanded that you shut the window, then it's cold outside.

It's not cold outside.

Therefore, it is not commanded that you shut the window.

It's worth stressing again that such inferences make a kind of sense in English, since in our language “it is commanded that” can be added to a well-formed sentence to produce a new sentence that reports on aspects of the world—it's true just in case a particular command has been issued. But these inferences are bizarre given the above stipulation about Commando language: we can't reason with imperatives as if they were true; and we can't meaningfully embed them in the antecedent of conditionals.

As one might expect, a lot of philosopher Commandos came to worry about the nature of what is commanded. For those of a representational bent, there was a lot of work done hunting for truthmakers for these claims. But then one day along came a non-

representational Commando who suspected that this way of philosophizing about what is commanded was wrongheaded.

"Hunting for truthmakers for command sentences is wrongheaded," she reasoned. "What we need to do is give an inferential account of the meaning of 'it is commanded that' based on the function that command talk plays in Commando life. Let us be lightweight about commands." Command discourse, she reasoned, plays the role of making it clear to a listener that the utterer wants something to be the case, and more specifically, wants the listener to make it the case. Given the pragmatic role of this discourse, she thought, it makes sense that command claims follow certain inferential rules. We can embed propositions of many kinds within the scope of 'it is commanded that', because that expedites the role of command discourse. And we can use 'it is commanded that' in the consequent of a conditional, because it's useful to make commands that are contingent on certain states of affair for their force.

Our non-representational Commando was shocked, though, to realize that some of the inferential rules that her fellow citizens gave 'it is commanded that' didn't make any sense at all. The truth-predicate, for example, should never have been added to command sentences, because using 'truth' in this way instigates disputes about commands when no dispute is necessary. Two people can issue different commands without either of them being in error! And furthermore, there were ways that people used command language in reasoning that just made no sense at all—again, how can we make sense of a command embedded in the antecedent of a conditional? The conclusion she reached was that the ordinary ways that Commandos had of thinking and talking about 'is commanded that' were illegitimate. The discourse was diseased and needed to undergo radical revision.

The moral of the story is that for an inferentialist like myself, the project of accommodation is showing that moral language isn't faulty and misleading in the way that the Commandos' command language is. Just because something looks like a sentential operator and acts like a sentential operator, we can't rest assured that it's going to act like the other sentential operators we're used to. A thorough analysis of the inferential rules that 'it is commanded that' follows will show that some types of inference make sense, but others don't. To assure people that moral language is not illegitimate in the way Commando language is, it must be shown that it makes sense to treat moral claims in different embedded contexts, and one can't rely on the appearance of operatorhood or truth-aptness of the moral 'ought' to accomplish this task. We'll have to take the "slow-track" to accommodation, building up our account of moral inferences piece by piece.

For those attracted to the "fast-track" solution to the Frege-Geach problem, there are a couple of possible responses to this story. The first is to deny that I've offered a true translation of "it is commanded that". If these people are consistently using the phrase in the way I've described, it seems inconsistent to insist that the phrase is used in the same way we would use imperative language. Interpretative charity demands I give some other translation—perhaps "It is commanded that" is Commandese for "you ought to", or something related.

But if we take the project of accommodation seriously, we can't simply rest assured that all uses of language are sensible, just because a group commonly employs them. If this was our assumption, the Frege-Geach problem would have never gotten any traction: what's interesting about normative language used in force-stripping contexts is

that on at least some accounts of the semantics of normative language, it *doesn't* actually behave the way we might assume.

Another related option is to flatly deny that the story is plausible, because Commandos wouldn't be able to make sense of sentences that begin with "it is commanded that" as assertions. The problems I adverted to are exactly the sort of thing that would keep the language from ever getting off the ground. This seems like a reasonable response, but we want to be sure we don't put the cart before the horse. We'd need some plausible account of what can and can't count as assertions in the first place. And if we're non-representationalists, we'd better hope that the distinction doesn't have to do with reporting on truthmakers.

Perhaps such an account is available, though. Recall that on an inferentialist account, the meaning of an assertion can be read from the entitlements and obligations that follow. Now, imperative sentences have entitlements and obligations of their own; if I tell someone, "Shut the door!", she is entitled *ceteris paribus* to infer that I want the door closed, and I have an obligation to not then react with outrage when she shuts the door. But as we'll see, these "upstream" and "downstream" commitments will be different in kind from those that follow from assertions.

The commitments that attend an assertion will transferable, in a sense, because

Putting a sentence forward in the public arena [as an assertion] is something one interlocutor can do to make that sentence available for *others* to use in becoming entitled to further assertions. (Brandom 1983: 461)

Assertions are agent-neutral in this sense; whether or not a particular assertion is warranted doesn't depend on who has made it.

Again, the idea is that if you make an assertion, like, "Drinking from that stream will make you sick", then in doing so, you warrant other assertions (e.g., "If he's sick, it

might be because he drank from the stream”), but you also make yourself liable to justify your assertion if someone asks you to (“Why do you say that?” “Well, I saw Todd and Anna drink from it yesterday, and they’re both sick today”). And once you’ve justified the assertion, the warrant you’ve given it transfers to your interlocutor—she can then use the assertion to license inferences of her own, and she can transfer whatever warrant you gave the claim if another interlocutor demands it. “An assertion makes the same claim regardless of who is speaking and who is listening.” (Lance and Kukla 2010: 117)

But whatever warrant one has in saying “It is commanded that you shut the window” won’t carry over in this agent-neutral way. The import of such a command is deeply connected to who is making it, and who is hearing it. That’s why imperatives don’t have the structure of assertions.

An imperative must be issued to someone in order for it to count as an imperative at all. The idea of ‘translating’ an imperative into the third or first person while retaining its meaning or force does not even get off the ground conceptually. “Close the door!” makes a specific demand upon someone in particular (or some particular group of people). A ‘translation’ into the third person, such as “Mark ought to close the door”, is a declarative with a totally different pragmatic structure; it does not constitute an order at all. Nor is there any clear sense in which it ‘means’ the same thing as the original imperative. (Lance and Kukla 2010: 117-118)

This makes sense of why we can embed imperatives in the consequent but not the antecedent of a conditional: we can understand commands being conditional on a state of affairs that can be warranted from an agent-neutral perspective, but the command itself (as opposed to a state of affairs in which a command has been given) can’t count as a potential *situation* from which we can reason.

Such an approach should be consoling to the ethical inferentialist, since ‘ought’ claims are ones that have this sort of agent-neutral implications. Notice also that this puts us in a good position to return to the sort of explanation of the meaning of moral claims

that was inspired by Williams: we can make sense of the kinds of inferences complex moral claims license by reference to the pragmatic role that they expedite. Those drawn to a fast track approach might feel it's not necessary to do any more work at this point, but certainly no harm can come from soothing any qualms those who would take the slow track approach might have—by shoring up this approach with an extension of the kinds of inferential roles we should expect to follow from the role that 'ought' plays in our ability to reason to practical conclusions:

(2) Using 'ought' to reason to practical conclusions (continued): The moral 'ought' enables us to reason about the application of those practical forces associated with moral discourse—what situations call for these forces, given the aim of coordinated behavior. It also enables us to reason from and among our principles, determining what follows and which motivational forces to bring to bear, given other principles

(2.2) $[(\gamma \rightarrow O(\phi)) \& (\gamma)] \rightarrow O(\phi)$: We can use the term 'ought' in the indicative, and so inferences like *modus ponens* enable us to reason from a conditional principle (like, "If a man is wicked, he ought to be punished") to an application of the behavior-influencing forces limned above, given the appropriate situation.

(2.3) $[(O(\psi) \rightarrow O(\phi)) \& O(\psi)] \rightarrow O(\phi)$: We can also use 'ought' to reason between our principles, making inferences from conditionals involving two principles to a direct application of the practical forces associated with moral discourse, given the antecedent principle. This means for example that "If it's wrong to lie, it's wrong to get your little brother to lie" and "It's wrong to lie" taken together license the inference to "It's wrong to get your little brother to lie".³² (This also shows that we can make sense of the moral 'ought' being embedded in the antecedent of a conditional, in a way that we can't with the Commando's "It is commanded that".)

³² No doubt there are other primitive inferences that are licensed by the term 'ought', familiar to students of logic, like *modus tollens*, the conjunction introduction and elimination rules, and so on, but for this sketch of moral inferentialism, I'll limit myself to a few important inferential roles. As we'll see below, a complete account of the inferences licensed by the term 'ought' might need to include inferences involving quantification, predication, tense, modals, and so on, but such a treatment lies outside the scope of this chapter.

7. Inference from Function: Solving the Frege-Geach Problem

These inferential roles associated with the moral ‘ought’ are justified by the pragmatic aim of moral discourse; we can see how these roles associated with (1.1-1.3) and (2.1-2.3) above enable social-coordination:

(1') Obviously, for moral discourse to constrain our behavior, we will need at least one term that influences our actions. Because the moral ‘ought’ carries behavior-influencing force that is categorical, its effect on our behavior is stable—the import of an ‘ought’ claim is not contingent on any particular desires or authoritarian commands. Because the term influences moral emotions, and because it can inform liability regarding punishment and reward, it can weigh against the selfish motivations of those whose behavior is not swayed directly by moral considerations. Finally, ‘ought’ claims have implications for utterers’ behavior, and the expectations others have regarding this behavior is essential for coordinating behavior.

(2') This practical force should be situation-specific. In order for moral language to effect *complex* coordinated behavior, we need at least one term that allows us to reason from, between, and to action-guiding principles. For moral language to effectively facilitate coordinated social behavior, we need to know not only when to employ its typical motivational force, and which principles follow from which, but also when to not employ it—introducing the notion of permissibility into moral discourse makes available a fine-grained application of its practical import that is otherwise out of reach.

Because the moral ‘ought’ has these inferential roles whether it is embedded or unembedded, its meaning remains constant and so the charge of equivocation is inappropriate. Consider again the inference above regarding the punishment of a cheater. This is an instance of *modus ponens* that is the paradigmatic challenge of the Frege-Geach problem. Here, ‘ought’ has the same inferential roles in both premise and conclusion. Following the inferential role (2.2), it enables us to reason from a principle calling for behavior-influencing force in a particular situation (the conditional) that, combined with an understanding that one is actually in that situation, warrants an inference to a conclusion (thereby fulfilling function (2)) that brings the relevant practical force in to play (thereby fulfilling function (1)). The meaning of ‘ought’ is jointly constituted by its inferential roles; these stay constant, and so too does the meaning. Of course, which inferential role is most relevant in a particular inference might change, but this doesn’t affect the meaning of the term.

The inferentialist account of validity here is straightforward. Where the representational approach starts with the notion of truth and explains valid arguments as ones that are necessarily truth-preserving, and where the expressivist starts with the notion of attitudes being expressed and explains valid arguments as ones that are guaranteed to not produce jointly unsatisfiable or unrealizable mental states, the inferentialist begins at the outset with the notion of correct inferences—those that a competent language-user will be expected to have mastered in virtue of understanding the meanings of the premises and conclusion—and explains valid arguments as ones wherein one is committed to the conclusion simply in virtue of being committed to the premises^{33,34}.

³³ Cf. (Brandom, 1994, p.188)

The same sort of explanation is available for inferences from premises that are exclusively normative. For example:

If you ought to keep people from starving, then you ought to give to charity.
You ought to keep people from starving.
 So you ought to give to charity.

Here's another instance of *modus ponens*. The validity of the above inference here comes from the conceptual role the term 'ought' fulfills in (2.3) above—enabling reasoning not only from situation to principle, but also between principles themselves.

