



The Plausibility of Moral Error Theories

Item Type	dissertation
Authors	Knight, Casey Alton
DOI	10.7275/b3vc-2693
Download date	2025-04-19 20:47:48
Link to Item	https://hdl.handle.net/20.500.14394/39201

THE PLAUSIBILITY OF MORAL ERROR THEORIES

A Dissertation Presented

by

CASEY KNIGHT

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

May 2013

Philosophy

© Copyright by Casey Knight 2013

All Rights Reserved

THE PLAUSIBILITY OF MORAL ERROR THEORIES

A Dissertation Presented

by

CASEY KNIGHT

Approved as to style and content by:

Fred Feldman, Chair

Phillip Bricker, Member

Peter A. Graham, Member

Seth A. Cable, Member

Hilary Kornblith, Department Chair
Philosophy

For Mom

ACKNOWLEDGMENTS

I thank my friend Fred Feldman for reading and commenting on the chapters of this dissertation, and for instigating this project. I am in debt to Fred in virtue of his being a philosophical mentor to me during my time here at UMass. His influence on the way I do philosophy cannot be exaggerated. I hope that he is pleased with this work.

I am also grateful the other members of my dissertation committee: Phil Bricker, Pete Graham, and Seth Cable. Phil has played a crucial role in shaping every one of my philosophical views for the better. Pete's meticulous criticism of my ideas has helped me to become a far better philosopher than I otherwise would have been. I thank Seth for his service as an outside reader.

It has been a pleasure to be part of the philosophical climate at UMass. I hereby express my gratitude to everyone in the department, especially Chris Meacham, Maya Eddon, and Kristian Olsen. Chris has always provided speedy and insightful feedback on anything I ask him to help me with. I learned a lot about metaphysics from Maya by sitting in on her classes. I had many fruitful philosophical conversations with Kristian.

Before UMass, I had the privilege of being influenced by the climate at Western Washington University. I am grateful to all the philosophers at Western, especially Ned Markosian, Dan Howard-Snyder, and Hud Hudson. I give my heartfelt thanks to Ned for encouraging me to continue on in philosophy once I was done at Western. Dan was the first philosopher ever to sit down with me and help me through the rigorous construction of an argument. If Hud never gave me a prize for one of my papers in his class on the Empiricists, then I may never have been motivated to take philosophy as seriously as I do.

Without my family, I would not have been able to make it through the rough spots during the past few years. Kyle's upbeat attitude can always get me out of a slump. Whenever I

am bored, I can expect Dad to entertain me (for a while). Finally, I would not have been able to complete this project had Mom not instilled in me her perseverance.

ABSTRACT

THE PLAUSIBILITY OF MORAL ERROR THEORIES

MAY 2013

CASEY KNIGHT

B.A., WESTERN WASHINGTON UNIVERSITY

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Fred Feldman

The project that resulted in this work had two main goals. The first was to sort out the most plausible form of the moral error theory, the view made popular by J.L. Mackie in his *Ethics: Inventing Right and Wrong*. Second, I aimed to determine the extent of its plausibility. The first three chapters of this dissertation are the result of my attempt to accomplish the first goal, and the last two chapters are a consequence of the second. In the end, I argue that the most plausible version of the error theory is not true.

Along the way, I make three additional novel contributions to the meta-ethical literature. First, I provide a close reading of Mackie's work, and I argue that he has been misinterpreted thus far in the literature. Second, I attempt to clarify the linguistic and metaphysical components of the error theory. Third, I explicate a novel view in meta-ethics that isn't a direct descendent of Mackie's views, but which is, in many important respects, similar to his view.

Chapter 1, first, contains some preliminaries that are crucial for any work in meta-ethics. I discuss and explicate two distinctions involving the classification of meta-ethical views.

Realists and anti-realists about morality differ on whether the moral is in some sense “defective”. Cognitivists and non-cognitivists disagree over whether beliefs are essentially involved in moral judgments. In addition, I explicate the subject matter of the dissertation by explaining the concept of a moral judgment (or statement) and the concept of a moral thought (or attitude). Moral judgments are defined stipulatively *via* paradigm examples, and moral thoughts are defined as the psychological states expressed by moral judgments. The first chapter also contains a significant amount of historical exposition of Mackie, in order to set up a starting point for the project of the dissertation. I present and roughly explain six theses that he seems to accept and which seem to be foundational to his meta-ethical views. In addition, I go on to interpret him as holding a view that is a version of relativism about moral discourse, on which many of our moral judgments may come out true. My interpretation of Mackie differs from the standard interpretation of his view, on which most of our moral judgments are not true.

Error theories in morality are at least partially – if not primarily – views about moral language and discourse. Thus, with some important historical work as backdrop, Chapter 2 contains some discussion of the linguistic parts of these error theories. According to one of Mackie’s six foundational theses (what I call the ‘discourse thesis’), an illicit “presupposition” is embedded in moral judgments. I show that there are a variety of different ways of explicating the nature of this presupposition. It could merely be a pragmatic phenomenon, about what people do when they make moral statements. Or it could be a semantic phenomenon, about the linguistic requirements that must be met in order for moral statements to semantically express anything, or about conditions that the world must meet in order for moral propositions to be true. I argue that, on some renditions of the discourse thesis, the resulting error theories are either incoherent or too implausible to admit of much interest. I thus weed-out some of the less plausible versions of the view. According to the remaining versions, moral statements express propositions that entail the existence of defective kinds

of properties, and so, since those defective properties don't exist, our moral statements end up untrue.

In Chapter 3, I discuss a superficial but important metaphysical problem for these remaining versions of the error theory. I call it the 'problem of entailment'. The problem is that, in order for moral statements to entail anything, it seems that those defective properties must exist after all. Thus these versions of the error theory are *prima facie* internally inconsistent. My resolution of this problem involves some theorizing about the metaphysics of properties. I appeal to the idea that 'property' in ordinary and philosophical discourse is equivocal, and so the problem of entailment is merely superficial because it suffers from the fallacy of equivocation. At the end of the chapter, I explain in more detail the nature of the metaphysical theses that are natural for error theorists to accept.

With these important preliminaries in the background, in Chapter 4 I discuss Richard Joyce's version of the error theory. I interpret him as holding what I call 'the most plausible form of the error theory'. I explicate two arguments that he presents in favor of his view, and I show that both arguments are unconvincing. On the first argument, moral statements require that individuals have objective reasons, but there are no such reasons, and so moral statements are untrue. In reaction, I argue that talk of 'having objective reasons' is equivocal, and, as a result, the premises of the argument are plausible only if the argument is invalid. According to the second argument, moral judgments require the existence of a mysterious force that Joyce calls "practical oomph", but since there is no such force, our moral judgments are untrue. I reject the argument on the basis that Joyce's notion of practical oomph is too inchoate for the premises to be rationalized. A main upshot of the chapter is that the literature does not yet explicitly contain any compelling reasons to accept the plausible versions of the error theory.

Undermining arguments for the error theory does not, however, suffice to undermine the entire view. In Chapter 5, I attempt to provide a general consideration against the most plausible varieties of the error theory. My premise is that moral discourse contains too much

disarray for the truth of the error theory. I discuss some relevant literature in meta-ethics that seems to agree with my main premise, and I present a barrage of considerations in favor of the idea that moral language is semantically undecided. In addition, I explicate a framework on which moral language can be modelled as semantically undecided. So, while the dissertation concludes that the error theory is not true, this conclusion is premised on another somewhat extreme view about moral discourse. For it seems that I am in agreement with an important idea expressed by David Hume when he writes, “The chief obstacle ... to our improvement in the moral ... sciences is the obscurity of the ideas, and ambiguity of the terms.”¹

¹Hume, 1777: 61.

TABLE OF CONTENTS

	Page
ACKNOWLEDGMENTS	v
ABSTRACT	vii
 CHAPTER	
1. MACKIE'S ERROR THEORY	1
1.1 Introduction	1
1.2 The Uncontroversial Elements of Mackie's Position	2
1.2.1 The Questions of Realism and Cognitivism	2
1.2.2 Mackie's Major Theses	5
1.2.3 Mackie's Minor Theses	10
1.3 Moral Judgments and Moral Thoughts	14
1.3.1 Judgments, Evaluative Judgments, and Moral Judgments	15
1.3.2 Moral Thoughts	21
1.4 Mackie's Arguments for the Metaphysical Thesis	24
1.4.1 Objective Values and Unnatural Epistemology	25
1.4.2 Objective Values as Motivating	26
1.4.3 Objective Values as Supervenient	27
1.4.4 Objective Values as Reason-Giving	28
1.5 Mackie and Relativism	29
1.5.1 What is Explicit in Mackie's Writings?	30
1.5.2 Relativism about Mass and Morals	32
1.5.3 Mackie's Construction of Right and Wrong	35
1.5.4 Analogies and Disanalogies	39
1.6 Conclusion	41

2. ERROR THEORIES AND PRESUPPOSITION	43
2.1 Introduction	43
2.2 The Discourse Thesis in the Literature	44
2.2.1 Interpretations of Mackie in terms of Presupposition	45
2.2.2 Other Philosophers and Presupposition	45
2.2.3 Some Potential Caveats	46
2.3 Pragmatic Presuppositions	48
2.3.1 Stalnaker on Pragmatic Speaker Presupposition	49
2.3.2 The Discourse Thesis in Terms of Pragmatic Presupposition	51
2.3.3 Against Theses 2P-3P	53
2.4 Semantic Presuppositions	57
2.4.1 Presupposition: Semantic or Pragmatic?	57
2.4.2 The Discourse Thesis in Terms of Semantic Presupposition	59
2.5 Objections to the Fregean and Strawsonian Interpretations	63
2.5.1 A Worry for Theses 3F and 3S	63
2.5.2 Against Theses 2S-3S	66
2.6 Discourse Presuppositions and Entailment	72
2.6.1 Analytic Entailment and Non-Negotiability	72
2.6.2 The Discourse Thesis in Terms of Entailment	74
2.6.3 A Worry for Theses 3E and 3F	76
2.7 Conclusion	78
3. ERROR THEORIES AND ONTOLOGY	80
3.1 Introduction	80
3.2 The Metaphysical Thesis in the Literature	82
3.3 The Problem of Entailment	84
3.3.1 Formulations of the Discourse and Metaphysical Theses	84
3.3.2 Two Semantic Constraints	85
3.3.3 Statement of the Problem of Entailment	87
3.4 Why Rescue the Error Theory?	87
3.4.1 Parallels in the Philosophy of Mind and Science	89
3.4.2 Parallel to the Problem of Reference Failure	90

3.5	Pluralism about Properties, and its Rivals	92
3.5.1	Lewis on Pluralism	93
3.5.2	Terminological Foundations for Pluralism	94
3.5.3	Some Roles for Abundant and Sparse Conceptions	96
3.5.4	More Precise Characterizations of Pluralism and its Rivals	99
3.6	A Pluralist Solution to the Problem of Entailment	100
3.6.1	Four Versions of the Metaphysical Thesis	100
3.6.2	The Solution	102
3.6.3	The Epistemological Ordering of the Metaphysical Theses	104
3.7	Conclusion	108
4.	JOYCE'S ERROR THEORY	109
4.1	Introduction	109
4.2	Some Joycean Theses	110
4.2.1	Joyce's Metaphysical Theses	110
4.2.2	Joyce's Discourse and Error Theses	112
4.3	Joyce's First Argument	114
4.3.1	Formulation of the First Argument	115
4.3.2	Explanation of the First Argument	117
4.4	Criticism of Joyce's First Argument	121
4.4.1	Two Incomplete Objections	121
4.4.2	A Complete Objection	123
4.5	Joyce's Second Argument	128
4.6	Criticism of Joyce's Second Argument	133
4.6.1	Vagueness and the Sorites Paradox	133
4.6.2	An Application of this Approach to Joyce's Argument	135
4.7	Conclusion	138
5.	AGAINST THE ERROR THEORY	140
5.1	Introduction	140
5.2	Pieces of My Argument in the Literature	142
5.2.1	Kirchin's Tension	142

5.2.2	Burgess and Schiffer on Widespread Moral Indeterminacy	144
5.2.3	Lewis: Earlier and Later	149
5.3	A Framework for Indeterminacy	152
5.3.1	The Linguistic Account of Indeterminacy	153
5.3.2	Some Brands of Moral Indecisionism	157
5.4	Against the Error Theory	159
5.5	A First-Person Consideration	162
5.6	Some Inferences to Explanations	164
5.6.1	The Supervenience of the Moral on the Descriptive	165
5.6.2	Interest in Hume's Law	167
5.6.3	Open and Closed Questions	169
5.6.4	Rationally Irresoluble Fundamental Moral Disagreement	172
5.7	Conclusion	177
BIBLIOGRAPHY		181

CHAPTER 1

MACKIE'S ERROR THEORY

1.1 Introduction

In this chapter, I will explicate J.L. Mackie's "error theory" about morality, which he presents mainly in *Ethics: Inventing Right and Wrong*. The chapter is meant to serve as an introduction to the main metaphysical, psychological, and semantical issues that arise in meta-ethical discussions.

This work is a thesis in meta-ethics, so it is about the nature of moral reality, language, and thought. The first goal of the chapter is to situate Mackie's views, and views like Mackie's, in the broader meta-ethical terrain. A rough divide in meta-ethical topics is one between metaphysical questions – whether there is a moral reality, and what is its nature – and linguistic questions – whether moral discourse is descriptive and cognitive, and what are the crucial relations between moral discourse and reality. In section 2, I'll distinguish various general ways to answer these questions.

There is a need to clarify some of the subject matter of meta-ethics. The second goal of the chapter is to stipulate some technical terminology, which will help to clarify my topic, and which will be useful throughout this work. On the linguistic side, the main goal will be to explain the meaning of 'moral judgment', in order to explain what is moral discourse (section 3). On the metaphysical side, the main goal will be to explain, to a certain extent, the meaning of 'objective value' (section 4).

Mackie is the arch-error theorist about morality. The third goal of this chapter is to explicate his views in meta-ethics. I take a historical approach in my exposition of Mackie. One part of the approach will be to say in what respects we can come to determinate con-

clusions about what his views are. The upshot of this will be that, to a large extent, his positions in meta-ethics are under-specified (section 2). The second part of the approach will be to say, with regard to the under-specified parts of his views, what are the admissible specifications. The upshot of this will be that his views should be taken as roughly compatible with a wider range of meta-ethical views than is commonly supposed in the literature (section 5).

1.2 The Uncontroversial Elements of Mackie's Position

In this section, I'll present what I take to be relatively uncontroversial elements of Mackie's position – not uncontroversial in the sense that almost anyone will readily agree that they are *true*, but in the sense that almost anyone will readily agree that they are *his views*. Before I do that, though, I'll roughly explicate two main questions in meta-ethics.

1.2.1 The Questions of Realism and Cognitivism

Meta-ethicists hold views on the nature and existence of the elements of reality. Included in these broad metaphysical views are claims about the putative moral elements of the world and their natures. According to some, either there is no place in reality for anything moral, or else the moral is in some way defective. Perhaps morality isn't properly "natural" or "scientific"; perhaps it is in some way metaphysically "queer". These philosophers are anti-realists about morality; for short, 'anti-realists'.¹

According to others, there is a place in reality for the moral elements, and the moral is not in the suggested ways defective. Perhaps morality has some of these allegedly defective features – being unnatural, unscientific, or queer – but this is no defect. Or maybe morality

¹See, e.g., Ayer, 1952; Mackie, 1977, Joyce, 2001; Dreier, 1990, 2005; Harman, 1996; Lewis, 1989; and Schiffer, 1990.

doesn't have these features at all. These philosophers are realists about morality; for short, 'realists'.²

Thus, anti-realism is the view that moral reality is in some way defective. Realism is the view that moral reality is not in this way defective. There are two important, perhaps idiosyncratic, things to note about the way I am using this terminology. First, the anti-realist's crucial claim need not be about the non-existence of elements of moral ontology. She may claim that some elements of moral ontology – moral facts, properties, standards, propositions, or whatever – don't exist. Perhaps this is, in her view, the defect of moral reality. Alternatively, she may claim that, although there is a moral reality – there are moral propositions and properties, say – it bears some crippling defect.³

Second, the realist need not claim that the moral part of reality is "sui generis" or "free-standing". She may claim that there is no overlap between moral reality and the rest of the world; that moral facts do not reduce to some more fundamental part of reality. But she may, alternatively, claim that moral reality overlaps with, say, naturalistic elements of the world. The former sort of view is, broadly speaking, a brand of 'anti-reductionism'; the latter is a brand of 'reductionism'.⁴

Meta-ethicists also hold views about the nature of language and thought. These views help to fix their standpoints about ordinary moral discourse – about the nature of moral language and thought. According to some, it is (in some sense) "essential" to any sincere moral

²See, e.g., Moore, 1903; Huemer, 2005; Shafer-Landau, 2003; Brink, 1989; Balaguer, 2011; and Boyd, 1989.

³This distinction in kinds of anti-realism is especially important in order for an understanding of Schiffer's (2003, 1990) position.

⁴Moore (1903) is standardly thought to be an anti-reductionist; see also Huemer, 2005; and Shafer Landau, 2003. Jackson's (1998) view seems to be a brand of reductionism. The views of Brink (1989) and Boyd (1989) are supposedly "naturalistic" versions of anti-reductionism.

judgment that it involves the expression of a belief about moral reality. These philosophers are cognitivists about moral discourse; for short, ‘cognitivists’.⁵

According to other philosophers, a sincere moral judgment need not involve the expression of belief about moral reality. Different philosophers rationalize this position in different ways. Perhaps moral judgments are inherently “prescriptive”, and thus are never expressions of belief. Or maybe there is nothing – no moral reality – for moral judgments to be about, so it makes no sense to talk of moral belief. These philosophers I will call ‘non-cognitivists about moral discourse’; for short, ‘non-cognitivists’.⁶

There is one important, perhaps idiosyncratic, thing to note about the way I am using ‘cognitivism’: there are roughly two ways to be a cognitivist. First, one might claim that moral thoughts – the psychological attitudes that are expressed by people’s moral judgments – are naturally tied with certain beliefs, but do not reduce, even partially, to those beliefs. Such a view might hold that moral thoughts are, as a matter of psychological law, tied to states of belief. Second, a philosopher might claim that the beliefs that are necessarily tied to moral thoughts reduce (partially or wholly) to those beliefs. This view may hold that part of the analysis of moral thought will involve reference to beliefs of a certain kind. Either way, in my terminology, the philosopher qualifies as a cognitivist.

It is worthwhile to consider an example. Suppose I hold a view according to which, associated with any imperative judgment, one must hold a particular belief and a particular desire. On this view, whenever someone sincerely issues a command, to, say, shut the door, that individual, as a matter of psychological law, believes that the door is not shut. The proponent of such a view might simply hold that this is a contingent matter of fact, falsified in worlds where the laws of psychology are different. If she holds this claim, she might go on to analyze the thought expressed by a demand that the door be shut in terms of the

⁵See, e.g., See, e.g., Moore, 1903; Huemer, 2005; Shafer-Landau, 2003; Brink, 1989; Boyd, 1989; and Schiffer, 1990.

⁶See, e.g., Ayer, 1952; Hare, 1999; Gibbard, 1990, 2003; and Kalderon, 2005.

speaker's desire for the door to be shut. Even so, this theorist qualifies as a cognitivist about imperative judgments, in my terminology.

The cognitivist might instead hold that belief is more strongly essential to the thoughts expressed by imperative judgments. On this alternative kind of cognitivism, when we write out the correct analysis of an imperative thought, the analysis must refer to a belief of this type. Put another way: as a matter of absolute metaphysical necessity, if someone commands that the door be shut, then she believes that the door is not shut. Someone simply doesn't count, in any possible world, as genuinely commanding that the door be shut, unless she also believes that it is not shut. This view qualifies as the stronger form of cognitivism about imperative judgments. (Of course, it seems not to be a very plausible view.)

So we have a pair of questions: what is the nature of moral reality, and what is the nature of moral language and thought? And we have contradictory positions about each question, resulting in four packages of views: realist cognitivism, anti-realist cognitivism, realist non-cognitivism, and anti-realist non-cognitivism. In the following sections, I will put these distinctions to work in explicating Mackie's views in meta-ethics.

1.2.2 Mackie's Major Theses

I prefer to distinguish between Mackie's major theses and his minor theses. In this subsection I explain his major theses: the parts of his position that, as I see it, play the biggest role in shaping his overall meta-ethical view. In the next subsection, I present some of his minor theses: claims that play less of a role, but an important one nonetheless.

I can discern three major theses in Mackie's *Ethics: Inventing Right and Wrong* (hereafter: *EIRW*; unless otherwise noted, citations in this chapter are to *EIRW*.), which he summarizes, somewhat sporadically, in various places. He writes,

There are no objective values. (15)

...although most people in making moral judgements implicitly claim, among other things, to be pointing to something objectively prescriptive, these claims are all false. (35)

The first part of this citation involves an ontological claim, and so I take it to be partially determinative of Mackie's views on the debate between realism and anti-realism. The second part is a thesis about moral discourse, about what someone does when he makes a moral judgment. So it helps to determine his position on the debate over cognitivism. The final phrase in the passage seems vaguely to suggest what is Mackie's position about the crucial semantic relations between moral discourse and reality.

Consider, first, the ontological claim. It is this claim that fixes Mackie as an anti-realist:

Thesis 1: The Metaphysical Thesis. There is something seriously defective with moral reality, viz., it is supposed to be objective.

My statement of this thesis may seem odd to the reader. Doesn't Mackie clearly commit himself to the denial of an existence claim? Doesn't he explicitly say that objective values do not exist? I do not present Mackie's metaphysical thesis simply as the denial of the existence of objective values because it is somewhat controversial whether his statement should be interpreted in this way.⁷ Insofar as I wish to present Mackie in as uncontroversial a manner as possible, I avoid this natural interpretation. In addition, I would like Thesis 1, in this work, to be partially definitive of error theories, and so it is useful *not* to have it strictly a non-existence claim: there are error theories that seem not to reject moral ontology, but instead merely claim that it is defective.⁸

There are two important parts of Thesis 1 that require immediate attention. First, it is unclear what, on Mackie's view, would be the fundamental ontological categories of moral reality. What are the ontological categories of the elements into which moral reality decomposes? Is it comprised fundamentally of moral objects and properties; moral states-of-affairs; moral propositions; moral standards? It is my opinion that there is no clear, determinate answer to this question. However, Mackie suggests that objective values would have to be properties. I retain this supposition throughout this work.