The sketch above fits the desiderata I listed at the beginning of the last section, for an adequate response to the challenge of the Frege-Geach problem:

1. It does not presuppose any domain of elusive moral facts that are waiting to be described or responded to—it is a lightweight account. Instead it makes use of a naturalistically respectable understanding of the behavior-guiding function we might expect organisms like ourselves to use our moral language to fulfill.
2. The inferential rules licensed by the conceptual role played by 'ought' gives us the meaning of the term in such a way that it is clear this meaning remains constant even in force-stripped contexts, so that the charge of equivocation is evaded.
3. Identifying the inferential role of moral terms helps us understand the validity of inferences that we make involving the moral 'ought'. These inferences are not

³⁴ For those who insist that an adequate account of validity must involve truth-preservation, the inferentialist can be confrontational or conciliatory. Confrontationally: There is good reason to think that the relationships of consistency and inconsistency that are fundamental to an explanation of validity needn't be underwritten by a truth-theoretic account; following Portner (2010) and using a meaning-from-use approach not dissimilar to my own, Charlow (2013) argues persuasively that we can (and should) extend a non-representational account of (in)consistency between imperative sentences to declarative normative claims. Conciliatorily: This shouldn't be taken as an outright denial that validity (in moral arguments like the above, at least) entails truth-preservation, though: inasmuch as the inferentialist makes use of a minimal approach to truth, she can grant that valid arguments with true premises are guaranteed to have true conclusions.

understood as valid solely because of the truth-conditions of the premises and conclusions. Instead the inferences are licensed by the use we make of moral language.

Having outlined how an inferentialist account can solve the Frege-Geach problem, I want to argue that it also gives us the tools to deal with at least some of the worrying criticisms that have been leveled at the standard expressivist solution. Anyone familiar with contemporary literature on expressivism knows these are numerous; I'm focusing on two that have proven particularly malignant: the problem of permissions, and the expressivist commitment to what Schroeder calls 'mentalism'. I'll briefly explain each, and show how an inferentialist account can either avoid the problem altogether, or at the very least render it more manageable.

8. The Problem of Permissions

The standard expressivist response to the Frege-Geach problem involves understanding the validity of an inference in terms of the logical relationships that inhere between the mental states expressed by the different parts of the inference—an inference is valid if the mental states expressed by the premises in some way entail the mental state expressed by the conclusion. Likewise, two sentences contradict one another when the mental states they express are inconsistent. So consider for example two moral claims that are obviously contradictory:

- 4) Voting for Obama is wrong.
- 5) Not voting for Obama is wrong.

(4) expresses some mental attitude towards voting for Obama—let’s say a plan to disapprove of voting for Obama. And (5) represents the same attitude towards *not* voting for Obama. The two claims contradict one another, on the expressivist account, because one cannot consistently plan to *p* and plan to *not p*; one could not consistently have the *same sort of attitude toward contradictory content*.

Such an explanation has an advantage: the notion of contradiction between mental states here is one that is not specific to expressivism. We can think of other instances where commonsense dictates two mental states contradict one another *because they are the same sort of attitude toward contradictory content*—a belief that *p* contradicts a belief that *not p*, an intention to *q* contradicts an intention to *not q*, and so on.

Nicholas Unwin (1999, 2001) points out there is more than one way to contradict (4); how are we to make sense of the contradiction between (4) and the following?

6) Voting for Obama is not wrong.

Obviously, the expressivist must account for the contradiction between (4) and (6) in terms of the mental states the two claims express, but (6) does not express the same mental state as (5). And of course, we cannot simply explain the mental state (6) expresses as a kind of refusal to accept (4), since, as Unwin points out, “[T]here is certainly a very real difference between not accepting (or refusing to accept) something and actually accepting its negation.” (Unwin 1999, p.341) It’s very tempting to respond that (6) doesn’t express the same kind of mental state as (4) and (5)—two mental states of the same kind with contradictory content—but instead expresses a different type of mental state altogether: *tolerance*. The idea is that the contradiction between (4) and (6) doesn’t come

from having the same kind of attitude towards contradictory content, but from having *different kinds of attitudes to the same content*.

We saw in the last chapter that there is a problem with explaining this kind of contradiction with B-type inconsistencies between mental states. The difficulty with using such an approach is that it leaves unexplained *why* different attitudes like tolerance and disapproval would be related to one another logically. Common sense gives us a model of logical inconsistency between type-identical mental states that have inconsistent content, but we'd need a new model to explain mental states that are inconsistent because they involve different attitudes towards the same content. The expressivist could of course just stipulate that claims expressing attitudes of disapproval and tolerance have the inferential relationships they do, but making such a decree doesn't do any work to explain why those relationships inhere, or what sort of logical relationships we should expect them to have. "It becomes at best a mere brute fact that the attitudes conflict with each other, with no internal complexity that could explain why." (Ibid, p.342)³⁵ The worry is that "psychological states are not endowed with the right sorts of properties to offer self-sufficient explanations"(Charlow 2013, p.20) of why these states are inconsistent.

At the beginning of my account of inferentialism, I argued that one of its virtues is its simplicity; all things being equal, one should prefer a semantic account that directly explains those inferential roles that are constitutive of the meaning of a term over an account that explains meaning by reference to mental states which themselves provide the appropriate inferential roles. Now we can see that such a virtue might pay dividends:

³⁵ Alternately, one may try to account for the distinctive inferential role of tolerance by giving a "Higher Order Attitude" account (see Blackburn 1984, chapter 6) of the state in terms of disapproval, so that to tolerate ϕ is to disapprove of disapproving of ϕ . Such an account is troublesomely *ad hoc*, and has perhaps rightly come into disrepute (see Schueler 1988, van Roojen 1996.) See Dreier (2006) for a different approach and its attendant problems.

because the inferentialist does not rely directly on (in)consistency between mental states to account for inferential relationships, she doesn't need to posit tolerance as a special sort of mental state with its own unique permission-giving inferential features. She can instead account for the inferential role of 'permission' itself to explain its fundamental inferential features. Here the inferential role (2.3) outlined above is fulfilled. To permit behavior is (among other things) to withhold the practical forces associated with 'ought'; with negation this can be expressed by saying that "is it not the case that you ought not p ", or simply " p is permissible".

With this account of the inferential role the term plays in hand, we can understand the different kinds of negation that have vexed expressivist accounts. (4) and (5) above contradict one another because (4) invokes behavior-influencing force against voting for Obama whereas (5) invokes the same force against not voting for Obama. But (4) and (6) contradict one another because (6) undermines the very force that (4) invokes. Where (4) invokes the characteristic behavior-influencing force of 'ought' against voting for Obama, (6) invokes another sort of behavior-influencing force, which can at least in part be understood in terms of the kind of typical reaction someone voicing such a permission could reasonably be expected to have to someone who says we shouldn't vote for Obama. More generally, if U asserts "it's permissible to φ ", this implies (among other things) that U lacks dispositions to feel guilt upon φ -ing, disapprove, punish, or ostracize someone who has φ -ed—and furthermore this implies U has the second-order commitment that one ought to not disapprove of φ -ing, where such a second-order belief would have its own dispositional consequences (e.g., to chide those who say things like "one ought to not φ ").

The notion of negation here relies on how the two claims reflect dispositions that pragmatically clash.³⁶

This explanation is not a bare stipulation about the inferential role of “is permissible”, and it’s not an *ad hoc* response to the problem. Rather, it follows naturally from our considerations of the pragmatic function of moral discourse. We use moral language to coordinate effective social behavior with terms that influence behavior with motivating force. The inferential role of ‘is permissible’ gives us fine-tuned control of this force—it allows us to indicate where such force is to be used, and where it is to be withheld. Because she grounds the semantic content of moral claims in the inferential role played by the terms themselves—rather than in the mental states the terms serve to express—the inferentialist doesn’t have to stipulate that tolerance is a mental state with unique inferential properties, and so doesn’t have to account for such properties.

9. The Commitment to ‘Mentalism’

Recall from Chapter Three that an expressivist will want to give a univocal treatment of the sentential connectives:

The inconsistency between ‘P’ and ‘~P’ should be guaranteed by the meaning of ‘~’, not by the joint fact that if we first interpret ‘P’ as normative and ‘~’ as normative-sentence-negation, they are inconsistent, together with the fact that if we interpret ‘P’ as descriptive and ‘~’ as descriptive-sentence-negation, they are inconsistent, and if we interpret it in any other way it is not well formed. So if we hope to get very far at all, the pressure to have a uniform account of the sentential connectives is very high. (Schroeder 2008a, p.22)

³⁶ For more on cashing out negation in terms of pragmatic conflict, see Dreier (2009), especially p. 103-106.

That led us to conclude that these connectives play the *same* fundamentally expressive role wherever we find them—whether they connect evaluative claims or purely representational claim. And as we saw, this lends plausibility to the notion that the meaning of all claims—normative or representational—must be understood as the expression of the same fundamental type of mental state.

Inferentialism distinguishes itself from expressivism on precisely those grounds that lead to this troubling charge of mentalism, so one might suppose that the charge has no weight here. After all, if the inferentialist doesn't account for the meaning of moral claims in terms of the mental states they express, there's no reason to think that inferentialism needs to account for the meaning of logical connectives or descriptive claims in the same way. But even though the specifics of Schroeder's charge get no traction in regards to inferentialism, we may be worried that an argument of the same general form applies.

On the inferentialist account, the meanings of moral claims come from their inferential role. It follows that the meanings of the sentential connectives would also be given as a function of (not the mental states but rather) the inferential roles played by the claims that they combine. This in turn implies that the meaning of all atomic claims—normative or not—that can be combined with moral claims by logical connectives are to be given by their inferential roles! Though such a view wouldn't rightly be called 'mentalism' it is nonetheless a non-trivial position about fundamental semantic issues.

The conclusion is compelling, but not nearly so troubling as the charge of mentalism. It may be non-trivial, but it certainly isn't entirely novel; there already is a plausible, independently motivated theory of linguistic meaning committed to such a

conclusion: conceptual (or inferential) role semantics (see for example Brandom 1994, Horwich 1999). As we've seen, according to this semantic approach, the content of an utterance is determined by its use—that is, the inferential role it plays. It may come as a surprise that metaethical inferentialism entails a general inferentialist account of meaning for the logical connective and for non-normative claims, but even if it does, the account finds itself in good company. Schroeder's semantic challenge may not be entirely undermined, but it is defanged.

10. Where This Leaves Us

It has been my intention in this chapter to sketch a new way to respond to an old problem, in the hopes that this will motivate a more complete project of explaining the semantics of moral language in inferential terms without depending on the expression of mental states. I've very briefly sketched two of the problems that expressivism inherits because of its dependence on such states, and inferentialism's advantages in dealing with these. A point-by-point reckoning with all of the semantic issues plaguing expressivism lies outside the scope of this chapter, but is certainly needed for a full account.

One would obviously also need an account of the myriad other moral concepts we use, concepts like 'good' and 'bad', but also thick moral concepts—those which include both evaluative and descriptive elements—like 'brave' and 'miserly'. Astute readers might also notice that my treatment of 'ought' is incomplete, for two reasons. First, 'ought' isn't always used morally; presumably a complete account will integrate our

moral and non-moral uses of the term. Second, we often use the moral ‘ought’ in contexts where the behavior influencing force of moral language doesn’t obviously have a place. What sort of behavior are we trying to effect when we say, “Stalin oughtn’t have been so devious”? And what are we to make of moral sentences that don’t directly involve any agents at all, as in, “This shopping mall ought to be more accessible to the disabled”?³⁷ Speculatively, it seems the inferentialist has two options: Either seek a reduction of such claims in terms dissectible via the pragmatic function of moral discourse I listed above, or expand on the list to account for such uses. Assessing the relative plausibility of these options is a project for another time.

And of course we can’t forget those other semantic problems expressivism faces that I referenced earlier but haven’t addressed: G.F. Schueler argues that non-cognitive mental states like approval and disapproval just aren’t plausible candidates as manufacturers of genuine logical contradiction. Mark van Roojen points out that the expressivist account of validity in an argument seems to overproduce valid arguments; it gives us a model that attributes validity to arguments that are obviously invalid. Mark Schroeder points out that even if the expressivist were able to give an adequate account of simple moral inferences, she would still owe an account that captures all of the ways we use moral terms—in conjunction with attitude verbs, binary quantifiers, tense, modals, and so on. Perhaps some of these issues will dissolve when we take up the inferentialist perspective; all of them are critiques aimed at expressivism, after all, and involve problems that spring from the mental-state-expressed account of content that

³⁷ For Wedgwood’s suggestions on giving a more general treatment of ‘ought’, see (2006, p.151-157). As I’ve explained in Section 6, my account of ‘ought’ differs from his in important ways; nevertheless, I see no reason that his account of the various implications of ‘ought’ couldn’t be favorably extended to a non-representational approach such as my own.

inferentialism rejects, so one might hope they simply don't apply (this seems especially promising regarding Schueler's objection). Still, one could be worried that some or all of these criticisms can be leveled in an inferentialist key—indeed, that they might arise for *any* non-representational account of content.³⁸ Responding to all of these concerns on behalf of inferentialism is another deep project.

I think these projects hold real promise, but of course the devil is in the details. Much work needs to be done to convincingly show that the approach resolves these issues. Because inferentialism offers us a simpler view of moral semantics, and because it promises to skirt many of the semantic issues that plague expressivism, I believe these are projects that are worth our attention.

The inferential account I argue for above also sheds light on my approach to objectivity in Chapter 2. Again, one advantage of inferentialism here is that we needn't concern ourselves with an account of mental states that are “the kind of things we can disagree with”, because inferentialism looks past mental states as semantic *explanans*. But we still do need a feasible account of objectivity in terms of the inferential import of moral claims. This will be the same tune I offered in Chapter 2, but played in an inferentialist key; the challenge is to explain why there are genuine moral conflicts, in terms of the practical role moral discourse plays in our lives. I'll argue that here too non-representationalists should be concerned with making sense of a particular inferential

³⁸ Addressing these issues lies outside the scope of this chapter, but I expect that both Schueler and van Roojen's objections to expressivism won't get much traction regarding inferentialism, for just the reason I suggested. Mark Schroeder's issues with the semantics of expressivism are a more serious challenge, though one promising line of response comes from embracing a generally deflationary theory of truth, meaning, and logic (see Brandom 1994), and then arguing that because moral claims can be seen as truth-apt statements on equal footing with non-normative claims, we are entitled to reason with them in all the standard ways. Those tempted by the fast-track approach to these issues will find this line of response promising, though of course the same sorts of considerations above will apply: we may want independent reasons for thinking that our ways of reasoning with moral claims in these contexts make sense.

norm that governs our moral discourse: that instances of genuine moral disagreement entail error on the part of at least one of the parties to the dispute. Here too we'll see that an inferential approach enjoys some advantages over its expressivist counterpart.

CHAPTER FIVE: OBJECTIVITY FOR LIGHTWEIGHTS

1. Overview

In Chapter Two, I argued that a non-representational approach to metaethics must take a *pragmatic* path to questions of objectivity. A ‘heavyweight’ representational approach can make sense of objectivity by describing some unified set of properties or relations that stand as truthmakers for our moral claims. On this account, morality is objective because for any instance of genuine moral disagreement, there must be some sort of error on the part of at least one of the disputants, and this feature of moral discourse is explained by the coherent set of heavyweight truthmakers that moral claims aim to represent; this is because (barring dialetheism) the realm of moral truths is governed by laws of non-contradiction and the excluded middle. That is one way to make sense of this feature, but as we saw, it is not the only way. The ‘lightweight’ non-representationalists instead explain why moral disagreement entails error by adverting to the pragmatic aim of moral discourse itself: we must treat moral claims objectively if moral language is to serve its role of driving coordinated social behavior. Moral disputes need resolution, not because they are disputes that are primarily about the way things *are*, but instead because they are disputes about what we are to *do*.

The advantage of taking this sort of pragmatic turn is that it allows us to avoid the ontological and epistemic debts we’ve seen associated with the heavyweight realist projects. The ontological debt is to give an account of the truthmakers in a given

discourse, hopefully in a way that comports with a plausible naturalism. How, for example, can moral facts be naturalistically respectable, when unlike any other natural facts, they seem to have a certain prescriptive ‘to-be-doneness’ built in to them? (Mackie 1977) If such an account can be given, the heavyweight must then explain how it is that primates like ourselves can come to develop a sensitivity to these special sorts of fact, and so come to know about such truthmakers. Because the non-representationalist doesn’t rely on heavyweight conceptions of the facts in question, she doesn’t incur the ontological or epistemic debts that plague traditional representational approaches. Being a lightweight has its advantages.

Nevertheless, the lightweight approach has its own distinctive debt to pay: the task of accommodation still remains. How can we establish that even without recourse to truthmakers, our ordinary ‘realist-seeming’ thoughts and practices regarding morality are legitimate? If she is to claim objectivity for her metaethics, the lightweight realist must make sense of a broad range of moral discourse. I argued that this task was particularly hard when we considered all the sorts of moral disagreements people have that don’t obviously cash out in terms of pragmatic conflicts about what to do. This is Jamie Dreier’s “Ibuprofen challenge”, which we saw in the final section of Chapter Two, and it extends to ethical disputes about impossible counterfactual choices and questions about the moral implications of past choices. For the expressivist lightweight, this challenge of accommodation must be answered by citing the mental states that are expressed by inconsistent moral claims, such that these mental states bring disputants into practical conflict. As we’ve seen, this is difficult to do.

But as we've also seen, one advantage of a lightweight *inferentialism* about morality is that we needn't concern ourselves with an account of mental states that are the kind of things we can disagree with, because inferentialism looks past mental states as semantic *explanans*. Now the task of accommodation will be to understand objectivity in terms of the inferential import of moral claims. In particular, it will be to show that the inferential norms that govern moral discourse are ones that *make sense* in our everyday life. In this chapter, I will show how this can be done.

I'll first offer an inferential role translation of some of the key expressivist moves towards objectivity we saw in Chapter Two. The most important of these will be an inferential rule that explains why there must be error in instances of moral disagreement: the fundamental task of justifying objectivity in moral discourse will amount to making sense of this inferential rule as it applies to a broad range of possible moral disputes. One way of doing this is by making use of a minimal account of truth; in Section 3 I'll argue this approach isn't adequate. In Section 4 I'll argue that our pragmatic considerations must go deeper—that we have to look at the pragmatic role of disagreement itself. This will put us in a position to make sense of problematic moral disputes like the Ibuprofen challenge, and it will give us insight into the nature of disputes in other discourses, like modality and aesthetics. In the penultimate section of this chapter, I'll leave the reader with some considerations about the limits of objectivity for a lightweight realist. I'll conclude that though there may remain fanciful counterexamples to a moral objectivity based on the pragmatic function of moral discourse, these don't pose a serious threat to the task of accommodation: our commonsense commitments to morality remain untouched.

2. Revisiting and Translating Expressivism

It's not difficult to read the expressivist arguments for objectivity we saw in Chapter Two in terms of Michael Williams' meta-theoretic framework, which we saw in Chapter Four. Recall that in this framework, the meaning-constituting inferential rules that govern our use of moral terms are given by the way those rules facilitate particular functions of ethical discourse, which is in turn explained by the pragmatic significance of that discourse in our lives. I have argued that an important practical aspect of moral discourse is its ability to press disputants into consensus on questions of immediate behavior—moral questions must be treated objectively when their answers bear directly on what we are to do. The challenge is to understand moral discourse in conversational contexts that don't have this practical feature, since some moral disagreements don't require an arrived at consensus for disputants to act in accordance with each other. Before we tackle this problem, though, let's revisit some of the lessons we learned from the expressivist project of accommodating objectivity, and transpose these into lessons about inferential norms.

Consider two friends, Connor and Duncan, who are engaged in an ethical debate. Connor argues that stealing is wrong; Duncan argues that it is not. It seems obvious that if ethical discourse is not relativistic, one of them must be wrong. Again, if moral claims functioned representationally—if they aimed to correspond to properties or states of affair in the world—the answer would be obvious: one party would be right if he succeeded in describing the moral properties or relations in the world as they actually are; the other party would be wrong in virtue of the fact that his claims failed to correspond to the moral

world. But because the lightweight realist remains silent on the topic of heavyweight moral facts, she can't hold out for these to explain why it would make sense for Connor and Duncan to treat the disagreement as one that has some objective resolution.

Recall the expressivist response to this charge. By rejecting relativism, the expressivist voices a general commitment to be resolute to particular commitments, to not be swayed by the mere fact that that others disagree or that she herself might have had different commitments. Connor doesn't need to make sense of Duncan's mistaken moral commitment in terms of some moral truthmaker that Duncan is failing to track—since the non-representationalist doesn't understand ethical discourse in terms of tracking truthmakers, she does not have to hold out for an algorithm to track them. Nor does she need to hold out for some principled perspective from which she can evaluate all of her moral commitments with absolute authority. For Connor, it should be enough to note that Duncan's commitments are shortsighted and inconsistent: allowing theft would undermine anyone's claim to property, including Duncan's. This kind of commitment to objectivity is a practical stance, a way of distinguishing between our moral commitments and mere whims. It is a commitment to stick to our guns even in situations when we see others do not share those commitments.

We can read this commitment in terms of inferential rules, and following the lessons of Michael Williams, make sense of them in relation to the pragmatic goal of moral discourse itself: our moral practice and thought cannot act as an effective chaperone of human behavior if it reflects commitments that are no more stable than whims. From this, we can understand the appropriateness of an inferential rule for a moral term like 'ought' like the following:

(3.1) *U*'s assertion "one ought to φ " implies that *U* has the attitudes in favor of φ -ing and against not φ -ing limned in (1.1) and (1.2) of Chapter Four, and that *U*'s commitment to these attitudes is not undermined sheerly in virtue of the fact that another agent has asserted that "one ought to not φ " or that "it's not the case that one ought to φ ."

In the last chapter, I focused specifically on the meaning-constituting inferential rules governing our use of the moral 'ought', in order to show how such an account can be used to solve the Frege-Geach problem. Because I'm addressing the objectivity of moral language in general this chapter, it will be helpful here to cast my net wider, and speak not only of specific inferential rules for particular moral terms, but for inferential norms that constrain our use of moral language in general. A more universal characterization of (3.1) gives us the inferential norm:

(Gun-sticking): In any moral disagreement, the mere fact that a disputant holds a different view does not compel one to infer that one's original position is undermined.

Recall from Chapter Two that for the expressivist, these considerations have a flip side: a spirit of resolve towards one's moral commitments doesn't guarantee that one is always right about any particular commitment. Expressivism doesn't give us license to be smug:

When I wonder how I might improve [my moral commitments], I have to think about it deploying my current attitudes—there is no standing aside and apart from my present sensibility. But that does not mean that I have to deem myself perfect, or incapable of improvement. It just means that whatever flaws I suspect are judged as flaws in the light of other concerns. (Blackburn 1993, 313)

The important point is that as a moral stance, relativism would render such reflective practice nonsensical. By endorsing the practice of reflection, and of stalwart opposition to the commitments of those like Duncan, the expressivist effectively rejects relativism.

Again, we can make sense of this sort of practical commitment in terms of an inferential norm that makes sense, given the role of moral discourse. If moral discourse is

to guide us into reasoned consensus about our behavior, it must include a norm that gives us room to change our minds:

(Reflection): For any moral conviction one might have, the mere fact of having it does not license one to infer that one's conviction cannot be improved³⁹.

It's worth pointing out the fundamentally pragmatic approach I'm taking here.

Those of a representational bent may feel like they have been tricked. They've come for a non-representational account of objectivity, but nothing has been said about the *objects*—the set of properties, facts, etc.—that are supposed to make moral claims true or false. No account has been proposed of how we can identify such objects; no procedure offered for resolving moral disputes to get to the truth of the matter. But this complaint misses the import of the pragmatic turn lightweights take: we aren't here to supply necessary and sufficient conditions for moral truths, nor are we in the business of supplying decision algorithms for determining who wins moral arguments. Remember: lightweights aren't hunting for an account of moral truth-makers. Instead of accounting for the nature of the objects of our ethical evaluations, the non-representationalist tries to account for the nature of the evaluations themselves—to show how various aspects of this nature function to encourage the pragmatic aim of moral discourse. The quest for objectivity doesn't end when we locate the objects of true moral judgments, the lightweight says, rather it ends when *we've made sense of being objective about our moral commitments*.