⁷See, e.g., Hare, 1981; 1985; also chapter 3 of this work.

⁸Cf. Schiffer, 2003 and 1990.

The second point that requires clarification is what is meant by ‘objective’ in ‘objective values’ and ‘objectively prescriptive’, both in Thesis 1, and in the above passage from *EIRW*. Again, I believe that it is impossible perfectly to explicate the meaning of Mackie’s use of ‘objective’.⁹ But there are several suggestions about its meaning, which I will turn to in Section 4 of this chapter, when I explicate Mackie’s arguments for Thesis 1.

Contained in the above passage is a claim about moral discourse, about what happens when someone makes a moral judgment. This thesis is repeated in various ways throughout the first part of *EIRW*. He writes,

...ordinary moral judgements include a claim to objectivity, an assumption that there are objective values... (35)

...a belief in objective values is built into ordinary moral thought and language... (48-9)

As I understand this thesis, it strongly suggests that Mackie is a cognitivist about moral discourse.

Thesis 2: The Discourse Thesis. All moral judgments involve a claim to the effect that morality is objective.

The discourse thesis, however, does not decide whether, in Mackie’s view, moral thoughts contain, as parts, beliefs in objectivity, or whether these beliefs are merely necessarily tied (perhaps as a matter of general psychological law) with moral thoughts.

Consider the following ways of filling out Mackie’s view; they correspond to the two ways of being a cognitivist that I previously attempted to explain. On the first way, each moral judgment just is the expression of a belief that an objective value holds of some object. If this were the correct way of interpreting Mackie, then he would be a cognitivist of a strong sort: moral thoughts just are beliefs. On the second way, each moral judgment is the expression of some feeling of approval or disapproval about some object, but associated with these feelings are beliefs that objective values hold of certain objects. If this were

⁹See, e.g., Sinnott-Armstrong, 2010; Joyce, 2001: 16-7, for an introduction to Mackie’s use of ‘objective’.

Mackie's view, then, again, he would be a cognitivist, but of a weaker sort: moral thoughts, although they always *actually* carry with them beliefs, are not themselves beliefs.

So the discourse thesis does not decide an important question in moral psychology. It does not decide whether moral judgments are, as a matter of analysis, expressions of belief; or whether they express some other kind of psychological attitude that is naturally tied to belief. It does not tell us what is the general nature of moral thought. I will turn to what Mackie has to say about this question below, in section 3 of this chapter.

There are two further important features of Thesis 2. First, the notion of a moral judgment, as it is used in Thesis 2, must be clarified. I will use 'moral judgment' as a piece of technical terminology, and so I will make some important stipulations about the term. These stipulations will be in play throughout this work. Section 3 is also partially devoted to providing a meaning for 'moral judgment'.

Second, the "claim to objectivity" that Mackie sees in ordinary moral judgments can be taken in, roughly, three ways. It may be a pragmatic phenomenon. Then Thesis 2 will be a claim about what people *do* when they make moral judgments.¹⁰ The claim can, instead, be given a strongly semantic reading. On this reading, Thesis 2 is naturally taken as a thesis about requirements that must be met in order for pieces of ordinary moral language have *meanings*.¹¹ Finally, the claim to objectivity may be taken as a conceptual claim. Taken in this way, Thesis 2 is a claim about the *concepts* that are associated with moral terminology.¹² I discuss at length the question of how to interpret the discourse thesis in chapter 2 of this work.

I turn now to the final major part of Mackie's position. This is what he calls his "error theory" about morality. He writes,

¹⁰Hare, 1999; and Huemer, 2005 seem to interpret the claim to objectivity in this way.

¹¹This is, roughly, the way Shafer-Landau (2003 and 2005) seems to read Mackie.

¹²This is the way Joyce (2001) seems to read Mackie, at least in certain passages.

...a belief in objective values is built into ordinary moral thought and language, but ... this ingrained belief is false. (48-9)

This is something of the capstone of Mackie's position. It depends upon Theses 1 and 2: it is, more or less, an inference based upon them.

Thesis 3: The Error Thesis. Since objectivity is a defect in morality, and moral judgments involve a claim to objectivity, moral judgments are infected with widespread error.

Thus Mackie believes that there is a mistake that people make when they are judging things morally. But it is very unclear what exactly Mackie thinks is the mistake. There is a weak reading, and there is a strong reading, corresponding very roughly with the various ways of taking the claim to objectivity.

On the weak reading, people merely make the error of trying to "point" to something that is not there, much in the same way that Macbeth tries to ostend his floating dagger. It is clear that Mackie believes at least that people make this error (see, for example, Thesis 6, discussed below). However, it is important to note that this weak reading of Thesis 3 is not a semantic claim about what is required in order for a moral judgment to be true. It is thus consistent with moral judgments generally being true. In Section 5 of this chapter, I will argue that this is indeed an admissible interpretation of Mackie's views about the error in moral judgment.

On the stronger reading of Thesis 3, people systematically make the error of uttering claims that are not true. According this stronger reading, there would have to be objective values in order for moral judgments to be true. The strong reading, of course, is consistent with a wide variety of different ways of failing to secure truth. Moral judgments may be false, indeterminate, contentless, or even meaningless combinations of words. In chapter 2 of this work, I'll further discuss these ideas.

1.2.3 Mackie's Minor Theses

In this subsection, I explicate a variety of minor Mackian theses that are important to a proper understanding of his general position. In addition, since contemporary error theories are rooted in Mackie's views, it will be important to keep them in mind for my discussion of, e.g., Richard Joyce's error theory in Chapter 4.

Mackie holds a general thesis about the meanings and contents of evaluative utterances. In *EIRW*, he writes,

We can then offer a general definition of 'good': *such as to satisfy the requirements (etc.) of the kind in question.* ... This general definition covers different uses of the word 'good' ... because it leaves open just how the requirements in question are specified or indicated; it leaves room for different interests to be fed in in different ways in different sorts of cases. (55-6)

This passage is an echo of his earlier work. In "Aesthetic Judgments – A Logical Study", he writes,

In making an evaluative judgement we are at once saying or hinting that the object judged has certain natural characteristics, and commending it, or perhaps condemning it, on that account. The good-making characteristics are natural features which we are using as criteria or standards of value: we commend something which satisfies these standards, and condemn something which fails to satisfy them. Evaluative judgement is relative to standards or criteria, and presupposes them. (65)

The following analogy, I believe, helps to elucidate Mackie's claim.

Consider our judgments involving tallness. At a basketball game, I utter, 'Mike is tall'. You might dissent, claiming that six-and-a-half feet is, all things considered, pretty short for a basketball player. At a cocktail party full of professors, you utter, 'Mike is tall'. It would be reasonable for me to agree, perhaps noting that everyone at the party is relatively short in comparison to Mike. Intuitively, since we utter the same sentence, and we are speaking the same language, the meanings of our utterances do not differ. However, intuitively, the contents of our utterances differ: our utterances express different propositions. A possible explanation of this is that, although the meaning of 'tall' doesn't vary between contexts,

its content does, because, in the different contexts, we have in mind different standards of tallness. Intuitively, it is harder to be judged tall relative to the standards I have in mind in the basketball context than it is to be judged tall relative to the cocktail context. ‘Tall’ is an indexical: the content of a tallness judgment is determined both by some salient standard of tallness, and by the meaning (or character) of the whole uttered sentence. As a result, the contents of tallness judgments can vary between contexts.

As I understand Mackie, he means to accept the following (admittedly rough) thesis about evaluative judgments:

Thesis 4: Evaluative Indexicalism. The content of an evaluative judgment is partially fixed, in context, by some salient standard of evaluation, and partially by the meaning of the the sentence that is uttered in making the judgment.

There are some important things to note about the notion of a standard of evaluation. First, in the first of the above passages, Mackie indicates that standards of evaluation involve, or are comprised from, certain natural features of the world. Second, these natural features need not be mind-independent features: some may involve our desires or interests, and some might not. Third, the natural features from which standards of evaluation are composed are supposed to serve as a basis for our affective attitudes, viz., commendation and condemnation.

The thesis of evaluative indexicalism involves some new terminology. The terms ‘content’ and ‘meaning’ require some explanation. The discussion in Section 5 will help to fix how these terms are used in my statement of evaluative indexicalism. I’ll explain the notion of an evaluative judgment in Section 3 as a generalization of the notion of a moral judgment.

Mackie holds some general views about the nature of reasons. In *EIRW*, he writes,

Confining ourselves to human agents and their choice of action, we might ... hope to determine what people ought to do by seeing what can count as reasons for action. There seem to be several kinds. Most obviously, we would say that there is a reason for *a*’s *G*-ing, or that *a* has a reason to *G*, if *G*-ing would lead to the fulfillment of some desire or purpose or ideal that *a* now has, and *a* knows this. (77)

Here he accepts that someone has a reason to do something if he knows that a present desire of his would be satisfied. He goes on,

Someone can have a reason ... for doing what will lead or is likely to lead or even is wrongly believed by him to be likely to lead to the satisfaction, perhaps in the remote future, of some desire (etc.) that he now has. (77-8)

Here he accepts that someone has a reason to do something if he has some positive credence that a desire of his would be satisfied. Mackie even allows that someone might have a reason to do something when that person has no positive credence that it will promote his desires. He goes on,

But what if he will have (and knows that he will have) some desire or purpose at some future date, and something that he can do now is likely to lead to its fulfillment; does this constitute a reason for his now doing this? ... Do the desires and especially the sufferings of other people, if known to me, constitute a reason for me to do something, if I can...? (78)

Since he answers ‘Yes’ to these questions, he accepts that others’ desires can be a basis for someone’s reasons to act.

As I understand Mackie, he means to accept some thesis to the following effect:¹³

Thesis 5: Humeanism about Reasons. Whenever a person has a reason to do something, there are, or he believes there are, some desires, interests, or ends that would be promoted by his doing it.

It is important to notice how loosely ‘promotion’ is construed in Thesis 5.¹⁴ In addition, the desires that ground a reason can come from just about anywhere: they may be future desires, the desires of others, or even the ends of some “institution”. The desires need not even actually exist; it may merely be that someone has a positive credence in their existence.

Thesis 5, which is part of the basis of Richard Joyce’s error theory about morality, will be important in chapter 4 of this work, where I explicate Joyce’s position.¹⁵

¹³Compare the formulation of Mackie’s position on reasons in Phillips, 2010.

¹⁴Compare Schroeder’s formulation of Humeanism about reasons in his 2007: 110-3.

¹⁵Joyce, somewhat deceptively, calls his view about reasons a ‘non-Humean’ account of reasons. He inherits this terminology from Smith, 1994. In my opinion, it doesn’t matter what name we attribute to the thesis; what matters is the content of the thesis, and whether it is true.

Finally, Mackie holds a general thesis about the psychological process involved in the making of moral judgments. In *EIRW*, he writes,

If we admit what Hume calls the mind's 'propensity to spread itself on external objects', we can understand the supposed objectivity of moral qualities as arising from what we can call the projection or objectification of moral attitudes. (42)

In earlier work, he writes,

...in using moral terms we are as it were objectifying our own feelings, thinking them into qualities existing independently of us. For example, we may see a plant, say a fungus, that fills us with disgust, but instead of stating that we have this feeling, or merely expressing and relieving it by an exclamation, we may ascribe to the fungus a semi-moral quality of foulness, over and above all the qualities that a physical scientist could find in it. ... The feeling and the supposed quality are related as a seal or stamp and its impression. (1946: 81-2)

As I understand Mackie, he is explaining what he thinks is the causal chain that leads to many of our moral judgments.¹⁶ First, we see something happen, or think about some possible state of affairs. Perhaps, in reaction to this, we experience a feeling of approval; perhaps a feeling of condemnation; perhaps we think that someone else might have such a feeling about the state of affairs. Finally, we project onto that state of affairs a certain feature, viz., an objective value, that is supposed to be internally related to that feeling of approval or condemnation.

Mackie calls the process that leads up to this projection the 'objectification of values':

Thesis 6: The Objectification of Values. When someone makes a moral judgment, he projects onto the world his, or someone's, affective attitude about something.¹⁷

¹⁶Joyce (2010 and 2008) has an interesting discussion of these kinds of moral judgments, which he calls 'paradigmatic moral judgments'. On his view, some of our judgments (the non-paradigmatic ones) have a testimonial causal heritage: perhaps a mentor simply impresses this judgment on us; perhaps we are brainwashed into having them. Other judgments (the paradigmatic ones) have causal heritages that are more internal to us: they are the product of our own moral sensibilities. In my opinion, there may be no actual clear-cut cases of Joyce's paradigmatic moral judgments. It seems to me plausible that the causal heritages of all of our moral judgments are a jumbled mishmash of many factors, at least some of which involve the conditioning of mentors, as well as other "societal pressures".

¹⁷Mackie discusses this further in his 1980; see also Joyce, 2010.

There are two important things to note about Thesis 6. First, the notion of ‘projection’ is unclear: is projection an irreducible psychological relation that bears between a person, an affective attitude, and an object in the world? Or can the notion of projection be reduced to more well-understood psychological notions? For example, perhaps someone projects a quality onto the world when he believes that the quality holds of a certain object. I’ll discuss this issue in section 3 of this chapter.

Second, I believe that the thesis of the objectification of values is importantly related to Thesis 2, the thesis about the claim to objectivity. However, they are not identical. Thesis 6 is more specified in two respects. Whereas Thesis 2 makes no claim about the relation between supposed objective values and the affective feelings of people, Thesis 6 entails that there are purported necessary connections between these items. Mackie writes, “The objective quality is not simply the feeling itself transferred to an external object, but is something that would inevitably arouse that feeling.”¹⁸ And whereas Thesis 2 isn’t obviously an empirical claim about moral psychology, Thesis 6 makes a claim about the causal process that results in a psychological attitude called ‘projection’. Thus, as I understand Thesis 6, it purports to give us important information about the nature of moral thought. I’ll discuss the implications of Thesis 6 further in the next section.

I have made many promises, and thus incurred many obligations. It is time to start fulfilling them.

1.3 Moral Judgments and Moral Thoughts

In subsection 3.1, I explain how I will use the terms ‘judgment’, ‘evaluative judgment’, and ‘moral judgment’, throughout this work. What I say in this regard should be thought of as mostly a matter of stipulation. However, my stipulations fit well enough with the use of these terms in the literature, and they circumvent certain possible technical problems for the

¹⁸1946: 81-2.

views I'll discuss in this work. (I hope they do not raise further technical difficulties!) In subsection 3.2, I clarify what I believe are Mackie's views on the nature of the psychological attitudes involved in making moral judgments. A general assumption of the literature is that Mackie is a cognitivist, but it is not at all obvious to me what is the justification for this assumption. I'll argue that he should determinately be understood as a strong sort of cognitivist about morality.

1.3.1 Judgments, Evaluative Judgments, and Moral Judgments

In this work, I will use 'judgment' to designate a broad class of phenomena. The class of judgments will include beliefs and assertions. Thus, when someone asserts something to be the case, he judges it to be true; when someone believes something, he judges it to be the case. While I take 'judgment' to refer both to these public linguistic acts as well as to private cognitive phenomena, I'll refer to the agent of the judgment, by default, as the 'speaker', though I may sometimes refer to the agent as the 'thinker'.

I'll also use 'judgment' to talk about particular utterances of sentences. When someone utters, 'It is raining', his uttering it on that occasion is a judgment. I'll assume this, whether or not he asserts that it is at the time raining, and whether or not he believes that it is raining in his circumstances. For example, the speaker might be joking, lying, or engaged in a game of make-believe. I'll also often use 'statement' to refer to these tokenings of sentences. Thus a statement need not be an assertion (or an assertive utterance), in my terminology.

I'll even use 'judgment' to refer to the assumptions of speakers in context, so that a speaker's pragmatic presuppositions, background beliefs, or tacit beliefs are included in his judgments. Suppose I ask you, 'How old is the present King of France?' In doing so, I most likely presuppose that there is a present King of France, and I therefore have judged that there is one.

Finally, I stipulate that 'judgment' is to apply to what are intuitively non-cognitive and non-propositional sorts of phenomena: viz., command sentences and expressions of affec-

tion. Suppose I command you to shut the door, uttering as I do, ‘Shut the door!’ I will refer to what I do in this context as a judgment. Likewise, suppose that I open my refrigerator, hoping to eat some of last week’s leftovers. Finding a certain foul-smelling fungus in my leftovers, I say ‘Ewww!’ and toss them in the compost. My expression of disapproval of the leftovers’ smell is, in my terms, a judgment.

There are evaluative judgments among our judgments. In making an evaluative judgment, a speaker purports to describe something as foul, desirable, good, wrong, fair, and so on. For each of these kinds of judgment, it seems that the speaker is committing himself to a certain “evaluation” of something or other.

Among the evaluative judgments are the moral judgments. An important question to ask is: what is it that delimits the moral judgments from the rest of the evaluative judgments; and from the rest of judgments, more generally? I’ll take the following approach in response to this question.¹⁹ There are some uses of evaluative terms that are paradigmatically moral uses. I will call these uses of evaluative terms ‘moral terms’.²⁰ Sentence utterances in which moral terms occur I will call ‘moral statements’. Moral judgments, broadly construed, are judgments that would naturally be made through the use of moral statements. Consider a few examples.

First example. A young woman who aligns herself with the feminist movement is parading outside a courthouse with other self-proclaimed feminists. She claims to believe that a woman has an inalienable right to bodily integrity, one that even trumps a fetus’s right to life. The sign she is holding has written on it some slogan to the effect that abortion

¹⁹Compare Jackson, 1998.

²⁰It is important to be clear here about how I’m individuating terms. On one way of thinking about this, there is just one term ‘bank’, and two different ways of using it. This is not how I am conceiving the individuation of terms. Instead, I assume that there is a class of uses of ‘bank’ where it is used to refer to money banks, and a different class where it is used to refer to riverbanks. Each of these classes corresponds to a different term. Thus, I am distinguishing terms in a more fine-grained way than I would if I were just distinguishing between them typographically. Nothing of importance hangs on this. I recognize the other more coarse-grained way as a perfectly good way of individuating terms, and it is useful for certain purposes. It is just that my way is more useful for my present purposes.

shouldn't be illegal; that it would be unjust for a government to take away this right. She is chanting, 'Abortion is permissible!' Her use of 'permissible', in this context, is a moral use of the term.

Second example. A middle-aged man is sitting at the breakfast table, reading the local newspaper. He sees a headline that reads 'Canine Killer Strikes Again', and he goes on to read about a police investigation into the mysterious drowning of dogs in the area. The police think there is a unique human criminal who has, thus far, drowned ten canines. Somewhat disturbed about the prospects for his own puppy, he says to himself, 'It's wrong for the murderer to kill those dogs – whether or not he first asks their owners.' His use of 'wrong' is a moral use of the term.

Third example. A high school student is learning about the persecution of the Jews during World War II. He finds out that they went through all sorts of painful experiences. They were starved, tortured, worked to death, beaten, and so on. He figures they must have felt a lot of pain. He disapproves of what happened to them; thinking about it invokes in him a strong feeling of repugnance. While giving a presentation on this topic, he announces to the class, 'All that pain the Jews experienced was bad.' His use of 'bad' is a moral use of the term.

These are moral uses of 'permissible', 'wrong', and 'bad'. Each such use corresponds to a moral term. I will not, however, restrict the application of 'moral term' just to the words in the above examples. There are also moral uses of the following terms: 'good', 'desirable', 'evil', 'should', 'ought', 'just', 'fair', and 'deserving'. In any case, the judgments of the individuals in the preceding examples are moral judgments, because they involve moral uses of evaluative terms.

There are many more moral judgments than the above paradigms. Utterances of the following sentences can be moral judgments:

What the murderer did was wrong.
You ought to save the drowning child, though he is an evil little boy.
Abortion is permissible.

Pain is bad.
Capital punishment is unjust.
If torture is an evil, then you shouldn't do it.
Either getting your little brother to steal is wrong, or stealing isn't wrong.

There are some important differences, however, between the first two of these sentences and the rest. Compare the first sentence with the third. *Prima facie*, the first sentence asserts of a particular action (what the murderer did) that it was wrong. Similarly for the third: *prima facie*, it asserts of an action (abortion) that it is permissible. But it seems that the third sentence has a feature that the first does not. It is natural to prefix the third sentence with 'sometimes', 'typically', or 'always', where this is unnatural for the first sentence. We get the following variants of the third sentence:

Sometimes, abortion is permissible. / Abortion is sometimes permissible.
Typically, abortion is permissible. / Abortion is typically permissible.
Always, abortion is permissible. / Abortion is always permissible.

Similar features hold of the fourth and fifth sentences. Thus, I take sentences like the third, fourth, and fifth, to be generalizations in a way that the first two are not. I take the difference between these sentences correspond more or less with the classical distinction between sentences whose main subject terms refer to act-tokens, and those whose main subject terms refer to act-types.²¹

Now compare the first two sentences with the last two sentences. *Prima facie*, the last sentence doesn't assert of anything that you shouldn't do it. Likewise for the final sentence. They are conditioned, or qualified, in a way that the first two sentences are not. The second to last sentence, intuitively, is compatible with the permissibility of torture. Likewise, the final sentence is compatible with it's being morally permissible to steal.

Finally, compare the first two sentences with each other. Whereas the first seems to make just one moral judgment (that a certain action was wrong), the second seems to make

²¹Note that, while it is unnatural to prefix the first sentence with one of these adverbs of quantification, it is not linguistically impossible. I might make perfectly good sense by uttering, 'Sometimes, what the murderer did was wrong', so long as we take the bound sentence to be general in the same way the third, fourth, and fifth sentences are general. It's just that I wouldn't naturally interpret that sentence in that way.

two (that you have a certain obligation, and that a certain little boy is evil). The first seems to express a simpler thought than the second. The second is a composite of more basic judgments.

I will call judgments like the first ‘basic moral judgments (or statements)’. Judgments like the last six sentences (conjunctive, general, or conditioned), I will refer to as ‘complex moral judgments (or statements)’.