³⁹ How can we make sense of the notion of improvement here? Again, we can look to Simon Blackburn for an account:

How can I make sense of fears of my own fallibility? Well there are a number of things I admire: for instance, information, sensitivity, maturity, imagination, coherence. I know that other people show defects in these respects, and that these defects lead to bad opinions. But can I exempt myself from the same possibility? Of course not (that would be unpardonably smug). So I can think that perhaps some of my opinions are due to defects of information, sensitivity, maturity, imagination and coherence. (Blackburn 1998: 318)

Of course, there are two different ways one could make sense of being objective. The heavyweight realist would make sense of it by locating that set of properties or relationships that make moral claims true, and then show that these truthmakers are ones that do not vary from culture to culture, or from temperament to temperament. Again, such an account is unavailable to the non-representationalist—our understanding of moral claims can't start with an account of moral truthmakers. For the lightweight moral realist, we can only make sense of being objective about our moral commitments if we can show that this way of approaching moral questions—of treating moral disagreement seriously, in a way to explained below—is a necessary stance, if moral discourse is to play its pragmatic role in our lives. I'll argue that the kind of objectivity that the lightweight realist is left with is enough to do this—making sense of objectivity is a matter of explaining why it makes sense to pursue ethical disputes as ones that need resolution.

We lightweights must step very carefully, though. If the pragmatic turn justifies some of the inferences we traditionally associate with heavyweight conceptions of moral discourse, this does not guarantee that all aspects of our everyday moral practice will be vindicated wholesale. We must be wary of a greedy non-representationalism, which would gobble up all the resources of a heavyweight approach without checking first to see that they're all palatable. One particularly important effect—perhaps the most important effect—our commitment to objectivity has on moral practice lies in its role in moral disagreements--particularly the relationship between disagreement and error:

(D→E) For any genuine (moral) disagreement, it is necessarily the case that at least one disputant is in error.⁴⁰

⁴⁰ This is the inferential norm that makes sense of the general criterion for non-representational objectivity I argued for in Chapter Two:

This principle is a tough one to explain, though. It's important to notice that the inferences licensed by (Gun-sticking) and (Reflection) aren't sufficient to secure ($D \rightarrow E$).

Consider again the troubling ethical commitments of the Taliban that we contemplated in Chapter Two. When I engage with the sincere Talibanista in ethical argument, my commitment to objectivity does, as Blackburn argues, keep my resolve firm; the fact that he has different commitments than me doesn't undermine my own position. But why does it make sense for me to insist not only that my position on the appropriate treatment of women is right, but also that he is in error? If his assertions, like my own, genuinely reflect the inferences he's willing to make, and if he stands ready to justify these assertions (that is, if his assertion qualifies *as* an assertion, by Brandom's light of it registering both "upstream" and "downstream" commitments to his interlocutors), where could he be going wrong? We can of course point out other ways in which he might be wrong: his approval of the Taliban way could be informed by non-normative beliefs that are just plainly false—he might support his repression of women by citing the mental and physical inferiority of the fairer sex to men, or by pointing out that God Himself ordained that women are to be subservient. These beliefs are plainly wrong, and if they are essential to his moral commitments, we cannot see our disagreement as being purely evaluative. Or perhaps we could imagine a fully informed Talibanista who remains committed to oppressing women, but does so only because he hasn't yet resolved that commitment with his affinity for other principles—of fairness towards equals, for example. Doxastic failures and incoherency between commitments are one way of being

(Crit) A non-representational discourse X can establish objectivity if and only if the practical goal of X is one that could not be fulfilled without pressure towards coherence.

As we'll see, the justification for both ($D \rightarrow E$) and (Crit) in moral discourse are quite similar, though in this chapter I'll delve more deeply into pragmatic considerations to make sense of ($D \rightarrow E$).

in error, but it's not clear that either kind of failure is a necessary component of any genuine dispute. Is it impossible for a fully informed, perfectly consistent Talibanista to maintain his attitudes? If you think so: how about a Democrat or a Republican? How about someone who thinks that property rights trump obligations to charity?

The worry is that the only way inferentialists can explicate the error part of (D→E) is in terms of coherence, and there is more than one way of being coherent. The standard lightweight response to charges of relativism involves taking a pragmatic turn to explain what we're doing when we take an objective stance. I'll argue that explaining (D→E) will involve going deeper into our pragmatic turn and explain why we're doing this.

But perhaps the simplest thing to say is that the Talibanista believes women should be oppressed and that this is false. How can we make sense of such an error if we can't rely on those heavyweight truthmakers his claims fail to connect to? How, for that matter, can we make sense of the truth of my claim that women should be afforded equal respect? One answer to that question—the “minimalist package” often accepted by expressivists—offers a tempting, though flawed response to the above questions.

3. Objectivity from Minimalism?

We've seen that in order to justify the use of the truth predicate in moral claims, modern forms of expressivism marry their metaphysical approach with minimalism about truth. Minimalism about truth is itself a pragmatic stance; instead of investigating the nature of truth itself, we ask what function it serves in our language to call something 'true'. And all we have to understand the truth of claims is a deflationary or disquotational

sense: “Oppressing women is wrong” is true iff oppressing women is wrong. The crucial idea is that I can move from the sentence, “Oppressing women is wrong”, to the sentence, “It is true that oppressing women is wrong” without taking on any heavyweight metaphysical burden of explanation.

The pragmatic turn involved in truth minimalism is an important part of the project of accommodation: without it, we’d be left with Ayer’s conclusion that moral claims can’t be understood as true or false, which would fly in the face of our commonsense intuitions about moral discourse. And it seems that it’s at a least necessary condition (if not a sufficient one, as we’ll see) for making sense of moral objectivity. It would be very strange, if not incoherent, to insist that moral claims are never true, but that nonetheless moral discourse is objective.

This path is available for other non-representational approaches; Amie Thomasson takes it for the modal normativist:

Adopting a deflationary theory of truth also enables us to classify modal statements as true or false without their having to be made true by modal features of the world... Once we have modal truths, we can introduce terms for modal facts and properties. For we can derive singular terms referring not only to modal propositions, but also to modal properties and facts, by way of trivial transformations from uncontroversial truths... So the normativist position, as I see it, is one on which we get a straightforward first-order realism about modal properties and facts. But it is not a realism that takes modal facts and properties as *explanatorily* basic: treating modal truths as descriptions made true by modal facts or properties, and modal knowledge as derived from detecting these modal features. Instead, we work the other way up: we begin by understanding the function of modal discourse, use that to build an account of the meaning of ‘necessity’ (in terms of its inferential and epistemic role), use that (combined with a deflationary approach to truth and ontology) to account for modal truth, and explain our talk of modal facts and properties in terms of hypostatizations from necessary truths. (Thomasson 2013: 151-152)

As we saw in Chapter One, this route from explaining the function of moral discourse to securing truth for moral claims, to speaking of moral facts and properties, is available to

the moral non-representationalist—all without involving ourselves in any heavyweight metaphysical commitments.

Perhaps then we can continue to move along Ramsey's ladder, and get to a rung claiming, "It is objectively true that oppressing women is wrong." The path is relatively straightforward: lightweight realism gives us the tools to understand moral sentences at the bottom rung, simple sentences like, "Oppressing women is wrong"; minimalism about truth shows us how to attach the truth predicate to this sentence without paying any metaphysical price; finally, because (barring dialetheism) there can be no true contradictions, there can be no contradictory true moral claims⁴¹. The pressure to resolve moral disagreements is therefore made sense of, and so with it the final challenge of objectivity is met. This is Blackburn's recommendation:

How does minimalism enable anyone to dub themselves... a defender of objectivity and an opponent of relativism? The answer is, by taking all the terms that might be used to frame metaethical debates onto Ramsey's ladder:... taking us from 'p' to 'It is true that P' to 'It is really a fact... that it is true that p'. We saw that for the minimalist Ramsey's ladder is horizontal. From its top there is no different philosophical view than from the bottom, and the view in each case is just, p. In other words, if it is minimalism that justifies the ascent, then the ascent gets nowhere that is inaccessible to anyone of decent first-order ethical views. To say that an ethical view is true is just to reaffirm it, and so it is if we add the weighty words 'really', 'true', 'fact', and so on. To say that it is objectively true is to affirm that its truth does not vary with what we happen to think about it, and once more this is an internal, first-order ethical position." (Blackburn 2001: 295-296)

Sadly, this argument from minimalism won't do. You can't have your truth at a minimalist price, and then try to shoplift objectivity in with the package. As Blackburn says, a central point of minimalism about truth is the denial that adding a truth predicate to a claim adds any extra philosophical height. But adding objectivity does just that:

⁴¹ (Cf. Kölbel 2003: 64-66)

specifically it entails not only the content of the original claim *sans* truth predicate, but furthermore entails that any contradictory claim must be in error.

Perhaps an illustration is in order; I can put the story of the Commandos to work at this point. Recall that theirs is a society much like our own, except that they use the sentence-operator-like construction “it is commanded that” where we simply use imperative grammar, in order to issue commands. Let’s say that the Commandos are minimalists about truth, and so see no philosophical problems with adding the truth predicate to their odd imperative assertions.

Sharon the Commando wants Bob to come hither, and so says to Bob, “It is commanded that you come hither”. Reaffirming her commitment, she adds, “It is true that it is commanded that you come hither.” Meanwhile, Dan doesn’t want Bob to come hither, and so says to him, “It’s not true that it is commanded that you come hither”. Now, if the objectivity-via-minimalism argument above works for ethical claims, it should also work for Commando claims: Dan and Sharon have made contradictory claims; because there can be no contradictory true claims, at least one of them must be incorrect.

From this, we conclude that Commando discourse is objective. Of course, if we had an analysis of the operator, “it is commanded that”, such that we knew how it operated on a well-formed formula to stand for a particular property, this would not be so mysterious; we could justify claims to objectivity along the same lines as the heavyweight realist. The objectivity would come from the unified, consistent world of commanded properties to which these claims refer (or fail to refer). But by stipulation, Commando

claims don't function to represent properties⁴²; they function to give commands. The problem is that even with truth predicate in hand, the objectivity of Commando assertions makes no sense. In what sense have Sharon and Dan actually disagreed with one another? Why would such a disagreement need any sort of resolution, when one of them is trying to accost Bob and the other one is not⁴³? The moral of the story is the same as it was last chapter: in order to accommodate our normal ways of talking and thinking about moral discourse, we can't simply note superficial truths about the surface structure of moral talk. We have to dig deeper, and look to the goal that moral talk serves in our lives.

Now in the last chapter, we considered the worry that the force of this point stems from the artificiality of the Commando operator. We looked at reasons to suspect that Commando language would never get off the ground as assertive discourse, because their "it is commanded that" don't qualify as assertions. Perhaps we can rest assured that for any of our actual ways of using the truth predicate, this objectivity-via-minimalism path won't be so problematic, because any form of language that we use in conjunction with the truth predicate will have some feature that importantly differentiates it from Commando language. But we needn't concern ourselves with imagined languages to conclude this isn't the case. Securing objectivity via minimalism would allow us to do far too much, to claim a cost-free objectivity for any discourse—moral, modal, epistemic, aesthetic, etc.—simply by pairing a non-representational account with a disquotational approach to truth. This would impose objectivity on non-representational accounts of

⁴² We must avoid the temptation to think that "it is commanded that you ϕ " is true just in case someone has given the command to ϕ . This would be to conflate the appropriate conditions for uttering such a command with its truthmakers.

⁴³ Again, the parallel with Dreier's 'hiyo' example is obvious. Where the Commando accosts Bob by saying, "It is commanded that you come hither", Dreier's hiyo-user would do the same thing by saying, "Bob is hiyo." There is no difference between the two claims, except that one stipulated use acts like an operator, and the other acts like a predicate. The implications of truth minimalism are the same with both.

discourses, like aesthetics, for which many feel objectivity would be an awkward fit. And that would be greedy.

Consider again simple judgments of taste: one might plausibly hold that such claims serve the function of expressing pleasure or displeasure. “Chocolate ice cream is tasty”, then, would simply express the pleasure one takes in chocolate ice cream. Nonetheless, one can happily use the “tasty” predicate in sentences with minimal truth ascriptions, as in “It’s true that chocolate ice cream is tasty.” This on its own, though, won’t justify the assertion that such judgments of taste are objective; in fact we generally think such judgments are a hallmark of the subjective, whether we use them with the truth predicate or not.

Of course, truth may play a more robust role than mere disquotation. In the next section, I will look at Huw Price’s argument that “truth is normative, in a way not explained by the deflationary theory (Price 1998: 241)”. I don’t find his conclusions convincing, but he lays the grounds for a third pragmatic turn, wherein we look at disagreements themselves—not what disagreements are about, but instead what they’re for.