With this distinction between complex and basic moral judgments in mind, we are in a position to clarify further Mackie’s Thesis 2, the discourse thesis. According to that claim, moral judgments essentially involve a belief in the objectivity of values. Thesis 2 therefore implies that, for example, the young feminist, the balding man, and the high school student, in making their respective judgments, are all somehow committed to the objectivity of values. I believe that we should not immediately reject these implications as totally absurd.

Consider, however, a *Fourth Example*. David Hume, a skeptic about morality, is trying to decide what to do. His little brother could steal for him a book from Cleanthes’ library. Hume really wants the book. Another friend, Philo, tells him that, since stealing is wrong, getting your little brother to steal is wrong. Hume replies: ‘I guess I already knew that. If stealing is wrong, then getting your little brother to steal is wrong. But that doesn’t say much. There probably isn’t any real quality of wrongness. So stealing probably isn’t wrong.’ Given my stipulated meaning of ‘moral judgment’, Thesis 2 implies that, in making this moral judgment, Hume is in something of a paradoxical situation. He has made a moral judgment, and therefore committed himself to the objectivity of values. But, being a moral skeptic, he rejects any such commitment.

I find this troublesome. It seems more or less absurd to say that Hume has made any kind of “claim to objectivity”.²² Conditional claims like this often seem too weak to incur

²²I have similar intuitions about certain negative moral judgments and ones that involve propositional attitudes: e.g., ‘It’s simply not true that stealing is wrong’ and ‘Philo believes that stealing is wrong’.

such a commitment. Notice, however, that a revision of Thesis 2 that only quantifies over the *basic* moral judgments does not have this implication. So understood, Thesis 2 would be silent on whether people are committed to the objectivity of values in making conditional moral judgments. I take this as a reason to prefer this weakening of the discourse thesis. Moreover, I doubt that Mackie intended to imply that someone like Hume might have been committed to objective values in making conditional moral judgments.

Hereafter, in this work, I will use ‘moral judgment’ and ‘moral statement’ to mean what I have stipulated ‘basic moral judgment’ to mean, unless otherwise noted. Thus, unless I explicitly say that I am talking about complex moral judgments, the subject matter of my discussion should be understood to be in this way restricted.

I can think of two worries for restricting my discussion in this way. Both involve the idea that this understanding makes the discourse thesis objectionably weak. First, when someone makes a conjunctive or general moral judgment, it seems that he is committed to whatever he would have been committed to had he made a more specific judgment. For example, in uttering, ‘Abortion is permissible’, our young feminist seems committed to at least what she would have been committed to, in the way of objective values, had she judged of a particular act of abortion that it was permissible. But my understanding of Thesis 2 doesn’t allow us to derive this result.

Second, a philosopher might accept a principle to the effect that standard deontic notions are inter-definable. According to such a view, there may be a basic deontic notion, say, wrongness. Other notions, such as permissibility, are definable in terms of wrongness: for an action to be permissible is for it not to be wrong. Such a philosopher might claim that, therefore, a judgment to the effect that an action is permissible is a complex moral judgment. Then Thesis 2 goes silent with regard to all moral judgments of permissibility.²³

²³See the discussion in Pigden, 2010.

I have no direct reply to these objections. My indirect reply is that it is profitable to abstract away from the issues that they raise. My discussion can maintain a comfortable level of precision, given the way I have set things up. But to introduce a barrage of qualifications in reply to these problems would clutter my discussion beyond comprehension. Thus, I acknowledge that these issues must be dealt with. But not by me; not right now.

1.3.2 Moral Thoughts

Correlated with a typical statement is a psychological attitude, or thought, that the statement expresses. Intuitively, a statement expresses a certain thought if, under ordinary circumstances, it would be linguistically appropriate for someone to make the statement only if she had the psychological attitude. Thus an audience to the statement might reasonably infer that the individual making the judgment has the thought.²⁴ When I say, ‘Snow is white’, in a certain context, and you are my audience, it is ordinarily reasonable for you to think that I have a belief in the proposition that snow is white. When I say, ‘I am sitting’, you may reasonably come to think that I believe I am sitting. Or suppose I utter the following claim: ‘I hereby express my belief that Moriarty has arrived.’ You ought to take my utterance to be an expression of my beliefs about Moriarty’s location.

The psychological attitude need not be a belief. When I find the foul-smelling fungus in my refrigerator, and exclaim, ‘Ewww!’, perhaps you should think I disapprove of the smell of the leftovers. In this situation, my utterance is an expression of disapproval. If I claim, ‘There’s probably no God’, you should think that my credence in the proposition that there is no God is high. Here, it is natural to take my utterance as an expression of my credences about religious matters. But it would be going too far to interpret me as having expressed a belief – I may be agnostic about the matter.

I’ll use ‘moral thought’ to refer to the psychological attitude thus expressed by a moral statement. There is an important general question about moral thoughts: what sort of psy-

²⁴Compare Schroeder, 2008: 28-34.

chological attitude are they? I will not argue for an answer to that question here. However, there is restriction of the question: what sort of psychological attitude does Mackie think moral thoughts are? This question I will answer. In short, I believe Mackie thinks moral thoughts are beliefs. There are, however, other possibilities, which are suggested by Thesis 6.

Recall Thesis 6, the principle of the objectification of values. According to this principle, when someone makes a moral judgment, his judgment is a result of an internal process of recognizing his own affective attitude, and then projecting this attitude onto the world. There are thus apparently two psychological states that Mackie thinks are involved in the making of a moral judgment: an affection and a projection.

Should we read Mackie as accepting the idea that moral thoughts are mere affections? On this interpretation of Mackie, he would hold that moral thoughts are just the feelings of commendation or condemnation that are allegedly part of the process of objectification. There are two reasons why we should not read him in this way. First, the affective attitude that is part of the process of objectification is supposed to be causally prior to the moral judgment. That is, on Mackie's view, we first have a feeling of commendation; this causes us to project our feeling onto the world, thus giving rise to our moral judgment. Second, there is some direct textual evidence that he rejected this idea. In his earlier work, he wrote, "...we do not think we are merely ejaculating when we talk in moral terms."²⁵ I take him here to be rejecting the idea that a moral judgment is merely an expression of condemnation or commendation.

I believe that, on Mackie's view, moral thoughts are the final result of the process of objectification. I.e., moral thoughts are, or reduce to, the "projections" of our feelings onto the world. But it is not at all clear what is Mackie's notion of projection. There are two general ways to understand the idea.

²⁵1946: 80-1.

On the first way, the psychological attitude of projection is not reducible to more basic psychological notions. Thus, when someone projects his feeling onto the world, a certain irreducible attitude holds between an object in the world and his feeling. I will not interpret Mackie in this way. I have two reasons for rejecting this interpretation of Mackie. The first, in brief, is that I cannot understand what this attitude would be. I cannot interpret Mackie as accepting a certain thesis, when I cannot understand the thesis.²⁶ The second reason is that there is an easily graspable alternative interpretation of Mackie, according to which projection is not some kind of irreducible attitude.

I interpret Mackie in a second way, according to which the attitude of projection reduces to other psychological attitudes: projection, for Mackie, is supposed to reduce to belief. There is textual evidence for this. He writes, "...we are as it were objectifying our own feelings, *thinking them into* qualities existing independent of us."²⁷ On the most natural understanding of this quote, Mackie is claiming that, when we make moral judgments, we believe there is a property that is had by something we have an affective feeling about.²⁸ Our belief, in such a situation, is our moral thought, i.e., the psychological attitude that would be expressed by the salient moral statement. Thus, for Mackie, someone has a moral thought just in case she believes that a certain kind of affective attitude is internally related to an objective value, and that this objective value is exemplified in something.

When interpreted in this way, Mackie is a cognitivist about moral psychology. In addition, this is the only rationale that I can concoct in favor of the idea that Mackie really is

²⁶Joyce (2010: 39) takes the thesis of objectification to involve the idea that, when someone projects his feeling onto the world, he "experiences" some object "as" having an objective feature. I simply do not know how to engage with this idea. What is it to experience something as being good, if it is not just to believe that goodness holds of that thing? Perhaps, one might claim, it could be understood by analogy to our color perceptions: something might "strike me as" being red, even while I do not believe that it is red. But the sense in which something might "strike me as" being good is different from the sense in which something might "strike me as" being red. In the moral case, to say that x strikes me as being good is just to say that I am initially inclined to believe that x is good. In the color case, to say that x strikes me as being red is to make a claim about my experience of it: it appears to me as being red.

²⁷1946: 81, my emphasis.

²⁸See also Mackie, 1980: 150.

a cognitivist. On his view, moral thoughts just are beliefs, in certain complex propositions involving the existence of objective value properties and their internal connections with certain affective attitudes. But what are these objective values supposed to be? What is it that Mackie thinks would make these properties “objective”? I will begin to answer these questions in the next section.

1.4 Mackie’s Arguments for the Metaphysical Thesis

So far, I have not had much to say about Mackie’s notion of objectivity. In this section, I elucidate the meaning of Mackie’s use of ‘objective’ by quickly explaining the arguments from *EIRW* to the effect that there can be no objective values.

There are a few preliminaries to note. First, I will only discuss versions of Mackie’s “argument from queerness”. I’ll not discuss his “argument from relativity”.²⁹ My justification for this is that a discussion of his argument from relativity would not help to elucidate Mackie’s notion of objectivity. It does not seem to involve any assumptions about the nature of objective values.³⁰

Second, the main theme of the argument from queerness goes roughly as follows: “Objective values, if they existed, would have to be queer entities; but it is implausible to think there are such entities; and so, there can be no objective values.” Thus understood, each version of the argument from queerness involves two steps: first, a claim about the essence of objective values; and second, a claim about whether objective values, so conceived, exist. The first step in each version is non-substantive: it is merely supposed to be a partial defi-

²⁹1977: 36-8.

³⁰The argument from relativity is roughly this: “There is widespread moral disagreement, and this is best explained by the nonexistence of objective values.” You might think that this *does* involve some assumptions about the nature of objective values, viz., that they would *cause* us not to be in widespread disagreement about morality, if they existed. Thus understood, objective values would be causally potent. But I do not think Mackie could have assumed this, because he seemed to think that, if there were objective values, they would stand outside of the causal order. I think it is something of a puzzle how to charitably understand the argument from relativity.

inition of the notion of objective value. The second step is supposed to be more substantive: it is about what is in the world.³¹

Finally, I will not have much commentary on the soundness of these arguments, nor on their substantive premises. My main goal is to explicate Mackie's idea of an objective value. Suffice it to say that I do not think the arguments, taken either individually or jointly, conclusively establish Mackie's metaphysical thesis.³²

1.4.1 Objective Values and Unnatural Epistemology

Roughly, according to Mackie's first version of the argument from queerness, if there were objective values, then we could have no knowledge of them, since they would not be part of the natural world. Mackie writes,

...if we were aware of [objective values], it would have to be by some special faculty of moral perception or intuition, utterly different from our ordinary ways of knowing everything else. (38)

The reason why knowledge of objective values could only come by way of a faculty of intuition is that they would have to be non-natural, mind-independent properties. Mackie does not have much to say about what he means by 'non-natural'.³³ It is reasonable to suppose, however, that for Mackie, non-natural properties are ones that do not figure in the natural sciences.³⁴ He does, however, present some kind of explanation of his notion mind-independence. He writes,

The ordinary user of moral language means to say something about whatever it is that he characterizes morally, for example a possible action, as it is in

³¹Compare Joyce, 2001: 16-7; and Smith, 1994: 65.

³²For discussions of Mackie's arguments, see Sinnott-Armstrong, 2006; Huemer, 2005; Shafer-Landau, 2003; Joyce, 2001; and Smith, 1994.

³³See, e.g., Mackie, 1977: 31-2.

³⁴This is a somewhat standard rendering of 'non-natural'. See, e.g., Shafer-Landau, 2005 for references. I do not think this sort of account provides a sufficient explanation of the notion of non-naturalness, but I'll not press the issue here. In general, in this work, I try to steer away from issues directly pertaining to the naturalism/non-naturalism debate.

itself, or would be if it were realized, and not about, or even simply expressive of, his, or anyone else's, attitude or relation to it. (33; cf. 29-32, 35)

As I understand Mackie, his claim here is that, if there were objective values, then whether they held of certain objects would depend on the intrinsic natures of objects external to our brain states. They wouldn't depend on our psychological attitudes.

So, in the context of this first version of the argument from queerness, Mackie conceives of objective values as properties that the natural sciences don't purport to study, and that are external to our psychology. According to Mackie, properties like this could only come to be known *via* a "special faculty of moral intuition".

1.4.2 Objective Values as Motivating

On this second version of his argument from queerness, Mackie claims that objective values would have to be essentially motivating. He writes,

Plato's Forms give a dramatic picture of what an objective value would have to be. The Form of the Good is such that knowledge of it provides the knower with both a direction and an overriding motive; something's being good both tells the person who knows this to pursue it and makes him pursue it. An objective good would be sought by anyone who was acquainted with it, not because of any contingent fact that this person or every person, is so constituted that he desires this end, but just because the end has to-be-pursuedness somehow built into it. (40)

It is unclear how strongly to take Mackie's claim here.³⁵ On one interpretation, if there were objective values, then it would be impossible for them to be exemplified in something without people having certain motivational reactions to them. On a weaker reading, people would only have to have motivational reactions to their beliefs that objective values were instantiated. According to either of these readings, the notion of necessity involved is supposed to be broad, metaphysical necessity. But there is an even weaker reading, according

³⁵See Dreier, 1990, 2010; and Sinnott-Armstrong, 2010, for good discussions of this point. Lewis (1989) accepts a view on which there is an 'iffy' connection between values and motivation.

to which, as a matter of general psychological law, people have motivational reactions to their beliefs involving objective values.

In any case, Mackie attempts, while giving this argument, to draw a strong connection between objective values and motivation. He conceives objective values as (qualifiedly) essentially motivating properties. Thus understood, we can add to Mackie's definition of objective values: they are non-natural, intrinsic, essentially motivating properties, ones that we can only grasp via a faculty of intuition.

1.4.3 Objective Values as Supervenient

According to Mackie's last version of the argument, the queerness of objective values would consist in their supervenient nature. He writes,

What is the connection between a natural fact that an action is a piece of deliberate cruelty – say, causing pain just for fun – and the moral fact that it is wrong? It cannot be an entailment, a logical or semantic necessity. Yet it is not merely that the two features occur together. The wrongness must somehow be 'consequential', or 'supervenient'; it is wrong because it is a piece of deliberate cruelty. But just what *in the world* is signified by this 'because'? (41)

So, although objective values are, for Mackie, supposed to be distinct from natural properties, they are supposed to hold in virtue of these properties. There still are supposed to be necessary connections between the natural features of the world and objective values. I.e., for any objective value, there are some natural properties, whose instantiation would guarantee the instantiation of the objective value. These necessary connections would not hold in virtue of logic; Mackie thinks one cannot simply derive an 'ought' from an 'is'. Nor would they hold in virtue of analytic facts about the meanings of evaluative and descriptive terms.³⁶ They are, instead, supposed to be "brute" metaphysical necessities.

This alleged connection between objective values and natural properties provides us with another feature of Mackie's conception of objective values. In his view, an objective

³⁶See, e.g., Mackie, 1977: 64-73.

value is non-scientific, intrinsic, essentially motivating, brutally supervenient, and such that it can only be known through the lens of intuition.

1.4.4 Objective Values as Reason-Giving

There is another version of the argument from queerness, which seems implicitly to be in Mackie's work, although he does not explicitly endorse it in *EIRW*.³⁷ According to this line of argument, the queerness of objective values would consist in their being "objectively reason-giving". Mackie writes,

A categorical imperative, then, would express a reason for acting which was unconditional in the sense of not being contingent upon any present desire of the agent to whose satisfaction the recommended action would contribute as a means. ... So far as ethics is concerned, my thesis that there are no objective values is specifically the denial that any such categorically imperative element is objectively valid. The objective values which I am denying would be action-directing absolutely, not contingently ... upon the agent's desires and inclinations. (29)

On Mackie's view, then, if there were objective values, then whenever they held in the world, some reasons for action would also come into existence. Since objective values are mind-independent, on Mackie's view, the existence of these reasons would not have to depend on states involving our psychology.³⁸

Provided with these considerations involving Mackie's conception of objective values, we are now in a position to state a rough definition of 'objective value', at least as Mackie seems to use the word. It is useful to think of this definition as another one of Mackie's minor theses:

Thesis 7: Mackie's Conception of Objectivity. Objective values would have to be non-natural, intrinsic properties; nonetheless, they would have to supervene on natural properties; they would also have to be essentially motivating and reason-giving; and, to top it all off, they could only be knowable *via* a special faculty of intuition.

³⁷See 1977: 49, where Mackie quickly lists off his arguments for his metaphysical thesis.

³⁸This is roughly related to Joyce's argument for the error theory. I will discuss it (or something like it) more fully in chapter 4 of this work. See also Smith, 1994: 64-6.

I genuinely believe that Thesis 7 embodies Mackie's conception of objective values, at least as it is presented in the first chapter of *EIRW*. Hereafter, however, I will not use 'objective value' with the explicit meaning that we might take Thesis 7 to inject into it. This is because, usually, when I use 'objective value', what I will be trying to communicate will not depend on Mackie's apparent understanding of the notion. Rather, the information I'll convey through my use of 'objective value' will only depend on the assumption that objective values are a certain kind of property.

It is important to be clear about why I have gone through the process of explicating Mackie's notion of an objective value. First, it is important, for historical reasons, to get Mackie's views about objective values on the table. More recent versions of the error theory should be seen roughly as elaborations of Mackie's approach, and so his notion of an objective value helps to expose error theories in general. Second, it is important to see what sorts of features are supposed to be defective, according to the error theorist. By considering what Mackie took to be queer characteristics of morality, we are better able to understand why a more contemporary error theorist might reject morality.

1.5 Mackie and Relativism

In this section, I'll continue to explicate what I take to be a good candidate for Mackie's official position. I'll argue that Mackie's explicit views are under-specified in such a way that he can easily be understood as a sort of relativist about morality. This is not acknowledged in the literature – indeed, there is no commenter on Mackie's writings who interprets him in anything near the way I do. I therefore take it to be an interesting result.

According to a relativist about morality, we can understand moral judgments to be true only relative to moral standards.³⁹ Such a view holds that, even if ordinary speakers im-

³⁹As I will use 'relativism', I intend to denote a class of views that it would perhaps, in the present philosophical climate, be more appropriate to call 'contextualism' or 'indexicalism'. See, e.g., Cappelen and Hawthorne, 2009: 19-20. Thus my sentences that involve the phrase 'true relative to' should be taken as a sort of loose talk.

implicitly think they are describing an objective, mind-independent moral reality, it would be best to interpret our ordinary moral discourse differently. It is reasonable to interpret moral judgments as involving this sort of relativization, and when we do so, we are able to make many of them come out true. I believe that Mackie held such a view about morality.

1.5.1 What is Explicit in Mackie's Writings?

Recall the three main elements of Mackie's position: the metaphysical thesis, the discourse thesis, and the error thesis. Mackie sums up these elements of his view in the following passage, which I have already presented above. Still, it will be worthwhile to present it again:

...the denial of objective values will have to be put forward ... as an 'error theory', a theory that although most people in making moral judgements *implicitly claim*, among other things, to be pointing to something objectively prescriptive, *these claims* are all false. It is this that makes the name 'moral skepticism' appropriate. (35; italics my own)

Almost the same summary occurs a few pages later in *EIRW* at the close of the first chapter. He writes,

Moral skepticism must, therefore, take the form of an error theory, admitting that *a belief in objective values* is built into ordinary moral thought and language, but holding that *this ingrained belief* is false. (48-9; italics my own)

These are Mackie's most explicit full statements of his view. It is important to notice, however, that neither of them includes a claim about the truth-values of the moral judgments themselves. Consider the emphasized parts of the former passage. On a natural reading of this passage, 'these' is anaphoric on 'implicitly claim': an implicit claim is not the moral judgment itself. This is even clearer in the later passage. The ingrained belief, whose truth he denies, is not identical with any of the ordinary moral judgments of the average man on the street.

So Mackie's explicit formulations of his error theory lack any thesis about the truth-values of moral judgments. This extra thesis, however, is not missing from contemporary

explications of the error theory.⁴⁰ Perhaps this is because it seems like such an obvious consequence of what Mackie wrote in the earlier parts of *EIRW*. As a result, commentators on Mackie's error theory seem just to assume that Mackie held this extra thesis.

Mackie, however, is not so sanguine about this alleged result of his view. He writes, in the last paragraph of the first chapter of *EIRW*,

But what if we can establish this negative conclusion, that there are no objective values? How does it help us to say anything positively about ethics? Does it not at one stroke rule out all normative ethics, laying it down that all affirmative first order judgements are false, since they include, by virtue of the very meanings of their terms, unwarranted claims to objectivity? I will take up these questions in Chapter 5... (49)

So Mackie does not immediately suppose that our moral judgments are false as a result of his metaphysical thesis and discourse thesis. When we flip ahead to the fifth chapter of *EIRW*, expecting a clear and explicit statement of his view about whether moral judgments are true, though, we are left wondering what he really thought. About all he says on the matter is this:

I have argued in Chapter 1 that there are no objective values, and in Chapters 2, 3, and 4 that no substantive moral conclusions or serious constraints on moral views can be derived from either the meanings of moral terms or the logic of moral discourse. (105)

He never says, there or elsewhere in *EIRW*, whether moral judgments are generally true, false, or otherwise. At least, I can find no claim to that effect. Neither can I find any citation in the literature that references a part of *EIRW* where he clearly states such a view: every relevant citation that I know of references one of the passages I have already cited in this chapter.

This absence, I claim, is evidence that he thought his view allowed for the truth of positive first order moral judgments. I grant that it is not yet very strong evidence. But

⁴⁰See the explications of moral error theories in Huemer, 2005; Joyce, 2001; and Shafer-Landau, 2003; as well as just about any other explication you might find in the literature.

the systematic untruth of moral judgments would be an important and surprising result. At least, it would be important enough that we should expect Mackie to come out and say that he believed it, if he did. Since he never explicitly wrote that moral judgments systematically fail to secure truth, we have reason to reconsider the ubiquitous claim that Mackie accepted this thesis.

In the next section, I will draw an extended analogy between Mackie's position and certain versions of relativism. This is meant to further support my contention that Mackie's view is roughly compatible with relativism.