4. Disagreement from a Practical Perspective

In “Three Norms of Assertion”, Huw Price outlines three important rules that govern our use of assertions. These are the rules that we learn “when we learn how to acquire beliefs and how to reason with them” (Kölbel 2003: 68); they are the inferential norms that a speaker will have acquired in order to become linguistically competent.

These are the Brandomian norms that govern “which perceptual inputs justify which other beliefs, which beliefs justify which other beliefs, which course of action is motivated by which beliefs etc.” (ibid) Our practice of holding one another and ourselves accountable to these norms is crucial for establishing meaningful assertive communication within a group.

The first norm Price considers is the norm of subjective assertibility—it is “*prima facie* appropriate to assert p when and only when one believes that p.” (245) Such a norm is in place not only for assertions, but for meaningful utterances in general: “It is *prima facie* appropriate to request a cup of coffee when and only when one wants a cup of coffee”(ibid), even though such a request isn’t an assertion. There is a general expectation that speakers will meet this norm, though obviously it’s not one that can always be relied upon. The important point is that such a norm is a precondition for meaningful assertive communication—one may doubt that what he is hearing is sincere, but utterances get no traction at all if one doesn’t take it that she is at least meant to be hearing such utterances as sincere: “the hearer must have some basis for supposing [the speaker] must believe it to be true, and... the utterance is an attempt to inform the hearer and convince him of its truth.” (Searle 1965: 151)

The second norm is warranted assertibility—“‘p’ is objectively assertible by a speaker who not only believes that p, but is justified in doing so.” In assertion, this norm is important because it indicates to the listener that the speaker is not only sincerely expressing his belief, but that this belief is *warranted*—that the belief he expresses can give the listener a reason to also hold that belief. The norm of warranted assertibility is

also at play in non-assertive utterances, for example when I question your request for more coffee because I know you want to get to sleep early.

A linguistic community armed with only these two norms of assertion—a language that allowed what Price calls ‘Merely Opiniated Assertion’ (MOA)—would have a very limited notion of error available: one could be wrong in an assertion only if she was speaking insincerely (and so in violation of the first norm), or if her assertion was in some way unwarranted (violating the second norm). Price argues that neither of the two first norms can be the norm of truth, since “neither provides a norm that a speaker can fail to meet, even if she speaks sincerely, on the basis of a justified belief.” (245) But if the truth predicate were only disquotational, there would be no problem in using it in a linguistic community governed by MOA:

This speech community could quite well make use of deflationary truth, for example as a device to facilitate agreement with an expression of opinion made by another speaker. “That’s true” would function much like “Same again”, in the context in which a group of customers is placing orders in a restaurant... The crucial point is that if the only norms in play are subjective and objective assertibility, introducing disquotational truth leaves everything as it is. It doesn’t import a third norm. (Price 1998: 248)

This is not, however, how we in fact use the truth predicate; Price argues that this implies a third norm is at play:

There really is a third norm, we are inclined to think, even if these simple creatures [the Mo’ans] don’t know it. When they make incompatible assertions, at least one of them must be objectively incorrect—must have spoken falsely—even if by their lights they both meet the only norms they themselves recognize. (248)

This third norm is fundamentally the same as (D→E): both norms commit us to inferring that an error has been made when there is a genuine disagreement. And this is the norm that marks truth, according to Price. Now, I’m skeptical of this, because I don’t know that the truth predicate can be counted on to do such heavy lifting. Again, there are

topics that we are typically comfortable using the truth predicate with but for which $(D \rightarrow E)$ seems inappropriate; consider again the example from taste talk: how natural it could be to say, “That’s true” when another speaker claims, “Apples are tastier than pears”, if you share the opinion. The use of the truth predicate here, though, doesn’t seem to entail that anyone who prefers pears is wrong. One might argue that it’s technically wrong to mark your concurrence with the truth predicate in such a situation, though to my ear that sounds a bit pedantic. But my disagreement with Price here isn’t important; even if it’s not plausible to suppose $(D \rightarrow E)$ is a norm that governs all truth talk, it is a hallmark of *objective* talk.

What’s of interest in Price’s work on the subject is how he handles such a norm. Instead of asking, “What sort of property is it that we discover when we hit upon this third norm?”, he asks: “What might the invention of such a norm be useful for?” And his answer is pragmatic: the error that’s implied by disagreement makes for “unstable social situations” (251) which encourage argument and perhaps eventual agreement—“and it provides an immediate incentive for argument, in that it holds out to the successful arguer the reward consisting in her community’s positive evaluation of her dialectical position.” Disagreement of this type has obvious positive social implications: it encourages speakers to incorporate one another’s cognitive resources, to marshal evidence and reasoning in favor of their positions, which leads speakers to more informed, stable positions. By giving discourse a point of friction, adding a third norm “puts the cogs into cognition” (Price 2011: 165). It’s a good metaphor; the addition of cogs prevents two wheels from just spinning against each other—they introduce pressure that can get work done.

For lightweight realists, making sense of moral objectivity cannot be accomplished by explaining the robust ethical property that was discovered when our ancestors started speaking of moral truth. (Even if there were such a property, lightweight realists *qua* lightweights couldn't make use of it in their explanations.) What's needed instead is an explanation of why such a third norm should be at play in moral discourse. The question is not, "Does this norm pick out a real property in the world?" (to this question the lightweight realist should answer with silence), but instead, "Does it make sense for our moral discourse to be governed by this norm?"

The same lesson applies to non-representational approaches in other discourses, and may explain why (D→E) doesn't seem appropriate for some topics—questions of taste in particular (almost definitely regarding questions about tastiness as such, but also quite plausibly regarding aesthetic questions). For some kinds of discourse, the norms in place

involve the rule that different thinkers ought to agree, [but] sometimes they don't. That is, with some contents of belief ... there is a presumption that if one thinker correctly believes them, then no other speaker can correctly believe the negation of that content. For example, when I correctly believe that there are 25 chairs in this room then you can't correctly believe that there aren't 25 chairs in this room. Arguably, this is an *a priori* rule every competent thinker knows⁴⁴. With other contents, there is room for legitimate disagreement, and we learn that this is so when learning how to speak and think. For example, usually teachers will allow that learners believe that chocolate ice cream is delicious even if they themselves believe it isn't. (Kölbel 2003, 69)

The approach I offer makes sense of this phenomenon: the norms that govern our use of concepts like 'tastiness' only make sense in light of the functional role such concepts play in the lives of us hairless apes. As we've seen, the norms informing the meanings of terms make sense in light of the function of the discourse they serve.

⁴⁴ Kölbel is speaking of the *a priori* rules "governing which perceptual inputs justify which beliefs, which beliefs justify which other beliefs, which course of action is motivated by which beliefs, etc. (ibid, 60)" Again, these are just the inferential rules or norms with which the inferentialist is concerned.

Plausibly, we use ‘tastiness’ talk to warn others away from food they’ll find unpleasant, to implore them to try other foods, to indicate what the speaker herself likes, and also perhaps to promote common standards of taste in order to foster a sense of community. When a disagreement about tastiness arises in which both disputants satisfy Price’s first two norms of assertion, people are likely to think there’s no further fact that’s really being argued about. This makes sense, because once such a dialectical point has been reached, there’s nothing to be gained (from the perspective of the role ‘tastiness’ talk plays in a community) from a continued pressure towards resolution. The cogs of cognition should disengage so each disputant’s wheels can spin freely. If aesthetic disagreements function in roughly the same way we can shed light on the intuition that some deep disagreements about e.g. whether the Beatles are better than the Rolling Stones may sometimes have no resolution⁴⁵.

5. Inferentialism and Objectivity

What hope is there for robust moral objectivity, then, from an inferentialist viewpoint? To answer this, we must turn to the pragmatic aim of moral discourse. For a straightforwardly representational discourse (i.e., the type of discourse that marks sentences like, “There are 25 chairs in this room”), the pressure towards concurrence that is the hallmark of a robust objectivity comes from our understanding that the aim of that

⁴⁵ There may be times when context does drive such debates further than they’d normally go: “I’m more inclined to regard myself as disagreeing with someone whose favorite dessert is Rice Krispie Treats if his plan is to order it for the whole table than if he consumes it among consenting adults in the privacy of his own home.” (Field 2009: 277-278) It’s helpful to notice that the stakes are raised when the question at hand is about what we should do next. The import of this for my argument is clear: “If the person acts on moral preferences different from mine, this will usually affect things I care a lot about, which I take to be a large part of the reason that moral differences are usually disagreements whereas differences in taste often aren’t.” (ibid.)

discourse—accurate representation of a unified, coherent world—cannot be met if we have no pressure to resolve the tension between contradictory claims.

We've seen in Chapter Two that the non-representationalist must explain this pressure in pragmatic terms. It seems we must treat our ethical commitments as objective because doing so is the only way we can answer the moral question, "What do we do now?" We can see then, why the need might arise to treat our ethical commitments as objective. Accepting (D→E) in moral discourse ensures that the cogs will always engage, and so always drive us to answer that question.

But we also saw that this leaves the ethical expressivist with a debt to be paid: she must explain why our ethical disagreements still count as disagreements even in contexts where their resolution doesn't have any obvious behavioral payoff. For the expressivist, this explanation must work through an account of the attitudes moral claims function to express, such that those attitudes can be thought of as being in disagreement with one another. This is a difficult task—as we saw in Chapter 3, the kinds of accounts expressivists seem to offer face challenges in explaining the relatively straightforward disagreements between assertions of obligation and permission in terms of clashing attitudes, much less the attenuated cases of disagreements about who is to take the Ibuprofen, or how America should've behaved in World War II, or what Superman should do with his time.

Once again, the simplicity of inferentialism counts as an advantage: unlike expressivists, inferentialists don't need to advert to mental states to explain moral disagreements. They instead need to make sense of the inferential norms at play by identifying the pragmatic aspects of moral language that necessitate these norms. What's

needed is a more nuanced investigation into the pragmatic significance of moral discourse—and how it's served by our use of moral claims.

Allan Gibbard's account of the role of moral discourse (1990) is instructive. He sees ethical commitments as a subcategory of normative commitments in general—according to Gibbard, ethical questions are about the rationality of the moral sentiments of anger, guilt, shame, and resentment—and he argues that normative claims are best understood expressively. The naturalistic explanation of our normative commitments is found in the type of game-theoretic considerations found in the evolution of social behaviors. Coordination is crucial for social creatures such as ourselves, and the language-infused coordinating mechanisms we would expect to arise in such creatures is pretty much what we find in ourselves. Social behavior is coordinated by way of coordinating beliefs and the moral sentiments (and, the inferentialist might add, by way of coordinating the inferential roles that inform our use of moral terms).

Because we're intelligent language-using animals, and because preparation and rehearsal are often indispensable, we don't only bargain with one another about the appropriateness of specific acts or feelings, but also about the appropriateness of a broad range of possible acts or feelings. Speaking counterfactually gives us the ability to coordinate what we would do, and coming to such accord gives a group (and its members) a distinct selective advantage.

One way language influences actions and emotions is by letting people think together on absent situations. With language, people share not only their immediate situation, but past, future, and hypothetical situations as well... That opens a great new scope for coordination, and so a capacity for shared evaluation would be fitness-enhancing in a species with a complex social life. Those who can work out together reactions to an absent situation—what to do and what to feel—are ready for like situations. They are better prepared than

they would otherwise be to do what is advantageous in a new situation, and they can rely on complex schemes of interpersonal coordination. (Gibbard 1990: 72)

From this practice arises a commitment to systematic coherence, even in the context of disagreements over what should be done in hypothetical situations.

There is value in pursuing moral disputes objectively even when it's unlikely that one is rehearsing for some possible situation, and even when such disputes won't obviously lead to any immediate practical conflict, as with Dreier's Ibuprofen case. Disagreements in such cases may not be about what we should do immediately, but might instead be about the *principles* that inform what we would do.

An inappropriate commitment may be not only a behavioral serious liability in itself, but also a manifestation of a potentially more serious liability—a sign of a flaw in the system of background beliefs, rules or habits of inference, from which the commitment in question arose. It seems that disputes have the function not merely of correcting individual mistakes, but also of encouraging improvements in the general beliefs and dispositions that leads to those mistakes. (Price 1988: 158)

Dreier's dispute regarding Ibuprofen may not lead directly to any practical conflict, but it seems quite likely it could reflect disparate commitments, normative principles that could lead to future conflict in the long run.