1.5.2 Relativism about Mass and Morals

The following line of thought results in a relativist view about mass.⁴¹ On the resulting view, there are no intrinsic, non-relational, properties that it would make sense to call 'mass properties'. Instead, there are only relational mass properties, and so it is these that should figure in the truth-conditions of our mass judgments.

Contemporary scientific theories seem to indicate that there is not really any such property as mass *in rerum natura*. At least, if a mass judgment is to have any prospect of being true, then it can only be true relative to some frame of reference or other. Another way of putting this is that anything in the fabric of the world that somewhat satisfies our concept of mass is a relational property. There are no intrinsic mass properties.

Suppose that a farmer, while out in his hay field, says, 'My little red tractor is a thousand kilograms.' He seems to be expressing the idea that his tractor has a certain property, viz., being one thousand kilograms.

It seems that, when our farmer makes this judgment, the thought he expresses does not involve any relativization. His concept of mass is not of a feature that is relational to frames of reference. It would be better to describe his concept as an idea of something

⁴¹Note that, in this subsection, I do not intend to commit myself to any of the theses of the relativist about mass. The sentence to which this footnote is attached should be thought of as a preface that revokes any commitment to such views that I might have otherwise incurred.

that is intrinsic and non-relational. Unless he's informed of some internal change in the composition of the tractor, or he finds out that he's wrong about what it is made of, he will stick to his guns about its mass. If we removed the tractor's engine, he would agree that its mass changed. If we replaced the hard metal seat with a bright red cushioned one, he might say that it lost a little mass. But he would be just as inclined to judge that his tractor is a thousand kilograms, even if it were placed on a high-speed train, or in a near warp-speed spaceship. Insofar as the farmer tacitly accepts this, he presupposes that intrinsic non-relational mass properties are part of the fabric of the world. Most people probably have a similar "tacit theory" about mass.⁴²

So far, this line of thought is in exact parallel with Mackie's explicit statement of his "error theory". We could even paraphrase the view as follows, mimicking Mackie's formulation of his own view: "*Mass* skepticism must take the form of an error theory, admitting that a belief in *intrinsic, non-relational masses* is built into ordinary *mass* thought and language, but holding that this ingrained belief is false."

Many philosophers, when confronted with this triad of claims, are loathe to conclude that our mass judgments are systematically false.⁴³ This leads them to suggest semantical accounts of our mass discourse that allow mass judgments to be true, even if the judgments must be, in some way, relativized. Thus, on this relativist line of thought, although there's nothing that exactly corresponds to the ordinary concept of mass, there are the materials, *in rerum natura*, with which we can construct good (enough) substitutes. The good substitutes are those relational mass properties. On these accounts of our mass discourse, the relational mass properties take the place of the nonexistent intrinsic ones in the truth-conditions of our mass judgments.

⁴²That is not to say that anyone's theory is developed; or that the average man believes in an abstract plenitudinous realm of properties; or that he has a full understanding of the notions of intrinsicity and relationality; or that his tacit theory is fully coherent. See, e.g., Dreier, 2005: 261-2.

⁴³See Dreier, 2005; Harman, 1996; Lewis, 1989. But compare Boghossian, 2006 and 2011.

What makes them good substitutes? One thing is that ordinary speakers never explicitly said they weren't talking about relational mass properties. The ordinary man in the street is unaware of any metaphysical implications of contemporary scientific theories, and is equally unaware of the metaphysical implications of his inchoate tacit theory of mass.⁴⁴ Moreover, as it happens, our ordinary mass judgments track sufficiently well these relational mass properties. They are eligible, simply in virtue of the way we use 'mass' and related terms, to serve as the contents of our ascriptions of mass in ordinary discourse.

However, according to the relativist, none of these substitutes is, strictly speaking, fit to play the role of the meaning of 'mass' (or the concept of mass), since the ordinary concept of mass is an idea of a non-relational feature.⁴⁵ Instead, they are merely fit to be the contents of our mass predications. As Harman puts it,

Einstein's Theory of Relativity does not involve a claim about meaning or about what people intend to be claiming when they make judgments about an object's mass. The point is, rather, that the only truth there is in this area is relative truth. (1996: 4)

So, according to this line of thought, although the meaning of 'mass' may be infected with an inaccurate conception of the world, for the purposes of assigning truth-conditions, our mass judgments are to be taken as relativized. An important aspect of this final step is that, although our ordinary concept of mass helps to determine certain aspects of the "meanings" of our mass claims, it doesn't fix their truth-conditions. So there is a disconnect between, on the one hand, our ordinary concept of mass and the conventional meaning of 'mass'; and on the other hand, the contents of mass judgments and their truth-conditions.

⁴⁴Thus, according to this line of thought, these properties might not be admissible substitutes for someone who explicitly accepts a non-relativistic ontology of the natural world (say, a present day advocate of Newton), and who invokes that ontology in his ordinary utterances about mass.

⁴⁵See Harman and Thomson, 1996: 4-5; compare Boghossian's (hostile) formulation of relativism in his 2006: 16-7; and 2011: 55.

Some relativists about morality have argued for their brands of relativism by analogy to this line of thought.⁴⁶ As Dreier puts it,

For pre-theoretic moral ideas to be all they aspire to be, ... there would have to [be] absolute standards for moral concepts to latch onto. But since there aren't any, relativism suggests, why not make do with the relative standards that we actually do have? There is no need to abandon moral judgment altogether, so long as we are willing to tone down its aspirations. (2005: 261)

This conception of moral relativism may thus be seen as subscribing to a line of thought that consists, roughly, in steps I have described: there are no objective values; but people implicitly presuppose that there are such values when they make moral judgments; there are plenty of admissible substitutes; so, we can take these substitutes to be fed in to the truth-conditions of our moral judgments. As a result, we may allow many of our moral judgments to come out true.

1.5.3 Mackie's Construction of Right and Wrong

Mackie, I have noted, accepts a principle to the effect that evaluative terminology is indexical. Roughly, according to the thesis of evaluative indexicalism, in order to find out whether an evaluative judgment is true, we must first figure out what standard of evaluation is salient in the context.

Thus, in order to determine what Mackie would take to be the truth-values of moral judgments, we must determine what standards of evaluation are salient in moral contexts (i.e., contexts in which moral judgments are made). It is difficult to say exactly what is Mackie's view about this issue. In the second chapter of *EIRW*, he writes,

Someone who uses the concept of objective moral value will suppose that there are requirements which are simply there, in the nature of things, without

⁴⁶Harman's brand of moral relativism may not quite fit the mold that I have created. It seems that he might not accept the third step, according to which people make the presupposition that morality is objective. It is, however, unclear to me, since he *does* try to motivate his moral relativism by appeal to the analogy to relativism about mass and simultaneity. Compare Harman, 1996 with his 1978 and 1975; it may be that his views about relativism have warped over the years.

being the requirements of any person or body of persons, even God. To be morally good will then be such as to satisfy these intrinsic requirements. (59)

In this passage, he seems to be suggesting that, in any moral context, the salient standard is determined by “the intrinsic requirements”. No doubt Mackie thinks it would be a mistake to express a judgment that could be true only relative to such standards. Elsewhere in *EIRW*, however, Mackie suggests that the salient standard is not determined by “the intrinsic requirements”. He writes,

...‘ought’ seldom, if ever, in ordinary use, refers to such supposed intrinsic requirements alone; it typically refers also to reasons or requirements of at least one of the other sorts, the intrinsic requirements being seen as backing them up. (76)

It seems to me that, in these passages, Mackie is acknowledging the possibility that we may take the truth-conditions of our moral judgments to be determined partially by whatever standards “the intrinsic requirements” are supposed to “back up”. We should, therefore, recognize that Mackie thought there are perfectly good, non-defective, subject matters for our moral judgments, which are fixed in context by these standards of evaluation that need to be “backed up”. It seems that Mackie would allow that these non-defective subject matters are comprised of properties that may be substituted into the truth-conditions of moral judgments. In what follows, I’ll call these alleged substitute properties ‘subjective values’, and their standards of evaluation I’ll call ‘subjective standards of evaluation’.

I believe that, on Mackie’s view, we may take the truth-conditions of our moral judgments to be determined by these subjective standards of evaluation, even though the pre-suppositions of ordinary speakers seem to dictate otherwise. Thus, although our ordinary moral concepts, and the conventional meanings of our moral terms, involve some defect, the truth-conditions of our moral judgments need not involve this defect. When writing

out the truth-conditions of moral judgments, we need merely make reference to subjective values and their standards of evaluation.⁴⁷

On this interpretation of Mackie, then, many of our moral judgments may come out as true. The truth-conditions of moral judgments need not mention any philosophically problematic objective standards of evaluation, for, in any moral context, there are salient subjective standards. These subjective standards will be contained in the truth-conditions of our moral judgments. And if the world is the way the truth-conditions demand, then our moral judgments will come out true.

Aside from the textual evidence that I cited above, I have a further reason for interpreting Mackie in this way. If we don't understand him in the way I suggest, then there is a puzzle about the relation between the first and second parts of *EIRW*. The puzzle is this: if Mackie systematically rejected the truths of moral judgments, then there is no good way to interpret him without attributing to him blatantly inconsistent commitments.

Suppose that, in the first part of *EIRW*, Mackie's main conclusion is that moral judgments, by and large, fail to secure truth. Notice, now, that in the second part of *EIRW*, he makes a number of claims that imply the truth of moral judgments. For example, he writes, at the end of the eighth chapter,

This chapter, and indeed the whole of this part of the book, has done no more than sketch the outlines of a first order moral theory. ... No doubt my approach could be called, in a very broad sense, a rule utilitarian one ... but it would be utilitarianism without its characteristic fictions. (199-200)

We should here take Mackie as expressing his *belief* in some kind of utilitarian theory. On any version of utilitarianism, there will be principles that, roughly, connect facts about the happiness of people to moral facts about what we ought to do. Insofar as he believes this sort of utilitarianism, and insofar as he accepts that utilitarianism implies the truth of some moral judgments, he is blatantly committed to the truth of moral judgments. Therefore, if

⁴⁷See the next chapter, at the end of section 3 where I attempt to state this interpretation of Mackie's view in a bit more detail.

Mackie rejected moral judgments as untrue in the first part of *EIRW*, then he is committed to the truth of a contradiction.

I cannot, in good faith, interpret Mackie in this way. He did not publish a book wherein the main message of the first part was a certain proposition (that all moral judgments are untrue), and a main message of the second part was its denial (that some moral judgments are true). Even if there are some inconsistencies around the edges of Mackie's work, such a blatant inconsistency should not be attributed to him. In any case, I strongly prefer not to do it.

Someone might respond that it is possible to interpret Mackie as expressing "something less than belief" in utilitarianism.⁴⁸ On such a view, we may say that he "merely accepts", or "pretends to believe", or "makes as if to believe", the utilitarianism that he outlines in the second part of *EIRW*. One might cite, as evidence for this claim, Mackie's famous quote, "Morality is not to be discovered, but to be made: we have to decide what moral views to adopt, what moral stands to take."⁴⁹ There is a natural interpretation of 'decide what views to adopt' that takes 'adopt' as expressing an attitude short of belief.

On such an interpretation of Mackie, however, we are unable to explain why he takes it to be so important to his positive moral system that it contain no reference to the "characteristic fictions" of utilitarianism. He writes,

My hope is that concrete moral issues can be argued out without appeal to any mythical objective values or requirements or obligations or transcendental necessities, but also without appeal to a fictitiously unitary and measurable happiness or to invalid arguments that attempt to establish the general happiness as a peculiarly authoritative end. (199)

If Mackie is merely pretending to believe the moral views that he adopts, then why does he take such great care to avoid attributing a mythical quality to them? Typically when we

⁴⁸This "fictionalist" alternative seems to be a somewhat standard interpretation of Mackie. See, e.g., Garner, 2007; Nolan, Restall, and West 2005; and Joyce, 2001.

⁴⁹1977:106.

pretend to believe something, we are not very worried about its truth. There is no good answer to this question, other than that Mackie is simply not pretending. If he were merely expressing a make-believe story about morality, then it is reasonable to think that he would have explicitly noted it, and explained himself. He didn't, so he's not pretending.

I would prefer to interpret the above famous quote as an expression the commonsensical view that, if we abolished moral discourse, society might not be able to function properly. So we have to decide what views to adopt: it's best if we keep on with moral discourse, but don't invoke objective standards of evaluation when we participate in the discourse, since that's somewhat misleading. In this way, we can say lots of true and useful things, but we avoid presupposing the existence of anything metaphysically problematic. That's all Mackie means to suggest when he says we must carry on with morality.

1.5.4 Analogies and Disanalogies

There are some important respects of analogy between this interpretation of Mackie's view and the relativist view that I outlined in the preceding section. First, Mackie's view shares the ontology of relativism in a negative respect. Neither view, presumably, accepts the existence of objective values. But both views hold that people systematically presuppose that there are objective values, and that the notion of an objective value is somehow involved in the meanings of moral terms.

Second, Mackie's view shares the ontology of relativism in a positive respect. Both views accept the existence of "subjective values", values that are determined merely by the interests and desires of speakers, or some contextually salient audience. Mackie and the relativist both think there are plenty of materials in the world with which we can construct these subjective values (viz., the desires and interests of the people in question).

Third, Mackie, like the relativists, thinks these subjective values are "good enough" for the purposes of figuring in the content of moral judgments, even if they don't adequately capture the meanings of moral terms, or our ordinary moral concepts. In a typical moral

context, even if we are invoking objective evaluative standards, there are also contextually salient subjective standards. A corollary of this third point is that there may be a disconnect between the thoughts we express with moral judgments, on the one hand, and the truth-conditions of those judgments on the other. It's not the case that every aspect of the meaning of a moral term helps to determine its content, and so the thoughts expressed by moral claims don't determine their truth-conditions.

On this understanding of Mackie's view, the truth of a moral judgment depends on whether there is a certain kind of value, one that is salient in the context of the utterance, and which corresponds with the moral judgment in the right way. When I utter, 'pleasure is good', my judgment will be true if there is a contextually salient subjective value on which pleasure ranks highly. Perhaps the salient value is fixed by my own moral sensibility; perhaps it is composed of some collective sensibility that is shared in my society; perhaps it is determined in some other way. This is exactly analogous to the above relativism about mass. When a farmer utters, 'my little red tractor is a thousand kilograms', his judgment will be true provided that there is a salient frame of reference relative to which his tractor has a mass of a thousand kilograms. Perhaps the salient reference frame is one in which the tractor is at rest; perhaps it is one in which all the conversational participants are at rest; perhaps it is some other salient reference frame.

There is an important respect of disanalogy, however, between the sort of relativism Mackie may have accepted and the sort of relativism some want to use to interpret mass discourse. Plausibly, no philosopher would recommend of ordinary speakers that they stop using a concept of intrinsic mass, and instead start to think in terms of a relational notion of mass. It would simply be too unwieldy; everyone would first have to learn complex physics. But Mackie seems to suggest that we should stop using the concept of objective value, and replace it with certain subjective ideas of value. This respect of disanalogy, however, does not suggest that Mackie did not hold the relativist view that I have attributed

to him. It merely suggests that he believed that, since we already have these subjective values in mind, we may as well do away with the defective objective concepts of value.

1.6 Conclusion

Where are we now? In this chapter, my main goal has been to explicate the views of J.L. Mackie, who is allegedly the paradigm example of an error theorist. We have seen that, in many respects, it is very unclear what views he really held. This is, in part, because the technical terms he used – among them, ‘moral judgment’ and ‘objective value’ – are notoriously slippery. I have tried to clarify, to a certain extent, their meanings. I have not provided perfectly precise meanings for these terms. But I believe I have done enough in order to conduct a profitable discussion of how to formulate the error theory.

Another main goal was to situate the subject matter of this dissertation into the broader philosophical terrain. I have argued that Mackie’s position is a version of cognitivism and a version of anti-realism. Thus, on my interpretation of Mackie, moral thoughts are beliefs, but these beliefs often involve propositions with a defective subject matter. Finally, I have argued that Mackie can be construed as a sort of relativist about morality. On his view, our moral judgments may be true, so long as we have an appropriate subjective value in mind – and we typically do. We should eradicate the defective ideas of objective value from our conceptual schemes, since they don’t correspond to anything in reality.

My discussion of Mackie has had a historical flavor. I have not engaged thoroughly with his arguments, and I have tried not to voice many opinions about what is the determinately correct interpretation of his various views. But I do not take this to be a defect of my discussion. Indeed, I believe that it would be naive to present any opinions about a view that is as under-specified as Mackie’s, other than that it is under-specified, and here are the possible specifications. I use Mackie’s writings as a starting point from which a rigorous discussion of the error theory may begin.

In this vein, I continue the discussion in the next chapter, where I make precise a variety of ways of understanding Mackie's Thesis 2, the discourse thesis.

CHAPTER 2

ERROR THEORIES AND PRESUPPOSITION

2.1 Introduction

I take Theses 1-3 of Mackie's position to be partially definitive of moral error theories. Here they are:

Thesis 1: The Metaphysical Thesis. There is something seriously defective with moral reality, viz., it is supposed to be objective.

Thesis 2: The Discourse Thesis. All moral judgments involve a claim to the effect that morality is objective.

Thesis 3: The Error Thesis. Since objectivity is a defect in morality, and moral judgments involve a claim to objectivity, moral judgments are infected with widespread error.

In order to qualify as an error theory, a position must imply those claims. This is not to say that the conjunction of Theses 1-3 is a sufficient condition on being a form of moral error theory. We have seen that a relativist about morality may accept the central ideas behind Theses 1-3. Nor is this to say that any error theorist must hold Mackie's ideas about objectivity. Instead, we would do well to construe 'objective' as a sort of black-box, waiting to be filled in by a particular error theory. Objectivity, then, is whatever the error theorist finds to be defective about morality. In this chapter, I'll suppose that the defect is the alleged existence of objective value properties, conceived, very roughly, in the way Mackie conceives them.

The goal of this chapter is to make precise Thesis 2, the discourse thesis. There are roughly three crucial ways of interpreting it. On a pragmatic interpretation, it is about what people do when they make moral judgments (section 3). On a more semantic interpretation, it is about linguistic requirements for moral statements (sections 4 and 5). On yet another

understanding, it is about the constituents of our moral concepts (section 6). Contemporary explications of the discourse thesis tend to involve a generic notion of presupposition. There are, however, a variety of specific concepts of presupposition. I will show that, given this pluralism about presupposition, we get hugely diverse readings of the discourse thesis, which correspond to these three rough interpretations. The different readings of the discourse thesis also yield very different versions of Thesis 3, the error thesis.

In the meta-ethical literature, although it has been pointed out that a notion of presupposition may play a role in defining the error theory, the differences between conceptions of presupposition have not adequately been explored. Importantly, the differences between pragmatic and semantic notions of presupposition have not been noted in the literature on moral error theories.

2.2 The Discourse Thesis in the Literature

In this section, I explain why I take the idea of presupposition to be a central part of moral error theories. In short, there are roughly two reasons. First, some find the notion of presupposition to be implicit in Mackie's writings. Second, others claim, without regard to whether it is found in Mackie, that a proper understanding of the error theory must involve presupposition.

It will be helpful to have a rough generic idea of the notion of presupposition. Intuitively, a claim is presupposed when it is taken for granted. Sometimes, we can tell that a person presupposes a claim when it is obvious that believes it is uncontroversial. Sometimes a theory presupposes a claim when it requires the prior truth of the claim. Sometimes a judgment presupposes a claim when we must accept the claim, if for only a minute, in order properly to interpret the judgment. In each of these kinds of situation, claims will be taken for granted – whether they are being taken for granted by people, judgments, or theories. In some sense, then, the term 'presupposition' applies to each of these kinds of "taking for granted".

2.2.1 Interpretations of Mackie in terms of Presupposition

Here are the passages where the discourse thesis is most clearly stated in *EIRW*:

...ordinary moral judgements include a claim to objectivity, an assumption that there are objective values... (35)

...a belief in objective values is built into ordinary moral thought and language... (48-9)

These passages are vague. It is unclear what it means for a judgment to “include a claim”. Likewise, there are many ways for beliefs to be “built into” a discourse. On one way of paraphrasing these passages, we can take Mackie to be using some ordinary generic notion of presupposition. On this reading, for moral judgments to include a claim to objectivity would be for them to presuppose that there are objective values. They would take for granted the objectivity of morality.

This is a common understanding of what Mackie had in mind.¹ David Phillips, in a recent paper, clearly interprets Mackie in this vein. He writes,

Mackie’s error theory is generated by the combination of two claims: (1) that ordinary moral judgments presuppose the existence of objective values, and (2) that there are no objective values. (2010: 95)

Thus, contemporary explications of Mackie interpret him as making some claim about presupposition when he states the discourse thesis.² If so, then it is reasonable to expect precisifications of the discourse thesis to involve the idea of presupposition.

2.2.2 Other Philosophers and Presupposition

This tendency to explicate error theories in terms of presupposition is not particular to Mackie exegesis. It spreads to general discussions of error theories. In a recent debate about

¹In a textbook on meta-ethics, Russ Shafer-Landau writes, “[On one view,] there is a massive error of presupposition that underlies [moral] vocabulary. If that were so, then all claims within the domain would be false – they would assume the existence of a kind of reality that isn’t there. This is an error theory; John Mackie (1977) is the most famous proponent of this sort of view in ethics.” (2003: 19-20; cf. 2005: 108) Joyce (2011) writes, “...Mackie claims that morality presupposes ‘objective values’ and ‘objective prescriptions’...” (521). See also Smith, 1994:11, 64-6.

²Compare Hare, 1981 and 1993.

the plausibility of error theories, Richard Joyce and Stephen Finlay each explicate the error theory in terms of presupposition. Finlay writes that, according to an error theory,

...moral judgments involve a particular kind of presupposition which is essential to their status as moral... (2008: 347)

He labels this commitment of an error theory '*Presupposition*'. This thesis obviously bears certain similarities to the discourse thesis. Finlay goes on,

...I argue against *Presupposition*. My attack has two prongs. First, I argue defensively that an assumption of the absolute authority of moral value may not be characteristic of moral discourse at all.... Second, I argue that even if such an assumption is ubiquitous, we still have good reasons to conclude that it does not contaminate the meaning or truth-conditions of moral discourse... (2008: 348)

In a discussion of Finlay's paper, Joyce replies to Finlay's arguments. Joyce writes,

...moral discourse presupposes non-institutional desire transcendent reasons and non-institutional categorical imperatives.... I am willing to attribute to ordinary thought a presupposition about non-institutional desire transcendent reasons and non-institutional categorical imperatives (though of course I don't claim that ordinary speakers think in these terms!). (2011: 323, 325)

The focus of this debate between Joyce and Finlay is about whether the mentioned presupposition really is part of moral discourse. Joyce claims that it is; Finlay claims that it isn't.³

Thus, it is a live debate in contemporary discussions whether there is a defective presupposition embedded in moral discourse. So the notion of presupposition is not only an artifact of Mackie's position.

2.2.3 Some Potential Caveats

I would like to avoid some confusion by mentioning some caveats. Not all interpreters of Mackie understand him as Phillips does. Some do not use the term 'presupposition' in their explications of Mackie. For example, Michael Smith writes,

³See also Sinnott-Armstrong, 2006: Chapter 3; Shafer Landau, 2005; Nolan, Restall, and West, 2005; and Garner, 2010.

...Mackie's argument from queerness consists in an analysis of the concept of moral value together with a pair of challenges which purport to show that we can make no real sense of how that concept could be instantiated. (2010: 121)

Nowhere in the article of this passage does Smith explicate the discourse thesis as involving any kind of presupposition.⁴ Instead, Smith's reading of the discourse thesis takes it as a claim about our moral concepts.⁵

Some philosophers are implicitly slippery in their explications of the discourse thesis. For example, Joyce writes,

...the error theorist focuses on a proper subset of sentences containing the problematic terms: those that imply or presuppose the instantiation of a moral property... (2007a)

...for something to be morally bad would imply or presuppose a kind of inescapable, authoritative imperative against pursuing that thing... (2007b)

In these passages, Joyce seems to suggest that we could read the discourse thesis in a similar way to Smith's above interpretation. Insofar as he accepts this alternative reading, his understanding of the discourse thesis is, to a certain extent, equivocal. In the above mentioned debate with Finlay, Joyce interprets the discourse thesis in terms of presupposition. Here, he possibly interprets it in another way, in terms of a logical notion of implication.

Thus, it is hard to pin down exactly how contemporary authors understand the discourse thesis. Sometimes, it is to be understood in terms of notions of presupposition; other times, in terms of logical implication; other times, in terms of conceptual analysis. There is no guarantee that these understandings collapse into each other.

Some philosophers are even explicitly vague on how they understand the discourse thesis. In a recent article, Simon Kirchin writes,

A moral error theory...consists of two main elements: (i) a characterization of everyday moral thought and language that picks out some commitment that is embodied by such thought and language (or that is based on such thought and

⁴See, however, Smith, 1994: 11.

⁵Cf. Sinnott-Armstrong, 2010: 67-9 and 2006: 32; Joyce, 2001: 16.

language, or is directly implied by it, or similar), and which, in turn, is claimed to be crucial (or central, or key, or fundamental, or similar); and (ii) some argument (or arguments) that purports to show that this crucial commitment is false (or fundamentally misguided, or completely unjustified, or wildly crazy, or similar) such that the whole of our everyday moral thought and language is flawed, not just the commitment in question. (2010: 168)

In a footnote to this passage, he writes,

I use 'commitment' throughout as a neutral catch-all. 'Assumption,' 'claim,' 'idea,' 'intuition,' and the like will work equally well. (2010: 168)

Kirchin seems to be endorsing a sentiment that it simply *doesn't matter* how we precisely understand the discourse thesis, and that the imprecise characterizations are good enough to conduct a proper discussion. It seems that this sentiment is contained in much of the literature on moral error theories.

I reject this sentiment. Precision matters a lot. Radically different views result from the different ways of making precise the discourse thesis. These views are so diverse that it would be wishful thinking to expect any general argument for or against Theses 1-3 to lead to anything more than confusion. In my opinion, this has resulted in a lot of confusion underlying many contemporary discussions of the error theory. We would do well to wash some of the confusion away with clear and precise understandings of the discourse thesis and the error thesis.

2.3 Pragmatic Presuppositions

In this section, I explain the notion of speaker presupposition, through Robert Stalnaker's well known account of the subject. With the concept of speaker presupposition in hand, I state a version of the discourse thesis, as well as a version of the error thesis. These renditions of Theses 2-3 embody a potential take on the moral error theory. Finally, I argue that the approach cannot be made to work.

2.3.1 Stalnaker on Pragmatic Speaker Presupposition

I begin by outlining the general role that Stalnaker takes presupposition to play in our discourse. This is necessary for a proper understanding of Stalnaker's notion of presupposition.

In virtue of his cognitivism, it is natural for an error theorist to hold that when we participate in moral discourse – when we take part in moral conversations – we communicate with others. That is, when you make a moral judgment, you impart information to an audience. Just as you might utter, 'The bathroom is down the hall and to the left' to tell someone about the world, you might say, 'What those little boys did was wrong' to tell someone what the world is like.⁶ An error theorist, it seems, should view moral discourse as roughly similar in nature to any other kind of giving and taking of information.

According to a standard view, communication is a kind of rational cooperative activity.⁷ On such a view, there are rules for appropriate communication, which anyone participating in a conversation tacitly knows. As long as the participants are properly cooperative, they desire to follow the rules, and so they have reason not to break them. These rules of communication tell people what they can and cannot say, given what has already been said in the conversation. So the rules of conversation require the existence of a conversational record, which keeps track of what information has already passed between the participants.

On Stalnaker's view, a conversational record determines what he calls a 'discourse context': a set of possible worlds.⁸ Intuitively, the discourse context tells the participants which propositions are entailed by what has already been said in the conversation, which propositions have been ruled out, and which propositions are still live options. The propositions that are true in all possibilities in the discourse context are to be taken for granted by the

⁶Cf. Mackie, 1977 and 1946; Joyce, 2001: 9-16.

⁷Stalnaker, 1970 and 1998; Lewis, 1973 and 1979; Grice, 1975.

⁸The term 'discourse context' comes from Stalnaker, 1998. He also calls discourse contexts 'context sets' in his 1978.

participants for the rest of the conversation. Propositions that are false in all these possibilities should be ruled out for the remainder of the conversation. Propositions that are true in some but not all worlds in the discourse context are still live options: they are “fair game”, “up for grabs”, ready to be asserted or denied.⁹ The goal of communication, then, is to try to determine, in the most economical way, which of these possible worlds is the actual world. When we exchange information in a conversation, we try to figure out what is actual.

Our principal tool for communication is the speech act of assertion. Since an assertion takes place within a discourse context, its content interacts with its context. Two superficially indiscernible assertive utterances may have different content depending upon the contexts in which they take place. Determinants of the content of an assertive utterance include what has been taken for granted from earlier in the conversation, and what must be taken for granted in order properly to understand the assertion. Thus, according to an error theorist who accepts this picture of communication, and who accepts that we communicate *via* moral discourse, what we end up asserting with our moral judgments depends upon what we take for granted in making the judgments. On Stalnaker’s view, these propositions that we “take for granted” are our presuppositions.

According to this view, presupposition is a pragmatic phenomenon; it is something people do when they assertively utter moral sentences. Roughly, a person *pragmatically presupposes* a proposition at a time if she is disposed, as of that time, to act as if she accepts the proposition, and she is also disposed at that time to act as if her audience also accepts the proposition. Another way of putting this is that, when a speaker presupposes something, she takes it as common ground in her conversation. Pragmatic presupposition, thus conceived, is primarily a relation between speakers and propositions. In addition, it is not necessarily realized as a mental attitude in the speaker, since it is fundamentally a complex disposi-

⁹Stalnaker, 1978: 84-5.

tion. Our dispositions are not wholly determined by our internal mental states.¹⁰ Finally, it is important to note that derivative notions of presupposition may be defined in terms of this pragmatic conception of presupposition. On one possible definition, a judgment may be said to presuppose a proposition when the person making the judgment pragmatically presupposes it.¹¹

2.3.2 The Discourse Thesis in Terms of Pragmatic Presupposition

On the imprecise version of the discourse thesis, each moral judgment involves a claim in the objectivity of morality. We are now in a position to inject some precision into the discourse thesis by making precise the nature of the “claim” it alleges of moral judgments.

Thesis 2P: The Pragmatic Discourse Thesis. Speakers, when they make moral judgments, pragmatically presuppose that there are objective values.

The discourse thesis, thus stated, is still vague, in virtue of the proposition it says is pragmatically presupposed by speakers. At the very least, let us say that the pragmatic discourse thesis has it that when speakers make moral judgments, they take for granted the existence of a certain kind of property.

Here is an illustration of how the proponent of Thesis 2P would view moral communication. Suppose that I walk around the corner and find some young boys who are torturing a cat. I, hero that I am, break up the gathering and let the cat free. Later that day, I describe the facts of the situation to you in some detail. “They were surrounding the cat; kicking it; preparing to hang it from a tree.” Then I say,

1. What those little boys did was wrong.

¹⁰Stalnaker, 1970: 38; 1974: 52. See also Stalnaker, 2002, for a more recent definition that diverges from the idea that presuppositions are dispositional. Compare also 1987: 68-71.

¹¹There are obviously many other interesting presupposition relations that we may define partially in terms of pragmatic presupposition. Here are some: a sentence (type), ‘S’ presupposes *P* when people must pragmatically presuppose *P* in uttering ‘S’ in order to be rational; or when we must assume that the speaker pragmatically presupposes *P* in order to understand ‘S’; or when we can reasonably infer that a speaker pragmatically presupposes *P* on any utterance of ‘S’. See Stalnaker, 1974: 50.

According to Thesis 2P, when I utter (1), I pragmatically presuppose that there are objective values. Presumably, according to this view, the reason why I am presupposing the existence of objective values is that I tacitly think that, in order for ‘wrong’ to apply to what the boys did, such values must exist.¹²

Of course, the discourse thesis is meant not only to range over communicative utterances. It must apply generally to moral thought and language, even in situations that do not involve the public exchange of information. If, before I have my conversation with you, I think to myself that what the little boys did was wrong, then I presuppose that there are objective values, according to Thesis 2P. It is, presumably, supposed to apply both to “public” and “private” moral judgments.

This interpretation of the discourse thesis leads naturally to a version of the error thesis. According to the imprecise version of the error thesis, since morality is not objective, but our moral judgments claim that it is, our moral judgments admit of some kind of error. On most understandings of the error thesis, the mistake in moral judgments is that they fail to secure truth. So here is a precise statement of the error thesis in terms of pragmatic presupposition.

Thesis 3P: The Pragmatic Error Thesis. Since there are no objective values, but speakers pragmatically presuppose that there are when they make moral judgments, none of these judgments is true.

Presumably, the idea underlying Thesis 3P is that when a person’s presupposition fails, she can’t, in the end, be expressing something about actuality. Since the presupposition is inconsistent with what is actually the case, the resulting discourse context is defective. If the actual world is ruled out in the context by presupposing something that is inconsistent with

¹²The plausibility of this claim may legitimately be called into question: do people really act in the way that Thesis 2P says they do? At first glance, you might think that ordinary people have no beliefs about the notion of an objective value. I grant that this idea has some pull, though I think that Thesis 2P is not indefensible. I will not, however, discuss in depth the question of how to determine whether Thesis 2P holds in general, nor whether Thesis 2P really does hold in general. Error theories that consist partially in Thesis 2P will not be my main concentration in this work.

it, then it is hard to see how the judgment could be about the actual world. Statements that are not about what is actually the case, it seems, cannot be true: if you assert the proposition that x is F , but x is not actually F , then your assertion isn't true. On this view, presupposition failure is, to a certain extent, catastrophic.¹³

Here is an illustration of how the proponent of Thesis 3P would think of moral discourse. Suppose I am gossiping with you about funny European royals. First we talk about the Queen of England; then the Grand Duke of Luxembourg. Then you ask, "What do you know about the King of France?" I say, excitedly,

2. The King of France is bald!

When I utter (2), I take for granted that there is a King of France: I accept that there is one, so do you, and each of us knows this of the other. We take it to be common ground in our conversation. But since there is no such king, my presupposition is false. It allegedly follows that my utterance cannot be about how the world actually is, and so it isn't true.

A proponent of the error theoretic view embodied in Theses 2P-3P, it seems, would think of (1) along the same lines as (2). Like my utterance of (2), my utterance of (1) involves presupposition failure, resulting in untruth.¹⁴ Since we tend to think that what we say is true, this would be a big error on our part in participating in moral discourse.

2.3.3 Against Theses 2P-3P

I believe that the error theory that we get by combining Thesis 2P with Thesis 3P yields an incoherent conception of the relation between presupposition failure and the truth-values of our judgments.

¹³This idea is expressed in Yablo, 2006: section 2. I do not use Yablo's sense of 'catastrophic, though. See Yablo, 2006: section 9. It is important to note that Yablo doesn't, as far as I can tell, really accept the idea expressed in this paragraph. Neither do I (see below).

¹⁴This seems roughly to be the error theory had in mind by Huemer, 2005: 4-5; and Shafer-Landau, 2003: 19-20 and 2005: 108.

The problem is that, contrary to what the pragmatic error thesis implies, presupposition failure need not result in untruth.¹⁵ It is easy to imagine counterexamples to the general idea behind Thesis 3P. Generalizing Thesis 3P, we get the following principle:

Generalized Pragmatic Error Thesis. For any proposition, *P*, and sentence, ‘*S*’, if a speaker presupposes *P* in judging that *S*, and *P* is not true, then her judgment that *S* is not true.

This generalized pragmatic error thesis, however, is false. There are many ordinary, everyday examples where participants in a conversation satisfy the antecedent of this principle, but make the consequent false.

Suppose, for example, that we are discussing presidential wardrobes. You claim that Bush was a sloppy dresser. I say that I am inclined to agree. Then you attempt to generalize this to all Republican presidents. I, skeptical of this, ask, “What about McCain?” You, who are not up on the times, think McCain is the current president, instead of Obama. You utter,

3. The president wears a suit and tie to work.

And you back off from your initial generalization. You revise to the claim that *most* Republican presidents aren’t well-dressed. In this situation, there are two plausible ways I could interpret your utterance of (3), and neither of them results in untruth. Even so, in this context, you falsely presuppose that McCain is the current president.

First, I could understand you to be talking about McCain. If I interpreted you in this way, I would be presupposing along with you that McCain is the president, perhaps to avoid diverting the discussion. Still, what you assert by uttering (3) is true: McCain does wear a suit and tie to work. Second, I could understand you to be talking about Obama, even though you didn’t intend to be talking about him. If I interpreted you in this way, I would be rejecting your presupposition. Still, I would be interpreting your utterance of (3) as true.

¹⁵See, e.g., Stalnaker, 1998: 100; cf. Yablo, 2006.

Thus, on either plausible interpretation of your utterance, I should take your judgment as true, contrary to what the generalized pragmatic error thesis implies.¹⁶

This is not, in any way, a “special case” counterexample to the principle. The counterexamples to the generalized pragmatic error thesis are so widespread that no amount of chisholming can retain a plausible understanding of the initial idea behind it.¹⁷ Taking a step back, one may wonder what is wrong with the initial motivation for Thesis 3P, if it is such an implausible claim.

Recall the rationale that I presented for the pragmatic error thesis: when someone presupposes what is actually false, it is hard to see how her judgment could be about actuality; but if it is not about actuality, then it must not be true. The problem with this idea, in my opinion, is that it assumes the existence of a strong semantic connection between what is presupposed in making a judgment and what the judgment is about. This is a mistake: a person’s judgment may be about actuality, even if the person making the judgment has a false presupposition. In the next section, I discuss a notion of presupposition that has more of a semantic underpinning, and thus may yield a more plausible error thesis.

Before I move on, however, I would like briefly to mention a possible revision of Thesis 3P, which fits better with Thesis 2P. According to this revision, the error that people commit when they make moral judgments is not as severe as Thesis 3P would have it.

Thesis 3P: The Pragmatic Error Thesis**. Since there are no objective values, but speakers pragmatically presuppose that there are when they make moral judgments, they make the mistake of presupposing a falsehood.

¹⁶There is another way of interpreting what you say, which may lead to falsity: I might understand you as asserting the proposition expressed by the sentence, “McCain, who is the president, wears a suit and tie to work.” If you assert this proposition, then I should interpret your assertion as untrue. Two things about this: (i) in the context as I have described it, I think it would be uncharitable to interpret you as asserting that McCain is president, instead of merely presupposing it, and (ii) even if I interpreted your assertion in this way, the semantic value of your utterance should be interpreted as true.

¹⁷According to the *Philosophical Lexicon*, to *chisholm* is to make repeated small alterations in a definition or example. I think the lexicographers who wrote this entry intended to allow ‘chisholm’ to apply to this sort of procedure with regard to principles as well.

This revision does not claim that presupposition failure results in untruth. But it implies that it is often a mistake to presuppose something that is false. If we do it too often, our discussion can go off the rails in various ways. Suppose, for example, that in the above conversation about presidential wardrobes, we change the focus of our conversation to presidential policies. If we did this, the initial mistake of presupposing that McCain is president could compound into further mistakes down the road (supposing that McCain's policies would have diverged enough from Obama's).

The reason I mention Thesis 3P* derives from my discussion of Mackie's position in the preceding chapter. If I am correct in thinking that Mackie did not necessarily think moral judgments are false, then it is plausible that he accepted this revision instead of Thesis 3P. He seemed to think that, even though our false presupposition that there are objective values doesn't affect the truth-values of our moral judgments, we nonetheless should stop presupposing this. For, perhaps, this has led many moral theorists to big mistakes down the road in their theorizing. In Mackie's opinion, it is a mistake for moral theorists to take seriously the ordinary presupposition that there are objective values. This, I believe, is a very close approximation of what Mackie really had in mind. Of course, I put it forward only as a close approximation. I have already expressed my belief that it would be naive to come down in any perfectly determinate way about what position is expressed in Mackie's writings.¹⁸

¹⁸I believe that Mackie may even have thought it was a *conventional implicature* of moral judgments that there are objective values. This is suggested in a passage on page 35 of *EIRW* where he writes, "And I do not think it would be going too far to say that this assumption [in objective values] has been incorporated in the basic, conventional, meanings of moral terms." But conventional implicature does not (as it seems standardly to be conceived) affect truth-conditions. So understood, there is another version of the discourse thesis which we may interpret Mackie as accepting: that moral judgments conventionally implicate that there are objective values. Cf. Grice, 1975; Potts, 2007; and Stalnaker, 1998: 61 on conventional implicature.

2.4 Semantic Presuppositions

In this section, I explicate two semantic concepts of presupposition. Roughly, a judgment semantically presupposes a proposition when the proposition must be true in order for the judgment to have a truth-value. With the new concepts, I state two new versions of the discourse thesis, and two new versions of the error thesis. In the next section, I raise some worries about the plausibility of the resulting takes on the error theory.

2.4.1 Presupposition: Semantic or Pragmatic?

Recall our rough definition of pragmatic speaker presupposition: a speaker presupposes a proposition when she is disposed to take it as common ground in her conversation. On this conception, presupposition is fundamentally a pragmatic phenomenon: it is something we do when we participate in conversations. As such, it is also fundamentally a relation between people and propositions. We saw that it is possible to define useful notions of presupposition in terms of speaker presupposition that would relate judgments or sentences with propositions. But, on Stalnaker's account, the notion of speaker presupposition is a sort of foundational concept, in terms of which we may understand many linguistic phenomena.¹⁹

Stalnaker's (1974, 1973, and 1970) account of presupposition was presented in reaction to a then-popular account of presupposition in terms of truth-value gaps. One of his goals was to show that it would be a mistake to understand all presupposition in terms of truth-value gaps. The goal was not, however, entirely to supplant those notions of presupposition. When appropriately understood, the fundamental concepts of Stalnaker's view and those of its predecessors may live together in peace.²⁰ An account of presupposition in terms of truth-values gaps defines a concept that may be useful in our theorizing about language and

¹⁹Cf. Stalnaker, 2002: 702-3, 713-4.

²⁰See Stalnaker, 2002; and 1974: 53, 61; 1978: 89-90.

thought. Thus, such a semantic notion of presupposition may be what the error theorist is looking for in explicating the discourse thesis and the error thesis.

On a rough definition of semantic presupposition, a judgment presupposes a proposition when the truth of the proposition is a requirement that must be met before the judgment can be said to have a truth-value.²¹ There are, roughly, two such semantic notions of presupposition, one corresponding to a theory of presupposition which seems to have been put forward by Frege, and the other corresponding to a notion allegedly noticed by Strawson. Following Soames (1999), I'll call these 'Fregean' and 'Strawsonian' presupposition.²²

A proposition is a *Fregean presupposition* for another proposition just in case, in order for the second proposition to be true or false, the first proposition must be true. This notion of semantic presupposition is fundamentally defined over propositions. But there is a derivative sense in which it can relate sentences with propositions. Thus, roughly, a sentence has a proposition as a Fregean presupposition when the content of the sentence cannot be evaluated for truth unless that proposition is true.

A proposition is a *Strawsonian presupposition* of a sentence utterance just in case, in order for the sentence to express a content, that proposition must be true.²³ This notion is fundamentally a relation between sentence utterances and propositions. There is no corre-

²¹There is another notion of semantic presupposition, which seems to be important in the literature. On this notion, presupposition is a relation between interpreted sentence types and propositions, where 'S' presupposes P when it is impossible for someone rationally to assertively utter 'S' without also presupposing P. This notion does not totally reduce to speaker presupposition however, because of the notion of rationality that is contained in the definition. According to Stalnaker, "In [these] cases, one may just have to write presupposition constraints into the dictionary entry for a particular word" (1974: 61). See also Yablo, 2006: fn. 3. I do not totally understand what Stalnaker and Yablo are getting at here. But I suppose that this notion of presupposition is sort of a semantic/pragmatic cross-breed, partially defined in terms of the pragmatic notion of presupposition, (possibly) partially defined in terms of the notion of conventional implicature, and partially defined in terms of some notion of rationality. [I am confused about this.]

²²See also Soames, 1989: section 2. Later in this chapter, I'll use the terms 'logical presupposition' and 'expressive presupposition' to refer, respectively, to Fregean and Strawsonian presupposition. I borrow this terminology from Soames, 1989.