We shouldn't fall into the trap of expecting some behavioral payoff to every dispute, immediate or otherwise. According to my incorporation of Williams' framework, the semantic properties of moral terms are to be explained by the functional role those terms play in the context of moral discourse's influence on social behavior. But in order to have such a function generally, it's not necessary that moral terms fulfill this function in every particular instance. For any non-representational account,

it is important that dissent be cued to something of behavioral significance. However, these behavioral consequences are long run affairs, and no doubt hideously complicated. It would be absurd to require that in individual disputes, speakers are cued by an understanding of the behavioral significance

of the opposing points of view... We need to separate the practice of dissent from the underlying facts that explain its development. (Price 1988: 147)

As I wrote in Chapter 2, on occasions of moral disagreement, it will be unhelpful for both parties to simply agree that each is right in his or her own relative way; such unchecked disagreement will leave unanswered the crucial question, “What do we do now?” We can now see why it also makes sense to also try to objectively resolve disputes about what we should have done, about what Superman ought to do, and about how, in principle, we should deal with the Ibuprofen in front of us.

Seen in this light, our commitment to an objective morality can be understood in terms of the inferential roles that—over the long run—serve the function of promoting social coordination in the ways outlined above. As an inferential rule governing a single moral term like “ought”, we get:

(3.2) *U*'s assertion “one ought to ϕ ” defeasibly implies that any assertion to the contrary—“one ought to not ϕ ”, or “it’s not the case that one ought to ϕ ”—is mistaken.

And as a general inferential norms on moral claims:

(Gen) For any pair of moral assertions *x* and *y*, if *x* and *y* are logically inconsistent, at least one of these assertions is mistaken.

And (Gen) is just a different way of expressing (D→E), which is the inferential norm I set out to justify.

Again, such norms comport with the meta-theoretic framework Williams provides:

The meaning of a claim is constituted by its inferential roles, which are in turn explained by the function those roles play in the discourse in question, which make sense in light of the pragmatic significance the discourse has. A discourse not constrained by inferential rules and norms like (3.2) or (D→E) will fail to establish meaningful disagreements in situations where the payoff of such disagreements doesn’t cash out in terms of a practical

clash, and as such will not exert pressure towards consensus. This means such a discourse would be ineffective in corralling the behavioral influences implicit in the inferential roles (1.1)-(2.3) suggested in Chapter 4.

Recall that these inferential rules are meaning constituting; as I argued in Chapter 4, the inferences licensed by moral concepts establish the meaning of these concepts in a way that remains constant even in force-stripped contexts. Similarly, the inferential roles (3.2) and (D→E) should guarantee that moral terms retain their objective status regardless of the practical context in which they're being used. In what follows, I will consider the implications of such an account of objectivity: does this approach promise us the same kind of objectivity (and the same notion of error) that could be secured with a representational approach?

This pragmatic approach gives us insight into other non-representational accounts. Consider for example epistemic expressivism. Matthew Chrisman (2007) argues that the function of epistemic discourse is to regulate and keep track of our "information economy"; claims regarding knowledge indicate, "who can be trusted about which kinds of information." (243) With this account of epistemic discourse at our disposal, one can see a feasible pathway to objectivity. If Chrisman's suggestion is right, then an epistemic disagreement may share the crucial objectivity-conferring property I argue we find in moral disagreements: inconsistency would frustrate the function of the discourse. Having inconsistent standards for knowledge attribution, on Chrisman's "information economy" model, would not be unlike attempting to make trades on the stock exchange without an agreement about the value of currency. Disengaging the cogs of disagreement would compromise our ability to keep track of whom to believe; it makes sense to demand

compliance with all three of Price's norms. As we saw in Chapter Two, this is the lesson generalized:

(Crit) A non-descriptive discourse X can establish objectivity if and only if the practical goal of X is one that could not be fulfilled without pressure towards coherence.

With a little imagination, we can see why some discourses might be good candidates for objectivity, but others might not. If the function of the discourse can only be met if $(D \rightarrow E)$ is a norm that guides that discourse, then we have good reason to think that such a discourse is objective. If, on the other hand, $(D \rightarrow E)$ is not a good norm—as I argued it isn't for talk about tastiness—we should expect for there to be points in disagreements at which the cogs disengage.

Critics may object. Such an account of objectivity “smells of sulphur” (BB Ruling Passions vi); it's left out the real heart of the matter. Some may object that I've simply left out any real sort of objectivity in my argument for objectivity. I want to take this possibility seriously—what does this pragmatic account leave out? The challenge for the critic is to put her finger on what precisely is so bothersome about the kind of objectivity presented here. It's worth our time to see what such a critic can come up with.

Talking about norms of discourse and their place in the moral domain is all well and good, the objection goes, but it leaves us with no there there. There is no account of what such an objectivity itself amounts to; I've failed to indicate any sort of foundations for objective moral truths. Most damningly: I have at best told a story of why it makes sense for us to behave as if there were objective moral truths, and this is decidedly not the same as explaining why there actually are such truths.

As I've argued, this line of objection is in a sense question begging. It amounts to insisting that the only way to get objectivity is with a heavyweight account—one that

offers some analysis of the truth-makers for moral claims. And this is precisely the method that is being rejected by the expressivist approach. If the objectivity depends on heavyweight realism, then of course no non-representational account will suffice. But that's exactly the question that's being considered, and it's unfair to insist it can't be answered from the outset.

Perhaps it will help to distinguish between two types of objectivity: the heavyweight notion of objectivity and the lightweight objectivity I've argued for here. The question is: What do we gain if we have a lightweight objectivity, and then move on to secure objectivity writ large? We already can explain, given minimal objectivity, why it makes sense to stick to one's guns in the face of moral challenge, why one should question one's moral commitments, why we can speak of moral truths with a straight face, and why it's reasonable to pursue disagreements even when a dialectical adversary isn't obviously being insincere or inconsistent in her commitments.

Does heavyweight objectivity give us anything more? It strikes me that I'm not offering an 'as-if' moral objectivity, but moral objectivity in the only sense that matters. Consider my friend Ryan, who is utterly convinced of the correctness of an expressive treatment of meta-ethical issues and has found this conviction hasn't changed his commitments regarding normative topics—he still reasons about moral issues, gets into heated arguments, and feels the force of his moral conclusions in his practical decisions. The challenge for the heavyweight realist at this point is to reveal some inconsistency on Ryan's part; the worry for the heavyweight realist is that they can't, and so they've incurred a debt (ontological and epistemic) for a product of little value.

6. Deep Impasse

Perhaps this invitation can be met, though: Imagine two disputants engaged in a disagreement. Furthermore imagine that both of these disputants are beyond reproach in their compliance with Prices' first two norms of assertion. Both are being sincere, and each has consistently worked out all the consequences of his or her commitments. And throw this in for good measure: they're both maximally informed about straightforward matters of fact that are relevant to those commitments. Call such a situation a "deep impasse". Now both the lightweight and heavyweight objectivist will have to explain why it is that at least one of these people must be in error. And I think the heavyweight explanation is at least initially more appealing: one of the two disputants, no matter how commendably consistent he may be, is guilty of one error in particular: he is wrong. How are we to understand this wrongness? Again, very simply: there is some mind-independent fact of the matter out in the world, waiting for a factually correct claim to correspond to it. And inasmuch as the disputants are engaged in a genuine disagreement, and facts follow the law of non-contradiction, then someone has gone awry. Compare: I believe that Tyche, the giant gas planet some astronomers believe orbits in the distant Oort cloud, really exists, and you disagree. We are both ideally informed and consistent in our claims, so in virtue of what must one of us be wrong? Obviously, in virtue of the truth-maker (or lack thereof) for our respective claims, the fact that the Tyche either exists—and so you are in error—or it doesn't—and so I am.

Like the heavyweight realist, the lightweight realist will also insist that one of the two interlocutors must be wrong, but what explanation of wrongness is available here? A very minimal one:

The prior question is whether quasi-realism offers an ‘account’ of moral error at all. What I do offer is an account of a state of mind and a process— the state of mind of worrying whether one is oneself in moral error, and of the process of seeking to root out any hidden error. But if some theorist bent on finding truth-conditions asks me what my account of moral error itself is, then I am not very forthcoming... ‘What it is’ for it to be a mistake to think that happiness is always good, is happiness being sometimes indeliberate or bad. What it would be for it to be a mistake to think that slavery is always wrong, would be slavery being sometimes just fine.... To recapitulate, error is as deflationary as truth. (Blackburn 2009: 207)

This is, at least at first pass, distinctly less satisfying than the truth-maker account of objectivity the heavyweight realist makes use of. The heavyweight realist may be frustrated with my refusal to treat questions of objectivity in metaphysical terms, but here we see a way to frame the challenge with a pragmatic twist: perhaps what’s wrong with Ryan is that he has faulty commitments.

Recall my considerations about objectivity in matters of taste above. I argued that (D→E) probably wouldn’t be the best norm for all disagreements about tastiness, because in situations where both disputants feel confident that they’ve both obeyed Price’s first two norms of assertion, continued pressure towards resolution would be counterproductive—it wouldn’t serve the broader function of discourse about taste. The worry then is that the same lesson applies to moral discourse, and in letting (D→E) govern his moral commitments, Ryan is being silly; he’s committed himself to a norm that just doesn’t make sense. In the face of a deep impasse, his commitments lead him to behave like a child who insists you’re wrong to say that peas are delicious, even when you’ve explained why you’re saying it. The child’s obstinacy reflects a lack of awareness that there are important norms about taste talk that “tie a belief to certain features of its

possessor.” (Kölbel, 69) Perhaps so too with moral objectivity: such a norm would have us maintaining disagreements even in situations when the cogs of disagreement have ground to a halt. The disagreement remains, but it does no good. Questions of heavyweight properties aside, it would be problematic to commit to moral objectivity if such a commitment were misleading.

I don't think that this pragmatic objection should be taken to undermine all non-representational claims to objectivity. Plausibly, some forms of discourse play a pragmatic role such that a set of norms including $(D \rightarrow E)$ governing those discourses is not misleading. Take for example a non-representational approach to modality. Remember that in Amie Thomasson's modal normativism, modal sentences like “It is possible for X to be Y” or “For any X, necessarily X is Z” make explicit the socially regulated linguistic rules of use that set the conditions under which our concepts can be appropriately applied: “Necessarily, all bachelors are men” makes explicit the metalinguistic rule: “Apply ‘bachelor’ only where ‘man’ applies.” (Thomasson 2007a) Modal claims reflect constitutive semantic rules for concepts, or the logical consequences of these. The broader purpose of modal discourse is the orchestration of communication; it is a way that we coordinate common meanings for our terms. In the same way that the moral expressivist points to the function of coordinating social behavior, and the epistemic expressivist points to the function of regulating a cognitive economy, the normativist can point to this function of modal discourse to give a pragmatic take on the appropriateness of a modal version of $(D \rightarrow E)$. This shows how modal discourse fulfills the master criterion that I argued in Chapter 2 was a precondition for any non-representational justification of objectivity:

(Crit) A non-representational discourse X can establish objectivity if and only if the practical goal of X is one that could not be fulfilled without pressure towards coherence.

A precondition for meaningful communication is that people are using their terms in the same way (rather than using different terms homophonically)—that they're using terms with the same application and co-application conditions (Thomasson 2007b).