²³It is controversial whether this was Strawson's view, and I make no claim as to historical accuracy. See Soames, 1989: subsection 2.2, where he argues that Strawson was more or less confused about what notion of presupposition he was trying to elucidate.

sponding derivative notion that can systematically relate propositions with other propositions, since, if a Strawsonian presupposition of a sentence is false, a semantic fault occurs before a proposition can be expressed by the sentence. It is incoherent to say that a proposition can have a false Strawsonian presupposition.

The reason why it is useful to call these “semantic” notions of presupposition is that whether a statement has a presupposition, of either type, has an effect on the truth-conditions of the statement. But there is a difference in which the two notions effect truth-conditions. If we think of the semantic evaluation of sentence as a sort of two-step process, then the distinction between Fregean and Strawsonian presupposition is vivid. Suppose you want to figure out the truth-value of a sentence token. First you should try to figure out what proposition is expressed by the sentence in its context: if there is no such proposition, you can go no further in retrieving a truth-value for the sentence. Next, you should evaluate that proposition in the actual world. The truth-value of the sentence is a reflection of the truth-value of the proposition: if the proposition has no truth-value, then neither does the sentence. Strawsonian presupposition failure corresponds to a hang-up at the first step of semantic evaluation. Fregean presupposition failure corresponds to a hang-up at the second step.

2.4.2 The Discourse Thesis in Terms of Semantic Presupposition

In this subsection, I outline two approaches to the error theory, each of which focuses on one of these notions of presupposition. The first uses the notion of Fregean presupposition.

Thesis 2F: The Fregean Discourse Thesis. A Fregean presupposition of each moral judgment is that there are objective values.

The Fregean discourse thesis provides us with the following picture of moral communication.

Suppose, again, that I see the little boys kicking a cat, and later, I tell you,

1. What those little boys did was wrong.

In such a context, we need not think that I, the speaker, have pragmatically presupposed that there are objective values. Thesis 2F entails no such thing. Instead, it suggests, first, that I have tried to communicate with you. I have asserted the proposition that what the little boys did was wrong. Second, it entails that if there are in fact no objective values, then I have told you something that is neither true nor false.

There is allegedly a crucial relation between semantic presuppositions and pragmatic speaker presuppositions. As Stalnaker would have it, part of the goal of communication is to make *determinate* assertions. I.e., if you fail to say something that is either actually true or actually false, you have violated some rule of what can appropriately be said in a conversation, unless you presuppose propositions that would make it have a truth-value.²⁴ Thus, when we utter a sentence with a semantic presupposition (of either the Fregean or Strawsonian variety), we make the mistake of violating a norm of conversation, if we fail to presuppose what is semantically presupposed by the sentence.

With this in mind, it seems that the proponent of Thesis 2F will want to accept a version of the error thesis along the following lines:

Thesis 3F: The Fregean Error Thesis. Since there are no objective values, but a Fregean presupposition of moral judgments is that there are, either (i) speakers make the error of violating a conversational rule in making moral judgments, or (ii) their moral judgments are neither true nor false, because they express gappy propositions.

The use of ‘or’ in Thesis 3F clearly must be understood as inclusive: whether or not speakers actually presuppose the existence of objective values, their moral judgments end up semantically expressing propositions that are truth-value-less. No matter what speakers presuppose when uttering moral sentences, their judgments still end up being gappy. Since speakers tend to think that they say things that are true or false, this would be a mistake on their part. And the mistakes may pile up, if speakers are not careful. A proponent of

²⁴See Stalnaker, 1978: 88-90. Compare Grice, 1975 on the “maxim of quality”. Soames discusses the plausibility of Stalnaker’s idea in Soames, 1989: 583-4. He ultimately rejects it.

this view may hold that if speakers fail to presuppose what they're supposed to, then the conversation may go off the rails in so many ways.²⁵

Consider, for example, one of the above cases of non-moral presupposition. Suppose that I tell you,

2. The King of France is bald!

But, in doing so, I do not presuppose that there is a King of France. And you, ever the good conversational participant, accommodate my assertion by presupposing that there is one.²⁶ Later on in the conversation, you ask me whether the King of France also has a harem. I, still not presupposing that there is a King of France, reply: "What are you talking about? There is no King of France anymore – his head is on a spike!" Something has clearly gone wrong here, and the culprit is my failure to presuppose what I ought to have presupposed.²⁷

The discourse thesis may also be interpreted in terms of Strawsonian presupposition.

Thesis 2S: The Strawsonian Discourse Thesis. A Strawsonian presupposition of each moral judgment is that there are objective values.

The Strawsonian version of the discourse thesis provides us with another picture of ordinary moral discourse.

Again suppose that I utter sentence (1) to you after explaining what the little boys did to the cat. Thesis 2S does not entail that I presuppose that there are objective values when I utter such a sentence. It does entail, jointly along with Stalnaker's account of conversational rules, that I must presuppose that there are objective values, if I am to be a good conversational participant. But does it plausibly imply that I have tried to assert a proposition? Presumably, I have uttered a sentence that I *believe* expresses information about the world. But it seems that the proponent of Thesis 2S must say that, even if I presuppose that

²⁵See Simons, 2003 for a discussion of this idea.

²⁶See Stalnaker, 2002; and Lewis, 1979 on presupposition accommodation.

²⁷See Stalnaker, 1978: 85 on defective contexts.

there are objective values, there is no information, semantically determined by the sentence I have uttered, that I have tried to pass on to you. My utterance expresses no proposition: presupposition failure, in this sense, really is catastrophic.

Thus, it seems that the proponent of the Strawsonian discourse thesis will want to accept the following version of the error thesis:

Thesis 3S: The Strawsonian Error Thesis. Since there are no objective values, but a Strawsonian presupposition of moral judgments is that there are, either (i) speakers make the error of violating a conversational rule in making moral judgments, or (ii) their moral judgments are neither true nor false, because they fail to express propositions.²⁸

According to the combination of Theses 2S-3S, there are no moral propositions out there to be believed, feared, regretted, anticipated, asserted, and so on. This view therefore entails that when a speaker makes a moral judgment, he commits an incredible error. Not only does a moral judge typically fail to act in accordance with conversational rules; not only does he mistakenly think he is uttering a moral truth (or at least something with a truth-value). He makes the mistake of thinking his moral statements communicate any moral information *whatever*.

The mistake, of course, is not isolated to public contexts, where people seem to be communicating with each other about morality. Take an example from the first chapter of this work, where a man, while reading a newspaper article about a dog drowner, mutters to himself,

4. It's wrong for the murderer to kill all those dogs.

On the view embodied in Theses 2S-3S, this man not expressing his belief in a proposition expressed by (4). He does not believe that the murderer shouldn't have killed the dogs, for there is no such proposition. Presumably, however, he would be perfectly willing to attribute to himself belief in such a proposition. So in private contexts, people think they are doing things which in fact they are not, according to Thesis 3S.

²⁸Interestingly, Schiffer (1990: 605) attributes this view to Mackie.

2.5 Objections to the Fregean and Strawsonian Interpretations

In this section, I raise two worries for the Fregean and Strawsonian embodiments of the error theory. First, I argue that both views have a problem explaining what it is we are doing when we participate in moral discourse. Intuitively, we pass information with moral statements, but it isn't clear that these views can accommodate this intuition. Second, I argue that there is some *prima facie* reason to reject Thesis 2S. As a result, unless some good rationale for the Strawsonian position can be found, we should reject it.

2.5.1 A Worry for Theses 3F and 3S

An implausibility in the Strawsonian view is that it isn't obviously able to explain what we are doing when we participate in moral discourse. First I'll explain how the problem arises for Thesis 2S, and then I'll generalize it to Thesis 2F.

As I understand the Strawsonian view, it is committed to the idea that there is no such thing as moral belief and communication. Since belief can only be belief in a propositional sort of entity, and moral judgments do not express propositions, ordinary people have no moral beliefs. Likewise, since communication of information works *via* propositions, it seems that ordinary people cannot communicate when they make moral judgments. I take this to tell against the Strawsonian position. Allow me to explain.

When I utter (1) after giving you a description of the acts of the little boys, I communicate something to you. I take this to be a trivial fact about the situation. It is something that cannot coherently be denied. It is perfectly natural to report what I have done by uttering,

5. Casey said that what those little boys did was wrong.

Likewise, when a man utters (4) to himself after reading about the dog drowner, he expresses a belief. Then it seems true to say,

6. The man believes that it's wrong for the murderer to kill all those dogs.

Now, if the Strawsonian error thesis is true, then it seems that neither of these claims can be true. Since there is no proposition my utterance of (1) expressed, (5) cannot be true in

virtue of meaning that I bear the *said*-relation to a certain named proposition. Likewise, since (4) expresses no proposition in the man's mouth, (6) cannot truly say of him that he bears the *believed*-relation to such a proposition. This is unacceptable; it is here that I dig in my heels. I take it as a datum that we may truly utter (5) or (6) in order to report what has happened in these situations. Unless there is an alternative explanation available to the Strawsonian view, it has met its end.

I believe the Strawsonian may give an alternative explanation: he must adopt a different view about moral belief and communication. According to this revised view, roughly, the information reported about in (5) and (6) come apart from what is semantically expressed by their sentential complements. I think the Fregean view must adopt a similar explanation of moral belief and communication.

According to the Fregean error thesis, our moral judgments never actually succeed in securing a truth-value. They are uniformly lacking in truth and in falsity because they require the existence of something that in fact does not exist. As a matter of fact, arguments against the existence of objective values tend to be metaphysical arguments. Properly understood, they should yield the result that, as a matter of necessity, there are no objective values. Thus, it seems that any proponent of Thesis 3F believes that the proposition that there are objective values is an impossible proposition. This supposition, when combined with Thesis 3F, has an interesting result.

On the Stalnakerian apparatus that I have been using during my discussion of presupposition, propositions reduce to sets of possible worlds; or, equivalently, functions from possible worlds into truth-values. If we continue to use the apparatus of possible worlds semantics, then we have an easy way of representing propositions with Fregean presuppositions. If a proposition, Q , presupposes P , then the function from worlds to truth-values that Q reduces to is undefined on worlds where P is false. Intuitively, we may think of Q as being a long list consisting in two columns: on the left, a column with names for worlds; and on the right, a column with names for truth-values. If a given possible world is one

where P is true, then the column on the left has an entry for that world, which corresponds to an entry in the right column with either ‘True’ or ‘False’ (whichever value Q has at the world). If a given possible world is one where P is false, then the list for Q contains no entry in the left column for that world, and thus no corresponding entry in the right column.

What happens when a proposition has, as a Fregean presupposition, a necessary falsehood? The partial function to which a moral proposition reduces will be undefined for all worlds, since its Fregean presupposition is false at all worlds. The list that intuitively represents the partial function will be entirely devoid of entries, since the partial function will be undefined for all worlds.²⁹

As a result, it seems to me that in important respects, the Fregean error thesis more or less collapses into the Strawsonian error thesis. There is little difference between saying that people believe moral propositions that can’t accurately describe or misdescribe the world, and saying that people don’t in fact believe moral propositions because there are none. Similarly, there is little difference between the claim that people communicate propositions with no informational content, and the claim that people do not communicate at all.

Now, I have expressed my belief that this worry about moral communication and belief can be met (at least to a certain extent). A foundational thought that can be used to reply to this objection is the common idea that the semantic contents of sentences can come apart from their asserted or believed contents.³⁰ The proponent of the Strawsonian view cannot claim that moral speakers and thinkers assert or believe moral propositions. If, however, when someone assertively utters a moral sentence, there is some systematic and natural way of associating a proposition with the assertion, then we may consider that proposition to be what is asserted. If this strategy can account for assertion, then it can also plausibly work

²⁹Of course, this is just an artifact of my choice to represent these partially defined propositions in this way. We could adopt a way of including entries for “undefined” worlds, by letting there be a third truth-value (‘I’ for ‘indeterminate’). If we went this way, the rhetoric that I use would be lost, but the philosophical point would remain.

³⁰See, e.g., Stalnaker, 1978, 1987, and 1988; compare Soames, 2008.

for belief. The result would be an explanation of what is communicated or believed when someone makes a moral judgment.

As I see it, many of the traditional non-cognitivist views in the recent history of meta-ethics should be understood as reacting to the problem of moral communication in this way. Emotivists, prescriptivists, and expressivists should be seen as supplying an account of what we assert with our moral utterances, since typical proponents of these views do not think moral statements semantically express propositions.³¹

2.5.2 Against Theses 2S-3S

In the end, I believe it would take us too far afield to assess the plausibility of those traditional versions of non-cognitivism. Each view must be tested against various empirical data, and their responses to the notorious Frege-Geach problem must be examined. This is something I am not prepared to do in this work.³² Instead, I'll argue directly against the Strawsonian account. I'll not argue against the Fregean account further in this section.

In what follows, it will be useful to adopt some new terminology. I'll call Strawsonian presupposition 'expressive presupposition', and I'll call Fregean presupposition 'logical presupposition'. This is new convention helps to emphasize the distinction between the two kinds of presupposition. An expressive presupposition of a sentence is a requirement for the sentence to *express* a proposition. A logical presupposition is a *logical* requirement in order for the proposition a sentence expresses to have a truth-value.

My worry for the Strawsonian view can be roughly expressed by a question: in virtue of what linguistic features of moral judgments do they expressively presuppose the existence

³¹Cf. Ayer, 1952; and Hare, 1999. Compare Gibbard, 1990 and 2003; Schiffer, 2003: 256-8 and 1990: 611-2; and Schroeder, 2008. Of course, some of these philosophers provide non-standard accounts of propositions, under which there may be moral propositions semantically expressed by moral statements. In my opinion, these views fudge the (admittedly vague) line between semantics ("What is semantically expressed by a statement?") and pragmatics ("What is pragmatically asserted by a speaker making a statement?"). While the line is vague, it still exists – and we ought to acknowledge its existence in our linguistic theories.

³²There is a huge literature on these issues. See Schroeder, 2010 and 2008 for sophisticated introductions.

of objective values? Generally, sentences that plausibly suffer from expressive presupposition failure are ones whose parts *fail to determine meanings*. It is only with these cases do we have any explanation for why the sentences might fail to express propositions. But what parts of moral statements fail to determine meanings? We will see that there is no plausible answer available to the proponent of Theses 2S-3S.

The expressive presuppositions of a sentence are, by definition, what must be the case in order for the sentence to express a proposition. It is easy to list off a couple of these requirements. First, it seems that all relevant parts of the sentence must mean something: if ‘Fido’ or ‘is a dog’ go uninterpreted in a context, then (7) fails to express a proposition.

7. Fido is a dog.

There are roughly two ways for (7) to suffer from presupposition failure in this way. It might be that (7) is uttered by someone whose language doesn’t interpret ‘is a dog’. Thus, if a Russian uttered (7), it presumably wouldn’t express a proposition, since Russians typically don’t speak English. Alternatively, (7) may go uninterpreted in a context simply because the speaker doesn’t use ‘Fido’ to refer to anything in particular. Indeed, since there is nothing that ‘Fido’ refers to in my present context, it seems that (7) suffers from expressive presupposition failure right now. I certainly don’t know what proposition (7) would express, when uttered by me right now. I know of no dogs named ‘Fido’, I don’t intend ‘Fido’ to refer to anything on my desk, and I didn’t pick up the name from anyone who was using ‘Fido’ to refer to some particular dog.

Another sort of requirement is that the sentence be well-formed. It must really be a sentence of the salient language. Thus, (8) doesn’t, properly speaking, express a proposition in my context, even on the supposition that ‘Fido’ refers to some particular thing:

8. A dog, Fido is.

Although we may be able to associate a meaning with all the relevant parts of (8), it would be a bit forced to say that (8) expresses a proposition, since it is grammatically ill-formed.

Rather, it seems more natural to say that (8) is a confused attempt to express a proposition that might, instead, be expressed by (7). Someone who utters (8) isn't really speaking English at all.

Now, back to the question of why someone might think a moral judgment like (1) expressively presupposes that there are objective values.

1. What those little boys did was wrong.

There seem to be two options: first, that the existence of objective values is among the requirements that must be met in order for some relevant part of (1) to have meaning; and second, that there must be objective values in order for (1) to be grammatical. I hereby rule-out the second option without any argument. Sentence (1) is clearly grammatically well-formed, regardless of the existence of objective values. I'll concentrate on the question of why there must be objective values in order for (1) to go fully interpreted in a context.

Since (1) is a basic subject-predicate sentence, there are two relevant parts of it that must have meaning in order for it to express a proposition: 'What those little boys did' must mean something, and 'was wrong' must mean something. Thus, (a-b) represent expressive presuppositions of (1):

- a. 'What those little boys did' means something in the context.
- b. 'was wrong' means something in the context.

It is here that we arrive at a problem for the proponent of Thesis 2S. Where does the existence of objective values fit into the list of expressive presuppositions of (1)? I can only think of one rough stab at a possible explanatory answer to this question: the proponent of Thesis 2S must claim that, in order for (b) to obtain, there must be objective value properties. If this is correct, then (1) will not express a proposition unless there are objective values.

This is not a plausible view. There is good reason to think 'was wrong' means something irrespective of whether there are objective value properties. We may abstract away from the question whether there are objective values, and then ask whether 'was wrong'

has meaning. If, under this idealization, there is some reason to think ‘was wrong’ is meaningful, then there is reason to think the existence of objective values is orthogonal to the question whether our moral predicates have meaning.

So, what independent reason is there to think (b) is true? It is uncontroversial that, in general, our ordinary reports about belief and assertion provide evidence for whether a predicate is meaningful. If we ordinarily report people as having beliefs that involve a predicate, then there is reason to think the predicate is meaningful. Likewise, if we ordinarily report people as making assertions that involve a predicate, then there is reason to think the predicate is meaningful. This idea is not uncommon. For example, Stephen Schiffer writes,

A sentence has cognitive meaning, I propose, if it can be used to assert something that can be *believed*, in the strictest, most literal sense of the verb ‘to believe’. Belief is the backbone of cognitive meaning. (1990: 602)

Since we may take our reports about moral belief and assertion as evidence that moral sentences are meaningful, and moral sentences are meaningful only if their parts are, we therefore have reason to believe that (b) obtains in my utterance of (1). In general, we have defeasible evidence that moral sentences express propositions, irrespective of whether there are objective values.

Of course, I do not want to beg the question against the proponent of Theses 2S-3S. A proponent of this view might attempt to explain this evidence away *via* his alternative account of these reports of moral belief and assertion. My goal here, however, is to argue that whether moral predicates are meaningful is independent of the metaphysical question whether there are objective value properties. I have provided one reason to believe that these questions are independent of one another, and I can think of no good reason to believe otherwise.³³

³³In the next chapter, I provide the materials for another argument to the effect that these questions are independent. I argue that there are conceptions of properties as “meanings” and conceptions of properties as more “fundamental” constituents of the world. I also argue that objective value properties are supposed to be the latter kind of property. But since the question of the existence of “meanings” is independent of questions about “fundamental” properties, there is very good reason to think moral predicates have meaning even if there are no objective values. See the next chapter for details.

A proponent of Theses 2S-3S may be unsatisfied with my arguments thus far. Recall that I have supposed that the only plausible cases of expressive presupposition failure are sentences with parts that *fail to determine meanings*. It is here that the error theorist may balk. Instead, he may claim, other plausible cases of expressive presupposition failure are sentences whose parts *fail to determine reference*. The theorist who distinguishes meaning from reference will be able to make this move. Thus, Thesis 2S may still be plausible, so long as moral judgments can be said to suffer from reference failure.

The proponent of this reply on behalf of Thesis 2S, then, will accept that (c-d) are further expressive presuppositions of (1).

- c. 'What those little boys did' refers to something in the context.
- d. 'was wrong' refers to something in the context.

Now, the proponent of Thesis 2S may go on to claim that, in order for (1) to express a proposition, each of (a-d) must obtain. Let us suppose that none of (a-c) can plausibly be taken to be that in virtue of which (1) suffers from presupposition failure. Instead, it seems that (d) must be the culprit. The problem, for this view, is to make explicit what it means for a predicate expression to "fail to refer". There are a couple of possible ways of fleshing out this idea, and neither of them, as far as I can tell, results in a plausible view.

First, one might take predicate reference to consist in the existence of a non-empty extension at the actual world. On this view, a predicate secures reference if and only if there are some actual things to which the predicate applies. Unfortunately, this view results in absurd semantic consequences, since it is committed to the following claim:

- ρ . Any sentence with a predicate that applies to nothing actual expresses no proposition.

The proponent of Thesis 2S, if he takes this route, must claim that (9) fails to express a proposition:

- 9. Emily is a witch.

I take this, in itself, to be a *reductio* of the view. Of course, this is not the only implausible consequence of the view. Any predicate that doesn't actually apply to anything will occur in many sentences which manifestly do express propositions, but (ρ) implies that they do not. The counterexamples to (ρ) quickly multiply.

Second, one might take predicate reference to consist in the existence of a non-empty extension in some possible world. On this view, a predicate secures reference if and only if there are some possible things to which the predicate applies.³⁴ As I see it, this view is just as implausible as the preceding one, because it is committed to a slight modification of (ρ):

ρ' . Any sentence with a predicate that applies to nothing possible expresses no proposition.

If he accepts this view, then the proponent of Thesis 2S must claim that (10) fails to express a proposition:

10. Casey is a round square.

This, again, amounts to a *reductio*. Although (10) couldn't possibly be true, since there could be no round squares, it clearly expresses a proposition, contrary to what (ρ') implies. This completes my quick survey of how a proponent of Thesis 2S-3S might make good on his claim that moral judgments suffer from expressive presupposition failure.³⁵ I conclude that the proponent of Theses 2S-3S cannot save his view by taking this route. If there is another alternative available, which doesn't fail for reasons similar to the above, then I am unaware of it.

³⁴Joyce suggests at this view in 2001: 7-8. He remains entirely agnostic, however, about how to conceive of reference for predicates. He says that it is a 'task for metaphysicians'; cf. 2007a: section 4.

³⁵One might take predicate reference to consist in the existence of a non-empty extension in some possible or *impossible* world. On this view, a predicate secures reference just if there are some possible or impossible things to which the predicate applies. Barring the problem of how to understand the notion of an impossible world, this view, it seems to me, is no better than the others. It would rule-out many perfectly meaningful, proposition-expressing sentences as failing to express propositions. In addition, if we were able to provide a concise conception of the notion of impossible world, it would also be unclear why moral predicates could not apply to things in impossible worlds.

2.6 Discourse Presuppositions and Entailment

In the preceding sections, I have argued that neither the pragmatic rendering of the discourse thesis, nor an expressive rendering of the thesis, can yield a plausible error theory. I have not argued against the Fregean interpretation of the discourse thesis, according to which our moral judgments logically presuppose the existence of objective values. I set that view aside until the end of the present chapter.

In this section, I explicate a final rendition of the discourse thesis and a final version of the error thesis. According to this new error theory, our moral judgments analytically entail that there are objective values. Near the end of the section, I suggest at a preliminary worry for the resulting view, which applies equally well to Theses 2F-3F. The worry raises important metaphysical issues, which leads us into the third chapter of this work, where I discuss how to understand the metaphysical component of these error theories.

2.6.1 Analytic Entailment and Non-Negotiability

I suppose that, in order for someone to be a bachelor, he has to be an adult male. Likewise, a person cannot literally be a bachelor unless he is unmarried. I take these constraints on bachelorhood to be absolutely necessary. There is no genuine, interesting sense of ‘possible’ on which it would be true to say that it’s possible for there to be a bachelor that is not an adult, not male, or not unmarried. Of course, a sentence of the form of (11) isn’t a logical contradiction, so what explanation could we have for this necessity?

11. x is a bachelor but x is not an unmarried adult male.

A standard answer to this question is that something about what we mean by ‘bachelor’ makes (11) impossible to satisfy. The first conjunct of (11) necessitates that the second conjunct is false, since it’s part of the meaning of ‘bachelor’ that (12) is true.

12. If x is a bachelor, then x is an unmarried adult male.

In the standard terminology, it is a *conceptual* (or *analytic*) truth that all bachelors are unmarried adult males; likewise, (11) is conceptually or analytically false. In my terminology, conditional claims like (12) express ‘analytic entailments’.

Richard Joyce explicates his version of the error theory roughly in terms of the notion of analytic entailment. On Joyce’s view, in order to tell whether a certain claim is a conceptual truth, it must pass what he calls ‘the translation test’. He also calls conceptual truths ‘non-negotiable’ propositions. He writes,³⁶

Imagine that we were to encounter a population speaking a quite different language to our own, most of which we have translated and tested to our satisfaction, and we find that they have a concept that appears rather like our concept of [*bachelor*] ... – call their term [“schmachelor”] – but we also find that they don’t endorse one of the three propositions about [bachelors (see above)]. If that would be sufficient for us to decide *not* to translate [“schmachelor”] into [“bachelor”], then the proposition in question must be a non-negotiable part of our concept [*bachelor*]. (2001: 3)

Joyce applies this notion of non-negotiability not simply to our concepts. Apparently, for Joyce, when part of a concept is non-negotiable, there are corresponding non-negotiable parts of any discourse that uses the concept. Thus, Joyce would say that the belief that there are unmarried adult males is a non-negotiable part of our talk about bachelors – “the falsity of this belief would be sufficient to sink the whole [bachelor] discourse.”³⁷ If we were to find out that nothing in fact is an unmarried adult male, then we would have to retract any of our claims to the effect that certain people are bachelors. Elsewhere, Joyce even applies a notion of presupposition directly to discourses.³⁸ He thus would say that our bachelor discourse presupposes the existence of unmarried adult males.

With Joyce’s terminology in mind, as well as the concept of analytic entailment, we are in a position to state definitions of ‘non-negotiable’ and ‘discourse presupposition’. First,

³⁶Joyce’s example uses ‘phlogiston’. In the following passage, I have changed it to ‘bachelor’.

³⁷Joyce, 2001: 4.

³⁸Joyce, 2011: 323-5.

a proposition is *non-negotiable* to our moral discourse just in case our moral discourse presupposes it. Second, a proposition is *presupposed by our moral discourse* if and only if our moral judgments analytically entail that proposition.³⁹ These definitions help to make apparent a connection between presupposition and entailment, viz., there is a concept of presupposition that is defined in terms of analytic entailment. As far as I know, this concept of presupposition is not one that is heavily studied by linguists under that name. It is, nonetheless, an ordinary notion of presupposition, which Joyce puts to an interesting use.

2.6.2 The Discourse Thesis in Terms of Entailment

With the notion of analytic entailment, as well as the attendant notions of discourse presupposition and non-negotiability, we are in a position to state a new version of the discourse thesis.

Thesis 2E: The Entailment Discourse Thesis. Every moral judgment analytically entails that there are objective values.

The thought underlying Thesis 2E is that upon correct analysis of the meanings of our moral terms, we find that they can apply to things only if there are objective values. Put another way, our moral discourse presupposes that there are objective values.

As an illustration, consider an analogy to our usage of ‘bachelor’. On Thesis 2E, part of the meaning of ‘wrong’ is the concept of an objective value, just as part of the meaning of ‘bachelor’ is the concept of being unmarried. If I point to a man and call him a ‘bachelor’, then what I say is false if he is not unmarried. Likewise, when I utter ‘what those little boys did was wrong’, the proposition expressed by my utterance is false, so long as there are no objective values. In general, whenever you make a judgment that analytically entails something that is not true, your judgment ends up being false.

This fact about analytic entailment, as I am understanding the notion, makes it distinct from the notions of logical and expressive presupposition. When a sentence expressively

³⁹Cf. Lewis, 2005: 317.

presupposes a proposition, and that proposition is false, the sentence fails even to express a proposition. This is not so with analytic entailment: strictly speaking, analytic entailment is a relation that bears between propositions expressed by sentences. If no proposition is expressed by a sentence, then it makes no sense to say that the sentence analytically entails anything. When a sentence logically presupposes a falsehood, the sentence fails to be true and fails to be false. This is not so with analytic entailment: if the proposition expressed by a sentence is false, then so is the sentence.⁴⁰

With this distinction between logical (Fregean) and expressive (Strawsonian) presuppositions and discourse presuppositions in mind, we are in a position to state a new error thesis.

Thesis 3E: The Entailment Error Thesis. Since there are no objective values, but our moral judgments analytically entail that there are, all our moral judgments are false.⁴¹

On Thesis 3E, speakers make the error of committing themselves to falsehoods whenever they make moral judgments: if you have a moral belief, then it is false; if you make a moral assertion, it is false.

There are some differences, and also some similarities, between Thesis 3E and Theses 3F-3S. One obvious respect of similarity is that each of these versions of the error thesis implies that none of our (basic) moral judgments is true. All our moral judgments end up failing to secure truth, either because they fail to express propositions, because they express gappy propositions, or because they express propositions that are false. This is where differences arise between the positions. Thesis 2S should be separated from the other two positions, since, according to it, the proper semantics for moral judgments will

⁴⁰There are broader notions of entailment that one might like to call ‘analytic entailment’, on which all logical presuppositions are entailments. But I hereby stipulate that ‘analytic entailment’, in my mouth, is to pick out non-presuppositional entailments. One might like to call them ‘mere analytic entailments’.

⁴¹Of course, this version of the error thesis must be restricted to moral judgments that do not suffer from one of the kinds of semantic presupposition failure. If ‘x’ doesn’t refer in ‘x is good’, then the sentence might be neither true nor false. This is a trivial technical detail, which I’ll not address in the main text.

never assign any of them a proposition. On Theses 3E and 3F, the proper semantics for moral judgments will always take them to express propositions.

But there is an obvious difference between Thesis 3E and Thesis 3F with respect to the propositions that are expressed by moral judgments. I noted above that error theorists tend to hold it to be metaphysically necessary that there are no objective moral values. On Thesis 3F, this results in moral judgments expressing necessarily gappy propositions, whereas for Thesis 3E, the result is that moral judgments express necessarily false propositions. As I have been conceiving propositions, they are to be individuated according to their truth-value assignments to possible worlds. On this assumption, Thesis 3E identifies all propositions expressed by moral judgments with a unique necessarily false proposition, while Thesis 3F implies that moral judgments one and all express the same necessarily gappy proposition.⁴²

2.6.3 A Worry for Theses 3E and 3F

In this section, I will sketch a problem for the coherence of Theses 3E and 3F and the positions that they represent. In the end, I believe that this is only an apparent problem. It is a worry that can easily be met, but only by first recognizing various metaphysical distinctions. Thus, I intend the worry of this section to lead us smoothly into the next chapter, where I discuss some issues involving the metaphysics of error theories. I begin by discussing the worry as it applies to Thesis 3E. The worry generalizes to Thesis 3F.

According to the proponent of Thesis 3E, morality is not objective, i.e., there are no objective values. On this view, an objective value would be a certain kind of property. Moreover, any of our moral judgments would analytically entail that there are these objective value properties. So far so good.

⁴²As I see it, the proponent of Thesis 3E does have the problem of explaining what we are communicating with our moral judgments. I take this to be a serious worry, although I see no reason why the proponent of this view cannot adopt a position on moral communication analogous to the position I said Theses 3F-3S would have to accept. Thus, it seems to me that any of these versions of the error theory are in roughly the same boat with regard to moral communication and belief.

Now, analytic entailment holds in virtue of the meanings of parts of sentences. The reason why (12) expresses an analytic entailment is that it is part of the meaning of ‘bachelor’ that if ‘bachelor’ applies to something, then so does ‘unmarried adult male’. In some common terminology, this is just to say that the property of being a bachelor contains the properties of being unmarried, being an adult, and being male. Underlying this terminology is a conception of predicate meanings as properties. The property of being a bachelor just is the meaning of ‘is a bachelor’ as we use the phrase. It follows that analytic entailment holds partially in virtue of the properties expressed by the parts of sentences. Thus, if there is no property that we can associate with a certain predicate within a sentence, then we will be unable to say that the sentence can stand in a relation of analytic entailment. No properties, no analytic entailment.

Therefore, it seems that the proponent of Thesis 3E must accept the existence of properties corresponding to the meanings of moral predicates. I.e., ‘was wrong’ in my utterance of (1) expresses the property of being wrong, since (1) stands in an analytic entailment with the existence of objective values.

1. What those little boys did was wrong.

There is a pressing question that the proponent of Thesis 3E must, however, answer: in virtue of what is it that (13) expresses an analytic entailment?

13. If what those little boys did was wrong, then there are objective values.

The entailment certainly doesn’t hold in virtue of the meaning of ‘what those little boys did’. Instead, it seems that it must be in virtue of the meaning of ‘was wrong’, viz., wrongness. So, on the picture of analytic entailment that I sketched above, the proponent of Thesis 3E must claim that wrongness is an objective value: it is part of our concept of wrongness that it is an objective value.

But now the proponent of Thesis 3E seems to be in something of a pickle. On the one hand, he wants to reject the existence of properties that would have the features of

objective values. On the other hand, he wants to accept that moral judgments analytically entail the existence of objective values. It seems that the two hands cannot be reconciled with one another. In order for the analytic entailment to hold, wrongness must exist, and so an objective value must exist.

The worry generalizes to Thesis 3F. The proponent of this view rejects the existence of objective values, and yet holds that there is a crucial entailment-like relation that holds between moral judgments and the existence of objective values. It seems that this entailment-like relation cannot hold unless wrongness, an alleged objective value, exists.

2.7 Conclusion

The main goal of this chapter has been to introduce a variety of precise ways of understanding the discourse thesis and the error thesis. In effect, I have presented two defective embodiments of the error theory and two more promising versions of the view.

I argued against the pragmatic version of the error theory on the grounds that its version of the error thesis is semantically implausible. From the mere fact that a speaker presupposes something, we cannot conclude that what he says is untrue. This version of the error theory, it seems to me, has no hope of being developed into a plausible account of our moral language and thought, unless it is developed in a way that rejects the idea that moral judgments are systematically untrue.

I argued against the Strawsonian, or expressive, version of the error theory on the grounds that it cannot make good on its version of the discourse thesis. I was able to find no coherent, plausible way of cashing-out the idea that moral judgments expressively presuppose the existence of objective values.

I find the remaining versions of the error theory much more worth developing than either of the former positions. However, I have mentioned some initial worries about both views. First, neither view comes pre-packaged with an account of what we are believing or

communicating when we make moral judgments. Second, both views seem to have some metaphysical problems that must be worked out. I now turn to those metaphysical issues.

CHAPTER 3

ERROR THEORIES AND ONTOLOGY

3.1 Introduction

In chapter 1, I introduced the metaphysical thesis as a claim about the alleged “objectivity” of morality:

Thesis 1: The Metaphysical Thesis. There is something seriously defective with moral reality, viz., it is supposed to be objective.

Throughout chapter 2, I supposed that, according to a standard error theorist, morality is objective just in case there are objective values. I also assumed that an objective value would be a certain kind of property. Since these assumptions are somewhat standard in the literature on the error theory, they were justified. I also argued that readings of the discourse thesis in terms of pragmatic presupposition and expressive presupposition result in implausible error theories. I tried to show that the best version of the discourse thesis would be formulated in terms of the notion of analytic entailment or logical presupposition.

At the close of the chapter, however, I introduced a problem for these versions of the error theory. Roughly, the problem is that, if there are no objective value properties, then moral judgments can entail nothing. It thus seems that the discourse thesis is incompatible with the metaphysical thesis. In this chapter, I’ll argue that this puzzle – the problem of entailment – arises from a confusion in how to theorize about properties. Error theorists need not worry too much about this problem.

In section 2, I outline how the metaphysical thesis is conceived in the literature. Although it comes in a variety of forms, I argue that it usually should be understood as rejecting the existence of objective value properties (or else moral properties in general).

In section 3, I explicate the problem of entailment in greater detail and I highlight some important semantical assumptions of the problem. I argue that it would be beside the point for the error theorist to reject any of these semantical assumptions in reaction to the problem. That leaves the error theorist no option but to pinpoint some logical fallacy in my formulation of the problem.

In section 4, I digress from the main theme of the chapter to address the question of why we should try to save the error theory in the first place. If it is so easy to show that it is internally incoherent, why not just toss it out the window? My lengthy reply comes in two parts. First, because the problem of entailment arises in other areas of philosophical and scientific theorizing, we have reason to hope for a solution to it. Second, I draw a parallel between the problem of entailment and what I call the ‘problem of reference failure’. I claim that, because we expect a solution to the problem of reference failure, we therefore ought to expect a solution to the problem of entailment.

In section 5, I claim that it is antecedently plausible to think that there are many different ways of thinking of the nature of properties. These are differences in conceptions of properties, not merely differences in theories about the properties. I call this view ‘pluralism about properties’. As a result of this pluralism, it is possible that the problem of entailment arises from something like an equivocation on ‘property’.

Finally, in section 6, I present my solution to the problem of entailment: it arises because ‘property’ does not have a uniform meaning throughout its formulation. The metaphysical thesis comes in two main forms, corresponding to two crucially different ways of interpreting our talk of properties. On neither version of the metaphysical thesis is there a proof that the error theory is internally incoherent. At the end of the section, I argue that a clear-headed error theorist ought to accept both of these formulations of the metaphysical thesis, therefore two versions of the error thesis.

3.2 The Metaphysical Thesis in the Literature

Mackie begins *Ethics: Inventing Right and Wrong* with a “bald statement” of his metaphysical thesis: “There are no objective values.” He calls this claim ‘moral skepticism’, and adds,

...moral skepticism is an ontological thesis, not a linguistic or conceptual one. It is not, like the other doctrine often called moral subjectivism, a view about the meanings of moral statements. Again, no doubt, if it is to be at all plausible, it will have to give some account of their meanings.... But this too will be a development of the theory, and not its core. (1977: 18)

That an error theorist will hold such an ontological thesis has been taken for granted in the literature. While comparing his account of values to Mackie’s, David Lewis writes,

Strictly speaking, nothing shall get the name [of ‘value’] without deserving it perfectly. Strictly speaking, Mackie is right: genuine values would have to meet an impossible condition, so it is an error to think there are any. Loosely speaking, the name may go to a claimant that deserves it imperfectly. ... Then is my position a form of realism about values? – Irrealism about values strictly speaking, realism about values loosely speaking. The former do not exist. The latter do. (1989: 93)

Nolan (et al.) write,

Moral discourse ... presupposes that there are objective moral truths, ... but the ontology this talk seems to presuppose is strange, non-natural, and Moorean. So we seem to face a dilemma: either embrace an alien ontology, which we have good reason to reject; or become eliminativists about morality and agree that all positive moral claims are false, since there are no queer objectively prescriptive entities of the sort such claims presuppose. (2005: 307)

This ontological version of the metaphysical thesis thus seems to be the main form in the literature. On this view, there are no objective moral value properties; or, more generally, there are no moral properties at all. On a stronger version of the thesis, this nonexistence claim holds necessarily: there can be no moral properties.¹

¹See also Garner, 1990; Joyce, 2001: 134; and Balaguer, 2011.

But there is another strand of thought that may be attributed to error theorists. According to this idea, the main metaphysical thesis of the error theory is that, although there are moral properties, none are (or can be) instantiated. R.M. Hare writes,

[Mackie's] view ... is that in ordinary use [moral] words connote objective properties of actions, etc.; the error ... is a factual one; people think that actions have these moral properties, but they do not have them, since there are no such properties (not in the sense that the words connote nothing, but in the sense that nothing actually has the properties which they connote). (1981: 78)

Similarly, Shafer-Landau writes,

Moral nihilism is the view that either there are no moral properties, or there are, but none that are ever instantiated. On this view, there is no moral reality at all. (2006: 210)

Shafer-Landau cites Mackie (among others) as being a primary example of a moral nihilist.²

I'll register two firm opinions about the metaphysical part of the error theory in the literature. First, the literature is in something of a disarray with respect to how to formulate the metaphysical thesis. Sometimes, it is an ontological claim, about the nonexistence of certain kinds of entities. Other times, it seems to presuppose the existence of these kinds of entities, and it asserts something about them to the effect that they are defective: in the above passages, the defect is that the moral properties are not instantiated. Second, although there is this variation in the literature, the dominant form of the metaphysical thesis is the ontological construal, according to which those bits of moral ontology do not exist. When I say it is 'dominant', I mean that it is the form most typically explicitly asserted to be the error theorist's metaphysical thesis. In what follows, I'll take this version of the thesis as the canonical version.

²See also Feldman's (1979) review of Mackie (1977), where he distinguishes these two disambiguations of the metaphysical part of Mackie's view.

3.3 The Problem of Entailment

In this section, I'll outline the problem of entailment in detail, I'll highlight some key semantical assumptions used in my formulation of the problem, and I'll defend them from some objections.

3.3.1 Formulations of the Discourse and Metaphysical Theses

What I consider the dominant version of the metaphysical thesis may be formulated as follows:

Thesis 1O: The Ontological Metaphysical Thesis. There are no moral properties; in specific, there are no objective values.

The specific part of Thesis 1O banishes only objective value properties from existence. Thus, if it were true, then there would presumably be no such properties as moral goodness and badness. The thesis taken generally, however, entails that no moral properties whatever exist: not only are there no value properties; neither are there properties of moral rightness, wrongness, and obligation; nor properties of moral virtue and vice; nor properties of being morally deserving or undeserving; nor justice or injustice. For the purposes of detailing the problem of entailment, the differences between the general claim and the more specific part of Thesis 1O will not matter. What matters is that the thesis rejects the existence of a certain kind of property.

The entailment formulation of the discourse thesis asserts that the semantic relation of entailment holds between moral statements and the existence of moral properties:

Thesis 2E: The Entailment Discourse Thesis*.* Every moral judgment analytically entails that there are moral properties.

Another possible way of formulating the discourse thesis would have it that moral judgments entail that there are objective value properties. This is how I construed the entailment version of the discourse thesis in the preceding chapter. Thesis 2E* is only a slight rewording of that construal: since objective values would be moral properties, Thesis 2E*

is slightly weaker than my prior construal. My goal in formulating it in this way is to make transparent the fact that the typical error theorist thinks moral judgments entail the existence of a certain kind of property.

3.3.2 Two Semantic Constraints

These two elements of the error theory jointly run up against a pair of independently plausible semantic constraints. The first constraint is that if a judgment or statement has entailments, then it must express a proposition. This is motivated by the semantic intuition that entailment is fundamentally a relation between propositions. That is, entailment holds in virtue of relations between the propositions expressed by statements.

SC1: No Propositions, No Entailment. If a statement doesn't express a proposition, then it entails nothing.

There are two kinds of counterexample to SC1. First, one might claim that statements can have entailments, irrespective of whether they express propositions. For example, in a formal language, a statement like ' $\exists x(Red(x) \wedge Round(x))$ ' may logically (or formally) imply ' $\exists xRed(x)$ ', and this will hold merely in virtue of the form of the statements. This is a counterexample to SC1 on the supposition that one of the predicates ('*Red*' or '*Round*') is uninterpreted in the formal language: if '*Round*' has no meaning, then the first sentence expresses no proposition (even if we suppose that '*Red*' means something). Although this is a genuine counterexample to SC1, it is permissible for me to ignore it. This is because no error theorist formulates his view in terms of formal logical implication. Thus SC1 should be read as restricted to entailments that do not hold in virtue of formal logical implication.

Second, one might claim that there are statements that have entailments, even though they are semantically indeterminate and so do not express unique proposition. For example, 'John is a bald older fellow' does not express a unique proposition, since it includes the vague predicate 'bald': at best, it expresses a range of propositions, corresponding to the various ways of making precise the meaning of 'bald'. Nonetheless, it seems that this

statement entails that John is a bald man. Although, again, this is a genuine counterexample to SC1, I will, again, ignore it. This is because, for purposes of this chapter, I'll suppose that moral statements are not indeterminate in this way. Issues involving semantic indeterminacy are orthogonal to the issues discussed in this chapter (though I will be concerned with semantic indeterminacy in later chapters).

The second constraint is that a statement will express no proposition if one of its constituent predicates fails to express a property. This is motivated by the standard idea that a principle of compositionality about content holds. The content of a statement depends upon the contents of its constituents, in such a way that if some of the parts have no contents, then neither does the whole itself. For example, if I utter 'Obama is bumbly-doodle', my statement will not express a proposition, so long as 'bumbly-doodle' is mere nonsense.

SC2: No Properties, No Propositions. If a predicate doesn't express a property, then no statement containing it expresses a proposition.

There are two worries for SC2, but they do not worry me. First, one might reject the strong form of compositionality that seems to be presupposed by SC2. There are a variety of words in English that seem to show it to be false. For example, 'damn' in 'Disney was a damn bigot' doesn't have any genuine content, since 'Disney was a bigot' intuitively has the same content. Although this may undermine a strong version of compositionality, it does not prove SC2 false. 'Damn' isn't functioning as a predicate, and so we have here no counterexample to SC2.