Given the purpose and nature of modal discourse, the problem of deep impasse shouldn't be so troubling. This becomes clear when we think about what reflecting on what's involved in complying with Price's 2nd norm of assertion: a fluency with the rules of use any speaker must master to competently use a term (e.g., speakers who don't know to apply 'bachelor' only where 'man' applies aren't competently using the former), an implicit understanding of fundamental modal inferences (e.g., $\Box p \leftrightarrow \sim \Diamond \sim p$), and knowledge of any of the empirical facts that might be relevant to our use of a concept⁴⁶. If two disputants have complied with these norms, but still have a disagreement, Thomasson argues that we can see this dispute as reflecting a pragmatic disagreement about how we should use the relevant terms:

such debates are clearly pragmatic debates about how we ought to revise the rules of use for our terms, not factual debates that can legitimately purport to yield discoveries about what the relevant detailed conditions [of use] really are. (Thomasson 2009: 452)

In such cases, the pragmatic dispute may yet yield results, inasmuch as there are compelling reasons to precisify a term or otherwise change the norms for its use, but if even this dispute stalls, we can opt for an easy pluralism—since terms are individuated by the rules of use they have, neither disputant is really disagreeing because they're both using different terms and can carry on as long as they understand they're doing so. Notice

⁴⁶ e.g., It's a conceptual truth that whatever water's chemical composition actually turns out to be, it has that composition necessarily. When we discover the empirical fact that water is actually H₂O, it's trivial to infer that water is necessarily H₂O (Cf. Thomasson 2007b: 62)

that this is still a model of non-representationalism: we shouldn't conflate the norms that guide modal claims with the truthmakers for those claims. "Necessarily, humans are apes" reflects our linguistic commitment to use "ape" to refer to any human, but that commitment itself doesn't count as a truthmaker for the sentence; humans would necessarily be apes whether we had that linguistic commitment or not.

7. Deep Moral Impasse

Unfortunately, the model of objectivity available in modality isn't as plausible for moral discourse. Deep impasse isn't so worrisome for the modal normativist, because a consistent application of the norms governing modal discourse will either lead to the resolution of disputes, or conclude with an easy pluralism. This means the modal normativist can stand by her commitment to $(D \rightarrow E)$. But deep impasses can't be dismissed so easily for the moral inferentialist. Simon Blackburn points out that "there are not so many livable, unfragmented, developed, consistent, and coherent systems of attitudes" (1984: 197), but if there are even two such systems that entail different moral conclusions, the disputes they engender should be worrisome: the possibility looms that we've gotten on board with a inferential norm, $(D \rightarrow E)$, that at least sometimes misleads us. It's worth our time here to seriously investigate the difference between the approach moral non-representationalism and its modal cousin take to deep impasses.

If there is such a difference, we're left again with the troubling comparison to the committed objectivist about matters of taste, who is adamant about pursuing disagreement even when the only possible bones of contention have been picked over—who, to return

to Price's metaphor, insists on engaging the cogs of argumentative discourse even when there's no hope of them getting unjammed. Such a taste objectivist is acting irrationally, not only because there is no heavyweight objective property of 'tastiness' to adjudicate disputes, but because his obstinacy can't be justified even pragmatically—whatever the function of discourse about tastiness is, it's not being served by cleaving to (D→E). The threat here is that on a non-representational account, moral objectivists are being equally bullish in pursuing disagreements past the point of deep impasse, and the only satisfying justification they could give for doing so—a pragmatic one—disappears in situations where the cogs of disagreement are locked, and hold no promise of moving. And unlike the modal non-representationalist, the moral expressivist doesn't have the 'out' of claiming an easy pluralism; since moral debates concern how we are to behave (as opposed to how we are to use a term), we cannot simply agree that we are to behave in two mutually exclusive ways.

This is an objection worth taking seriously, but thoughtful reflection on the nature of moral disagreement will—once again—take much (if not all) of the sting out of it. The notion of a deep impasse is helpful, but we should ask how much of a role it actually plays in our moral engagements. It may not be that difficult to conclude that you and interlocutor have both fulfilled Price's first two norms of assertion in discussions regarding the deliciousness of a dish (though it may take longer than one would initially expect—watch the judging portion of any competitive cooking program to see the myriad and subtle ways in which food critics invoke shared norms to argue for particular verdicts), but how often does this happen in moral discourse? When we consider all the elusive influences that shape our moral commitments and the clarity of mind and precision

of tongue (not to mention the free time) it would take to elucidate these, we might wonder if it has ever happened—even in the rarified air of philosophical discussion⁴⁷.

This has practical import. Remember that the question at hand is whether or not $(D \rightarrow E)$ is an unreliable conversational norm for moral disagreements. If a recognition that disputants are in deep impasse is the only threat to $(D \rightarrow E)$, we non-representationalists may take comfort in their infrequency, all the more so if we think this never happens for moral arguments. As long as there's no deep impasse, the cogs can get traction and do work that will on the whole serve the coordinating function of moral discourse.

The careful reader may object that I've conflated two issues: the problem of maintaining a commitment to $(D \rightarrow E)$ in the face of a deep impasse with the problem of maintaining a commitment to $(D \rightarrow E)$ when we *recognize* that a deep impasse has been reached—the former may happen much more often than the latter. I'm not sure this speaks to what ways of moving forward are reasonable, though. One may be in a position of deep impasse without knowing it, and be entirely justified in pursuing the disagreement; it makes sense, given what we've said about the role of moral discourse, to at most be agnostic about the possibility that any particular moral dispute is an instance of deep impasse. As long as one thinks it's plausible the disagreement is not in this sort of deadlock, one may continue fighting the good fight without being vulnerable to charges of myopic bullishness; agnosticism about deep impasses makes sense because one can never

⁴⁷ Timmons treats the political disagreement about property rights and welfare between Hilary Putnam and Robert Nozick (1999: 171) as a case of what I'm calling a deep impasse, and it's as good a candidate for the title as I've seen, though I'm still convinced that given the right argument, Nozick could have seen the error of his ways. Even if I'm wrong about this, though, it's plausible that many deep ethical conflicts—like that between Western liberalism and the Taliban, or the thorny abortion debate in America—would go a long way (maybe all the way) toward concurrence if only the relevant disputants were divested of some of their false religious assumptions.

be sure that a particular disagreement won't eventually (or even suddenly) reach resolution.

Further reflection on the nature of moral disputes offers a different defense of (D→E): Even if we do find a case of deep impasse, these do not happen in a vacuum. When we offer arguments and justifications, we often do it in the context of an audience. So even when our interlocutor cannot possibly be moved by reasoning, it may be useful to offer it nonetheless—because others can. (Those who debate trolls on the Internet often use this line of reasoning to justify the way they spend their time.) Even without a sympathetic audience, maintaining a disagreement in the face of a deep impasse can be like a Supreme Court justice recording a dissenting opinion: it may have no impact on settling the law at the time, but the record may one day be important.

Finally, as we've seen Gibbard point out, "we rehearse arguments as a part of refining our normative convictions" (1990: 199), so that even in contexts where my arguments are convincing no one, maintaining a pressure towards resolution can help me think through my own commitments, which can in turn be valuable for future moral engagements. So even if we admit the possibility of a deep moral impasse, the pragmatic treatment of disagreement is still relevant: when we disagree we position ourselves to pool our cognitive resources—to share knowledge and reasoning so that we can be epistemically better placed at the end of a disagreement than we were at the beginning. The problem of deep impasses doesn't undermine this practice so long as one understands the final "we" above in a broad sense, including in its scope not only the disputants themselves, but also the audience of the dispute, those potentially concerned with the debate in the future, and of course one's future self.

Perhaps an imaginative philosopher might be able to offer a counterexample that undermines all of these pragmatic considerations. Consider a test case: you and I are the last two people on Earth. I think we have an obligation to spend some of our remaining days building a monument to our fallen race; you think that would be a sentimental waste of time, and we should instead just spend the rest of our lives enjoying ourselves as much as we can. Furthermore, we're both enormously articulate and thoughtful (capacities which no doubt helped us survive the apocalypse), and with all the free time we have on our hands have reached a point in our dispute where each is convinced the other is unmovable in her convictions—we're each of us perfectly coherent, given our particular moral commitments, and we both know it. We are divorced from the normal context in which it makes sense to consider this disagreement one with a correct resolution: you know that none of your arguments will sway me; you know that there is no audience for our debate; you know that no amount of moral rehearsal will be of use in the future. Now, I set out in this chapter to defend objectivity without heavyweight realism, and have done so by justifying our commitment to $(D \rightarrow E)$ —by showing that it always makes sense to assume error in the situations where there's real moral disagreement. Do we not have here a counterexample to that claim?

Of course, it's not clear that such a situation is very plausible. Again, we might doubt that two humans could be possessed of such clarity, memory, imagination, free time, etc., to be sure that they've reached the deep impasse I've described above. I'll concede the point, though, because what we're left with is interesting; it gives a pragmatic voice to the heavyweight realist's worries about non-representationalism's objectivist inadequacies. Like Blackburn's quasi-realist, I set out to defend our everyday use of moral

discourse, to show that even without relying on truthmakers, it makes sense to keep $(D \rightarrow E)$ among the norms that guide moral conversation and thought. Even if we take such counterexamples seriously, what we're left with is still an objectivity worth wanting, one that makes sense of and justifies our actual practice of treating moral disagreements seriously. If the only objection to my defense of expressive objectivity stems from rarefied philosophical thought experiments like the one above, I'm comfortable in concluding that our ordinary practice is safe.

Remember that the worry that motivated this chapter was the possibility that moral discourse was in need of serious revision. We hoped to avoid the conclusion that our moral language resembled that of the Commandos—one in which a serious consideration of the inferential norms that govern a discourse and the aims of that discourse are utterly out of step with one another. Were we Commandos, we would have to reform our discourse fundamentally. To make the inferential norms that we follow comport with the aim of imperative language, we'd need to cut many of them out of practice: we could no longer embed "it is commanded that" in the antecedent of conditionals; we would also banish use of the truth predicate in conjunction with that phrase. We'd need to rethink the entire structure of the discourse.

If we take the apocalyptic counterexample above seriously, what sort of revision does it recommend to the inferential norms informing moral discourse? Perhaps we should replace $(D \rightarrow E)$ with $(D \rightarrow AAE)$:

$(D \rightarrow AAE)$ For any genuine (moral) disagreement, it is almost always the case that at least one disputant is in error. Disputants aren't necessarily in error in the context of a deep impasse in which no audience is privy to the dispute, no posterity stands to benefit—where there is no *possible* practical payoff to a protracted disagreement.

But in the context of normal (and even most of the very unusual) moral disputes people enter in, $(D \rightarrow AAE)$ would be functionally the same as $(D \rightarrow E)$; the two norms would play practically indistinguishable roles in the discourse. This *is* the distinction the heavyweight realist wanted between his account of objectivity and the lightweight realist's account. From a heavyweight perspective, $(D \rightarrow E)$ would be an appropriate norm to follow even in the fantastical apocalyptic situation we're considering. One of the disputants would have to be in error, because the heavyweight moral fact of the matter would be out there, like Tyche in the Oort cloud, waiting for a correct resolution to the dispute. But it is a distinction without a difference; if a revision is called for, it's one so picayune that none would notice its implementation.

8. Conclusion

In the first chapter of this dissertation, I argued that non-representational approaches to morality have epistemic and ontological advantages over their representational contenders: they sidestep the problem of placing moral truthmakers in the world in a way that comports with the restrictions of naturalism, and so too the problem of explaining how we come to know about these truthmakers. But non-representationalists also have a special obligation to discharge: accommodating the commonplace assumptions that underwrite ethical discourse—especially our intuitions about moral objectivity and reasoning. And because they cannot make use of heavyweight moral properties as explanatory, they must accomplish this task by identifying what it is we're doing when we engage in moral talk.

In Part One, I showed how expressivists have approached this task: they argue that what we do when we make moral claims is express evaluative attitudes. This pragmatic account of moral discourse can go a long way towards securing objectivity and meaning for our moral claims, I argued, but the expressivist reliance on attitudes as the central driving force behind moral discourse unnecessarily limits how much a lightweight realist can accommodate. In Chapter Two I argued that an expressive account of objectivity requires an account of moral discourse that shows why its characteristic function is one that cannot be fulfilled without a pressure towards coherence. But we saw that for the expressivist, we can only understand this pressure in terms of the mental states that contradictory moral claims express. We must come up with an account of these attitudes such that it's clear all moral disagreements express attitudes that undergird the right sort of pragmatic clash. And we left the expressivist with this challenge unanswered.

In Chapter Three I proposed a pragmatic solution to the Frege-Geach problem for expressivist. They can account for the meaning of moral claims in a way that addresses the problem of embedding by locating the moral sensibilities that such claims function to express. We saw that the meaning of a claim is given as an expression of the same sensibility whether it's embedded or not—so the expressivist is not vulnerable to charges of equivocation. Again, though, the expressivist's commitment to explaining the meaning of moral claims in terms of expressed attitudes complicates matters: in order to give a complete account of the inferential relationships that mental states have to one another, the expressivist must accept novel commitments about non-moral language.

In Part Two, I offered a different non-representational approach, moral inferentialism, which shares the lightweight metaphysical commitments of expressivism,

but offers a simpler account of the meaning of moral claims—one that adverts to inferential relationships between moral claims themselves, instead of the sensibilities expressed by these claims. I argued that an inferential approach gives us the tools to answer the semantic worries of the Frege-Geach problem in Chapter Three. An appropriate understanding of the use of moral discourse—the facilitation of coordinated social behavior—explains the kind of inferences that are licensed by moral concepts. But while the function explains why we would want a term with this inferential role, the meaning is given in terms of the inferential role—not the use. I offered a sketch of the inferential roles the moral term ‘ought’ plays, and argued that once we recognize that the relevant inferential roles are meaning-constitutive, we will be in a position to solve the Frege-Geach problem in a new way, for those inferential roles remain constant even in contexts in which the term is used differently. I argued that the inferentialist solution to the Frege-Geach problem gives simple answers to at least two semantic worries that have vexed contemporary expressivists—the “problem of permissions” and the commitment to “mentalism”, both of which are problems that don’t get traction with an inferentialist approach.

In this chapter, I’ve tried to show that the same approach can make sense of our commitment to moral objectivity. This is the sort of objectivity you can get without heavyweight realism: it comes (again) from reflecting on the pragmatic role of the discourse. Here’s how you can get moral objectivity by taking a series of pragmatic turns: First, recognize that the commitment to objectivity plays a valuable role in the discourse at hand—it renders speakers stalwart in the face of disagreement, and furthermore makes sense of inspecting one’s own convictions for fault. Both of these practices make sense in

the context of the pragmatic role that discourse plays in our lives. Second, adopt the minimalist package, so you don't have to take the awkward position that there are indeed no moral facts or truths, but that nonetheless morality is objective. Then explain why it should be the case that disagreement within the discourse should entail error—why this is an inferential norm that makes sense, given the practical role moral discourse plays in human life.

I've also proposed a way of accomplishing a lightweight objectivity, not only for moral discourse, but also for the modal, epistemic, etc. The path I introduce comes from serious reflection on the role that disagreement itself plays within a discourse: the moral inferentialist justifies her claims to objectivity by showing the kinds of inferential norms that moral discourse needs to do its job. This method also clarifies why a non-representational approach to other discourses (like aesthetics) might be implausible candidates for objectivity: the practical role that some subject matters play in our lives is such that there's no point in insisting that all instances of disagreement entail error. But this isn't the case with moral discourse: when deciding on what should be done, it's worth treating disagreements seriously.

References

- Alwood, Andrew (2010). Imperative Clauses and the Frege-Geach Problem. *Analysis* 70: 105-17.
- Ayer, A. J.(1952). *Language, Truth and Logic* (2nd ed.). New York: Dover Publications.
- Beall, J. C. (2009). *Spandrels of Truth*, Oxford: Oxford University Press.
- Blackburn, S. (1984). *Spreading the Word: Groundings in the Philosophy of Language*. New York: Oxford University Press, USA.
- _____ (1988). Attitudes and Contents. *Ethics*, 98 (3): 501-517.
- _____ (1993). *Essays in Quasi-Realism*. New York: Oxford University Press, USA.
- _____ (1999). Is Objective Moral Justification Possible on a Quasi-realist Foundation? *Inquiry*, 41: 213-28.
- _____ (2001). *Ruling Passions: A Theory of Practical Reasoning* (New Ed ed.). New York: Oxford University Press, USA.
- _____ (2006). The Semantics of Non-Factualism. In Devitt, Michael & Richard Hanley (ed.), *The Blackwell Guide to the Philosophy of Language*, Malden, MA; Oxford: Blackwell.
- _____ (2009). Truth and A Priori Possibility: Egan's Charge Against Quasi-Realism. *Australasian Journal of Philosophy* (87) 2: 201-213
- Bloomfield, P. (2003). Is There Moral High Ground?. *The Southern Journal of Philosophy*, 41: 511–526.
- Brandom, R. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.
- _____ (1998), Asserting. *Nous* (17) 4: 637-650
- _____ (2007, June 3). Lecture Four: Modality and Normativity: From Hume and Quine to Kant and Sellars. *John Locke Lectures*. Lecture conducted from Oxford University, Oxford.
- _____ (2010) Inferentialism and Some of its Challenges. In B. Weiss and J. Wanderer, Eds. *Reading Brandom*. New York: Routledge, p.159-180.
- Brink, D. O. (1989). *Moral Realism and the Foundations of Ethics (Cambridge Studies in Philosophy)*. New York: Cambridge University Press.

- Charlow, N. (2013). The Problem with the Frege-Geach Problem. *Philosophical Studies*: 1-31.
- Chrisman, M. (2007). From Epistemic Contextualism to Epistemic Expressivism. *Philosophical Studies*, 135: 225-254
- _____ (2008). Expressivism, Inferentialism, and Saving the Debate. *Philosophy and Phenomenological Research* 2: 334-358.
- _____ (2010). Expressivism, Inferentialism, and the Theory of Meaning. In Michael Brady (ed.), *New Waves in Metaethics*. Palgrave-Macmillan.
- Dreier, J. (2004). Meta-Ethics and the Problem of Creeping Minimalism. *Philosophical Perspectives* 18: 23-44.
- _____ (2006). Negation for Expressivists: A Collection of Problems with a Suggestion for their Solution. *Oxford Studies in Metaethics, Vol. 1*. Oxford: Oxford University Press, 217-33.
- _____ (2009). Relativism (and Expressivism) and the Problem of Disagreement. *Philosophical Perspectives* 23(1): 79-110.
- Elder, C. (2004). Realism, Naturalism, and Culturally Generated Kinds. *Philosophical Quarterly*, 39: 425-44.
- Field, H. (2009) Epistemology without Metaphysics. *Philosophical Studies*, 143: 249-290.
- Finlay, S. (2005). Value and Implicature. *Imprint* 5:1-20
- Geach, P. T. (1965). Assertion. *The Philosophical Review* 74: 449-465.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Oxford: Clarendon Press.
- _____ (2003). *Thinking How to Live*. Cambridge, Mass.: Harvard University Press.
- Hare, R. M. (1952). *The Language of Morals*. Oxford Clarendon Press.
- Hawley, K. (2004). *How Things Persist*. New York: Oxford University Press, USA.
- Horgan, T., and Timmons, M. (2006). Cognitive Expressivism. In Horgan and Timmons, Eds., *Metaethics after Moore*. Oxford: Oxford University Press, 255-298.
- Horwich, P. (1998). *Meaning*, New York: Oxford University Press.

- _____ (1999). *Truth* (2 ed.). New York: Oxford University Press.
- _____ (2010). *Truth – Meaning – Reality*. New York: Oxford University Press.
- Hume, D. (1888). *A Treatise of Human Nature*. Clarendon Press, 1968.
- _____ (1748). *An Enquiry concerning Human Understanding*, edited by Tom L. Beauchamp, Oxford/New York: Oxford University Press, 1999
- Jackson, F. & Pettit, P. (1998). A Problem for Expressivism. *Analysis*, 58: 239-251.
- Joyce, R. (2006). *The Evolution of Morality*, MIT Press.
- Kölbel, M. (2003). Faultless Disagreement. *Proceedings of the Aristotelian Society*, 104 (1):53–73.
- Kripke, S. (1980). *Naming and Necessity*, Basil Blackwell.
- Lance, M. and Kukla, R. (2010). Perception, Language, and the First Person. In B. Weiss and J. Wanderer, Eds. *Reading Brandom*. New York: Routledge, p.115-128.
- Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Blackwell.
- Mabrito, R. (2009, August 31). Notre Dame Philosophical Reviews, *Review of Mark Schroeder's Being For*. Retrieved October 18, 2012 from <http://ndpr.nd.edu/news/24142-being-for-evaluating-the-semantic-program-of-expressivism/>.
- Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*, London: Penguin.
- McDowell, J. (1981). Non-Cognitivism and Rule-Following. In Steven Holtzman and Christopher Leich, eds, *Wittgenstein: To Follow a Rule*. Routledge. 141-62.
- Moore, G.E. (1903). *Principia Ethica*. Cambridge University Press.
- Portner, P. (2010). Permission and choice. Manuscript, Available from semanticsarchive.net.
- Price, H.(1988). *Facts and the Function of Truth*. Oxford: Basil Blackwell.
- _____ (1998). Three Norms of Assertibility, or How the Moa Became Extinct. *Noûs*, Vol. 32, Supplement: Philosophical Perspectives, 12, Language, Mind, and Ontology. 241-254
- _____ (2011). *Naturalism without Mirrors*, New York: Oxford University Press.
- Price, H., & Macarthur, D. (2007). Pragmatism, Quasirealism and the Global Challenge. In C. Misak (Ed.), *The New Pragmatists*. Oxford: Oxford University Press.

- Ryle, G. (1949). *The Concept of Mind*, Hutchinson and Co.
- Schiffer, S. (2003). *The Things We Mean*. Oxford: Oxford University Press, USA.
- Schroeder, M. (2008a). *Being For*, Oxford; New York: Oxford University Press.
- _____ (2008b). Expression for Expressivists. *Philosophy and Phenomenological Research* 76(1): 86-116
- Schueler, G.F. (1988). Modus Ponens and Moral Realism. *Ethics* 98, 492-500.
- Searle, J. (1965). What is a speech act? In P. P. Giglioli (Ed.), *Language and social context* (pp.136–154). Harmondsworth, England: Penguin Books.
- Shafer-Landau, R. (2005). *Moral Realism: A Defence* (New Ed ed.). New York: Oxford University Press, USA.
- Sidelle, A. (1989). *Necessity, Essence, and Individuation: A Defense of Conventionalism*. Ithaca: Cornell University Press.
- _____ (1998). A Sweater Unraveled: Following One Thread of Thought for Avoiding Coincident Entities. *Nous*. 32(4), 423-448.
- Smith, M. (2005). *The Moral Problem*. Oxford: Basil-Blackwell.
- Suikkanen, J. (2009). The Subjectivist Consequences of Expressivism. *Pacific Philosophical Quarterly*, 90: 364-387
- Thomasson, A. (2007a). Modal Normativism and the Methods of Metaphysics. *Philosophical Topics*, 1/2 135-160.
- _____ (2007b). *Ordinary Objects*. New York: Oxford University Press, USA.
- _____ (2008). Non-Descriptivism about Modality: A Brief History and Revival. *Baltic International Yearbook of Cognition, Logic and Communication*. Volume 4: 200 Years of Analytical Philosophy. 1-26.
- _____ (2009). Answerable and Unanswerable Questions. *Metametaphysics*. Eds. David Chalmers, Ryan Wasserman and David Manley Oxford: Oxford University Press.
- _____ (2013). Norms and Necessity. *Southern Journal of Philosophy*, 51(2): 143-160
- Timmons, M., & Horgan, T. (2006). Expressivism, Yes! Relativism, No! In R. Shafer-Landau (Ed), *Oxford Studies in Metaethics: Volume 1*. (73-97) New York: Oxford University Press, USA.

- Timmons, M. (1999). *Morality without Foundations: A Defense of Ethical Contextualism*. New York: Oxford University Press, USA.
- Unwin, N. (1999). Quasi-Realism, Negation and the Frege-Geach Problem. *The Philosophical Quarterly* 49: 337–52.
- _____ (2001). Norms and Negation: A Problem for Gibbard's Logic. *The Philosophical Quarterly* 51: 60–75.
- van Roojen, M. (1996). Expressivism and Irrationality. *Philosophical Review* 105: 311-35.
- Wedgwood, R. (2001). Conceptual Role Semantics for Moral Terms. *Philosophical Review*, 110 (1):1-30.
- _____ (2006). The Meaning of "Ought". *Oxford Studies in Metaethics, Vol. 1*. Oxford: Oxford University Press, 127-60.
- Williams, M. (2010). Pragmatism, Minimalism, Expressivism. *International Journal of Philosophical Studies*, 18: 317-330.
- Wright, C. (1992). *Truth and Objectivity*. Harvard: Harvard University Press.
- Yalcin, S. (2012). Comments on Mark Schroeder, Being For. *Central APA*.