Second, one might claim that some predicates, for example, 'tall', do not semantically express properties, but rather, need to be enriched by context in order to be associated with a property. But we must be clear on the distinction between content and character.³ There are two ways of thinking of what a predicate "semantically expresses". We may say that a predicate expresses a content, e.g., a function from worlds to extensions at those worlds. Or we can say that a predicate expresses a character, e.g., a function from contexts of utterance

³For references, see Lewis, 1980.

to contents. Both of these concepts are equally useful in semantic theory, but it is only content with which I am here concerned.

3.3.3 Statement of the Problem of Entailment

Provided that SC1 and SC2 are true, Theses 1O and 2E* are incompatible. If Thesis 1O is true, then it would seem that all moral predicates fail to express properties. In specific, if there are no objective values, then it would seem that the predicates ‘morally good’ and ‘morally bad’ would not express any properties of moral goodness and badness. This result, when combined with SC2, implies that moral statements do not express propositions. As a result, SC1 then implies that moral statements do not entail anything. But if moral statements do not entail anything, then they cannot entail the existence of moral properties. Therefore, it seems that Thesis 2E* cannot be true if Thesis 1O is true.

What can we say on behalf of the error theorist? I firmly believe that it would be too costly for an error theorist to reject SC1 and SC2 in reaction to this problem. But these are apparently the only assumptions that are required to derive the result that Theses 1O and 2E* are incompatible. It seems that the only possibility available to the error theorist is to point out an equivocation on a crucial term in the derivation. I’ll propose that the term ‘property’, as it occurs in Thesis 1O, does not mean the same as ‘property’ in SC2. But first, I’ll digress to discuss the question of why we should try to rescue the error theory from incoherence.

3.4 Why Rescue the Error Theory?

Some readers may have doubts about the project of rescuing the error theory. The alleged proof that this version of the error theory is inconsistent is so simple as almost to be trivial. Whenever a view seems trivially provably inconsistent, a reasonable reaction is that it must not have been formulated correctly. So perhaps it was a mistake in the first place to think Thesis 1O could be true. There is the alternative of formulating the metaphysical

part of the error theory in terms of the non-instantiation of moral properties, instead of the nonexistence of moral properties. This seems to have been R.M. Hare's idea in the above passage.⁴ The following reformulation is not apparently trivially inconsistent with Thesis 2E*:

Thesis II: The Instantiation Metaphysical Thesis. There are moral properties, but they are not instantiated; in specific, although there are objective values, they are not instantiated.

If Thesis II is the chosen version of the metaphysical thesis, then we may move on to discuss the plausibility of moral error theories without having to introduce more complicated distinctions. So why not do that?

I have a quick reply and a long reply to this worry. The quick reply comes in two parts. First, if Thesis II were our only precise formulation of the metaphysical thesis, then we would knowingly be misinterpreting the dominant construal of the metaphysical thesis in the literature. Mackie and others explicitly chose to formulate the metaphysical thesis as a claim about the nonexistence of a bit of moral ontology, and so we ought at least to see whether there is a coherent way of making sense of this formulation. Second, in my opinion, if we merely chose Thesis II as our version of the metaphysical thesis, then we'd be ignoring important distinctions in the philosophy of properties. Any subsequent discussion would thus be at risk of running afoul of these distinctions. I take these to be reasons enough to pursue the project of replying to the problem of entailment on the error theorist's behalf.

One might, however, be willing to throw Mackie's explicit formulation to the dogs, and to reject the distinctions to which I aim to attend. For such a reader, I have a longer response, which, again, comes in two parts.

⁴Feldman (1979) also seems to express this idea, and Shafer-Landau (2006) seems at least to leave it open.

3.4.1 Parallels in the Philosophy of Mind and Science

First, there are various long-standing debates in the philosophy of mind about whether certain kinds of properties exist. Some have claimed that there is no such thing as mentality; others have claimed that there are no colors.⁵ Consider the following concrete example (which I have fabricated): an eliminativist about mental properties might claim,

“If there were mentality, it would have to be an irreducible property of a queer sort of mental substance, viz., a *mind*. But there are no such substances; the success of the empirical sciences has made it very plausible that we can have a complete explanation of the world that doesn’t require the existence of minds. So, there are no mental properties. But we make statements that entail the existence of mentality, e.g., statements that involve belief, intention, and the like. So, these statements are all false.”

While this view may not be all that plausible, it seems, at least, to be a coherent stance to take. For example, someone might react in this way to Descartes’ views about mentality. But the problem of entailment would have it that this eliminativism about mental properties is trivially incoherent. In addition, it would be uncharitable to claim that those involved in debates over the existence of mentality are really presupposing the existence of mental properties, while debating about whether these kinds of properties are ever instantiated. Insofar as we can make sense of these debates, we have reason to try to solve the problem of entailment.

There are also debates in the philosophy of science about whether certain kinds of properties exist. For example, some have claimed that there is no such property as absolute presentness (or absolute simultaneity between events).⁶ Consider another fabricated concrete example, this time about absolute presentness:

“The absolute present, were there to be such a property, would require that the laws of nature privilege a certain way of carving up the world into time-slices. But the laws of nature do no such privileging; the plausibility of Einstein’s theory of special relativity is evidence enough for that. So, there is no

⁵See, e.g., Lewis, 1995.

⁶See, e.g., Janis, 2010.

absolute presentness. But we make statements that entail the existence of such a property, e.g., statements to the effect that some event is happening *now*. So, these statements are all false.”

Again, although this view may not be all that plausible, it seems at least coherent, contrary to the upshot of the problem of entailment. And again, it would be uncharitable to suppose that all who attend to the debate over the absolute present are really presupposing its existence, while debating about whether it is instantiated. Therefore, the problem of entailment over-generates: it tells us that certain views are self-contradictory, when they seem not to be. Are we really prepared to throw these views in the philosophy of mind and science to the dogs?

3.4.2 Parallel to the Problem of Reference Failure

In the second part of my long reply, I’ll invoke a parallel between the problem of entailment and what I’ll call the ‘problem of reference failure’. It is part of ordinary adult thought that some names just fail to refer: ‘Santa Claus’ is a good example. If you ask an adult (who isn’t a philosopher!) whether Santa exists, he’ll tell you ‘No’; likewise if you ask him whether ‘Santa Claus’ refers to something. Of course, a child might answer ‘Yes’. You can explain this by reference to a condition the child probably satisfies: her parents told her that Santa comes down the chimney on Christmas Eve to deliver presents, and she believed them, even though what they said was false. This suggests that it is another uncontroversial part of ordinary thought that statements about Santa do indeed express propositions. If a child can believe that Santa comes down the chimney, then ‘Santa comes down the chimney’ expresses something that’s believable, i.e., a proposition. Similarly, if it’s false that Santa comes down the chimney, then that sentence expresses a proposition. Ordinary thought thus seems committed to the following:

OMT: Ordinary Metaphysical Thesis. Santa Claus doesn’t exist.

ODT: Ordinary Discourse Thesis. Statements in which ‘Santa Claus’ occur still express propositions (witness OMT).

But this combination of views results in a problem when we apply to it the semantical idea underlying our second semantic constraint. SC2 itself says that if a predicate doesn't express a property, then statements in which the predicate occurs don't express propositions. The semantical idea that underlies SC2 is that unless an expression has content, complex expressions made up from that expression do not express propositions. When applied to names, the idea becomes the following:

SC2: No Contents, No Propositions.* If a name doesn't have a content, then no statement containing it expresses a proposition.

The problem of reference failure, then, goes as follows. Whereas the content of a predicate would be the property it expresses, the content of a name would be its referent. It would seem to follow from OMT, however, that 'Santa Claus' doesn't have a referent, and so neither does it have content. But then, by SC2*, statements containing 'Santa Claus' don't express propositions, contrary to ODT. Thus ordinary thought about mythical characters like Santa seems to be internally incoherent.

There is a strict analogy between the problem of entailment and the problem of reference failure. Both problems arise for views that include a certain metaphysical thesis and a certain discourse thesis. The metaphysical part denies the existence of a certain piece of ontology; the discourse thesis makes a claim about statements that seem semantically to require that piece of ontology. The problem, for both views, arises from a semantic constraint on what is required in order for the relevant statements to express propositions.

I now note that the reader who is skeptical of the project of trying to save Theses 1O and 2E* is probably not skeptical of the project of trying to save OMT and ODT. It is taken for granted by much of the literature on reference failure that there is some adequate way

of reconciling ordinary thought with systematic semantics.⁷ Most plausible solutions will save OMT by reinterpreting it as a genuinely true nonexistence claim.⁸

My goal here is to convince the skeptical reader that, because of the parallel between the problem of reference failure and the problem of entailment, it would be irrational to expect a solution to the former but not to expect a solution to the latter.⁹ So, if you think there just has to be a solution to the problem of reference failure, then you likewise should think there must be a solution to the problem of entailment. Put another way, if you are easily willing to throw the error theory out the window, then you ought also to be willing to throw ordinary thought out the window. If you're not equally willing to give up on ordinary thought, I would suggest that you reconsider your willingness to give up on the error theory that consists partially in Theses 1O and 2E*.

3.5 Pluralism about Properties, and its Rivals

The pluralist about properties thinks our usage of terms like 'property', 'feature', and 'attribute' is something of a mess. There are many different, not obviously compatible, roles that properties are called on to fill. As a result, when a pluralist attempts to systematize our discourse about properties, he leaves open whether it should be subdivided into many

⁷See, e.g., Soames, 2005: 349-53 and 2002: 89-95; van Inwagen, 1977; Salmon, 1998; and Caplan, 2001.

⁸Even direct reference views about names try to save ordinary thought. According to these views, if we interpret OMT strictly, it either comes out false or else doesn't express a proposition. That is, such views hold that when people utter claims like OMT, what they "really mean" is that *there is no* person who is identical with Santa; or else that *there is no* thing that satisfies the descriptive properties are generally associated with the name 'Santa Claus' (see the references to Soames, Salmon, and van Inwagen, above). Clearly, even these direct reference views hold that the idea that underlies OMT is true. That is all I need for my argument in this section, for, in what follows, I'll argue that while Thesis 1O is false on one interpretation, it is true on another.

⁹I am even willing to claim that you ought to expect the solutions to be parallel. My solution to the problem of entailment is roughly parallel to a Fregean solution to the problem of reference failure. One way of construing Frege's solution is to say that (i) he accepts OMT and ODT, but (ii) he also accepts two kind of content: sense and reference. And while he accepts SC2* on its sense-disambiguation, he rejects it on its reference-disambiguation. As he would put it, sentences that involve reference failure (e.g., 'Odysseus deeply asleep disembarked at Ithaca') still "express thoughts", even though they do not "have reference" (they do not refer to the True or to the False). We will see that our error theorist should accept correlates of OMT and ODT, and should draw a parallel distinction to the one Frege draws.

smaller, internally coherent, parts. Someone who doesn't acknowledge this alleged disorder in our talk of properties, or who doesn't leave open the possibility that property talk should be divided in this way, counts as a monist about properties.

In this section, I'll provide references to support the idea that we ought to take pluralism seriously, and I'll introduce a few important notions required for a proper understanding of the view. Then I'll more concisely distinguish versions of pluralism from its rivals. I'll also provide a brief exposition of some of the roles that properties are expected to fill. The upshot of the section will be that it is very plausible to think an equivocation occurs in the problem of entailment.

3.5.1 Lewis on Pluralism

The canonical texts from which property pluralism arises are to be found in David Lewis's work.¹⁰ He writes,

It's not as if we have fixed once and for all, in some perfectly definite and unequivocal way, on the things we call 'the properties' Rather, we have the word 'property', introduced by way of a varied repertory of ordinary and philosophical uses. The word has thereby become associated with a role in our commonsensical thought and in a variety of philosophical theories. To deserve the name of 'property' is to be suited to play the right theoretical role; or better, to be one of a class of entities which together are suited to play the right role collectively. But it is wrong to speak of *the* role associated with the word 'property', as if it were fully and uncontroversially settled. The conception is in considerable disarray. It comes in many versions, differing in a number of ways. (1986: 55)

He distinguishes a couple of conceptions:

Sometimes we conceive of properties as *abundant*, sometimes as *sparse*. The abundant properties may be as extrinsic, as gruesomely gerrymandered, as miscellaneously disjunctive, as you please. They pay no heed to the qualitative joints, but carve up things in every which way. Sharing them has nothing to do with similarity. ...There is one of them for any condition we could ever write down.... In fact, the properties are as abundant as the sets themselves, because

¹⁰See Lewis, 1983, 1984, and 1986: 50-65.

for any set whatever, there is the property of belonging to that set. ... The sparse properties are another story. Sharing them makes for qualitative similarity, they carve at the joints, they are intrinsic, they are highly specific, the sets of their instances are *ipso facto* not entirely miscellaneous, there are only just enough of them to characterize things completely and without redundancy. (59-60)

Finally, he offers a methodological suggestion:

I would not recommend that we enter into debate over whether the properties really are abundant or whether they really are sparse. We needn't choose up sides. Rather we should acknowledge that we have both conceptions, and an adequate account of what there is ought to accommodate both. (60)

Lewis here accepts a form of pluralism. He thinks our talk of properties is in “considerable disarray”. The properties are to serve in many different roles. There is no guarantee that the roles are coherently combinable into one complete conception of properties. I'll argue that his pluralism allows for a resolution of the problem of entailment on behalf of the error theorist.

3.5.2 Terminological Foundations for Pluralism

Pluralism, as I have characterized it thus far, is incredibly vague. I have described it more as a sort of stance that a philosopher may adopt with respect to our talk of properties than as a real view with a genuine subject matter. But it is a statable doctrine, provided we have an exposition of the correct terminology. In order properly to understand the view, I'll explain the notion of a property role, the notion of a conception of properties, as well as the notion of accepting a conception of properties.

A *role* for a kind is a condition that entities are expected to satisfy in order to count as members of the kind.¹¹ To take an example from the philosophy of language, there are at least two crucial roles that propositions are expected to satisfy: they are to serve as the

¹¹I here use 'kind' in a sort of vague, pre-analytic fashion. I'll speak as if “the properties” and “the propositions” delimit “kinds”. But there is another, post-analytic sense of 'kind', on which different conceptions of properties or propositions may correspond to radically different “kinds” of entity (for example properties on one conception may be best thought of as Aristotelian universals, while on another they may be best thought of as classes). It is important not to confuse these two uses of the word in what follows.

primary bearers of truth and falsity, and as the objects of our attitudes of belief, desire, and the like. A *property role*, then, is a condition that entities must satisfy in order to count as properties. It is important that roles, as I am conceiving them, are analytic of their subject matters: it's part of the concept of a proposition that propositions are the primary bearers of truth and falsity. It is also important to note that we should not expect any perfectly precise explication of these roles; at least, not at first. Property roles, both in ordinary discourse and in many philosophical theories, are sufficiently inchoate to evade precise characterization. Nonetheless, the various roles that are associated with a kind may not be satisfiable at once by any collection of entities from a single ontological category; they may be different enough so as to make this impossible. Once we start systematizing the property roles, we may find that some of them go together in a coherent fashion, while others do not.

A *conception* of a kind is a determinate collection of entities that jointly are fit to satisfy some subset (proper or improper) of the roles of the kind.¹² To extend our example from the philosophy of language, there is a conception of propositions as sets of possible worlds: relative to a certain kind of semantic framework, sets of possible worlds are fit to serve as the bearers of truth-value, as well as the objects of propositional attitudes. A *conception of properties* is then a precise collection of entities that jointly may satisfy some subset of the property roles. Unlike the property roles, a conception of properties is a result of philosophical systematization, so it is unacceptable for a conception to be inchoate to the same extent that the roles are. Intuitively, conceptions must be reasonably precise, whereas roles may be quite vague indeed.

Part of the precision that is imputed to a conception is manifested in *criteria of individuation*, which specify conditions of identity and difference for the entities of the conception.

¹²I do not intend to define Lewis's usage of 'conception' in the above passages. The intuitive difference between these two senses of 'conception' holds in virtue of the fact that Lewis uses the term vaguely to talk about a pre-analytic notion, whereas I use it to talk precisely about a post-analytic notion. This roughly corresponds to the difference between the two senses of 'kind' that I noted above. For a usage of 'conception' that is most similar to mine, see Bricker, 1996.

These criteria may be thought of as deriving from the vague roles that the conception is supposed to capture. Thus, on a certain possible worlds conception of proposition, propositions are identical when they are true at all the same possible worlds, and believed or disbelieved, to the same degree, by all the same individuals; otherwise, they are distinct. Different conceptions of properties will come prepackaged with different criteria of individuation, depending upon which roles the conception is meant to capture. It is crucial to note that, since roles are analytic of their kinds, the criteria of individuation are, as a result, also analytic: the criteria of individuation of a conception partially define it. A consequence of this is that different claims about different conceptions of properties do not comprise different theories about the same thing. They need not be in competition with one another; intuitively, there may be many equally acceptable conceptions of properties.

Finally, we may say that a philosopher *accepts a conception* of a kind when he accepts entities fit to serve in the roles of that conception, and so do not violate the criteria of individuation of the conception. If a philosopher accepts classes of possible worlds and they satisfy the criteria of individuation of a certain possible worlds conception of propositions, then he thereby accepts that conception of propositions. Likewise, a philosopher *accepts a conception of properties* when he accepts entities that together may fill the roles of the conception and that do not violate its criteria of individuation. Depending upon the categories and quantities of the entities that a philosopher accepts, he may accept none, one, or many conceptions of properties.

3.5.3 Some Roles for Abundant and Sparse Conceptions

In order further to clarify pluralism about properties and its rivals, I'll need to provide a brief summary of the property roles.¹³ Before I do so, it is instructive to reiterate that the roles need not be specified in a completely determinate manner; they need not be precise.

¹³For references, see Lewis, 1983 and 1986: 50-65; Swoyer, 2011; and Bricker, 1996.

In the above passages, Lewis conveniently divides conceptions of properties between sparse and abundant. The differences in the abundance of properties on various conceptions arises from differences in the roles the conceptions are expected to fill. There are two main property roles that naturally are filled by abundant conceptions of properties. On the first, there is a property for any class of entities. A conception that satisfies this role will come with ready-made criteria of individuation, as classes are identical just in case they have all the same members. Since the classes, as they are standardly conceived, are very abundant indeed, properties are likewise very abundant on this conception.

More important for our purposes is a second role, on which properties are expected to be the “meanings” of linguistic conditions that may be asserted to hold of individuals, e.g., ‘is red’, ‘is round’, and so on. The quantity of expressible conditions is large, so properties are also very abundant on a conception that satisfies this role. In addition, since there is not just one single notion of meaning, the criteria of individuation associated with such a conception will vary according to what notion of meaning is at use in specifying the role.¹⁴ In my opinion, different notions of meaning may be derived from different semantical frameworks, thus giving us different criteria for sameness of meaning, and so different criteria of individuation. For example, in an extensional Tarskian framework, abundant properties will be identical when they have the same extension (i.e., they apply to all the same actual objects). On a simple possible worlds framework, abundant properties will be identical when they have the same intension (i.e., they apply to all the same possible objects). In some frameworks, meanings may bear complex structure, and so sameness of application is insufficient to guarantee sameness of meaning. Some more complex criteria of individuation are required on such conceptions: properties will be identical when they coincide in application and have the same structure.¹⁵

¹⁴See, e.g., Lewis, 1980, on the many senses of ‘meaning’.

¹⁵See Lewis, 1986: 58-9, on how to construct structured meanings from unstructured ones.

At this point, it is instructive to note that ‘property’, as it occurs in SC2, seems to be used in this abundant sense. When a philosopher claims that, unless a predicate expresses a property, statements containing it don’t express propositions, he seems to be invoking a property role on which properties are to serve as meanings. Further reason to interpret SC2 in this way arises when we consider sparse conceptions of properties and their roles: properties sparsely conceived seem not to serve as meanings.

Roles that may be filled by sparse conceptions of properties tend to derive from metaphysical theories, or from other philosophical theories. They may be required to provide a basis for a complete qualitative specification of the world; to serve as part of a foundation for natural laws and facts about causation; to explain the existence of objective resemblances among individuals; or to provide objective constraints on the interpretation of language and thought.¹⁶ Since it is not obvious that there is exactly one conception of properties that is able to fill all these roles, it may turn out that there are multiple sparse conceptions of properties, each with its own criteria of individuation.

Although I’ll not try precisely to differentiate particular sparse conceptions here, it is crucial to note that the criteria of individuation of a sparse conception will be more strict than the criteria of an abundant conception. For example, on a standard sparse conception, the criteria of individuation will have to mention the notion of objective resemblance: properties F and G will be identical only if, whenever individuals x and y respectively have F and G, x and y objectively resemble each other in some respect. But there is no necessary connection to objective resemblance on an abundant conception: two objectively very different individuals might satisfy the very same disjunctive condition. It is natural to conclude from this that sparse properties may be put into one-one correspondence with some proper subset of abundant properties, and so the sparse properties are strictly less abundant. As a result of this difference between sparse and abundant conceptions, it is coherent to

¹⁶For a more thorough summary of these roles, as well as references, see Swoyer, 2011 and Bricker, 1996.

accept the existence of a certain abundant property while maintaining that it doesn't have any sparse correlate.

3.5.4 More Precise Characterizations of Pluralism and its Rivals

With these preliminaries in mind, we are in a position to distinguish pluralism from its rivals. First, there is the nihilistic philosopher, who accepts no conception of properties, sparse or abundant. According to this view, there is no collection of entities the members of which can satisfy any of the property roles. I mention this view only because it is in the space of possible views, and not because it is plausible. Second, there is the monistic philosopher who accepts only one abundant conception of property, and claims that this conception is able only to fill a single abundant role. According to this view, there simply aren't enough individuals of the right categories to fill any of the sparse conceptions of property; perhaps the criteria of individuation would have to be too vague or indeterminate. It is crucial to note that both of these philosophers – the nihilist and the monist – might acknowledge that there are the variety of property roles that I have suggested. It may be that they find these roles either too mysterious or incoherent to posit entities that might fill them.¹⁷

Opposed to both of these views is the monistic philosopher who accepts just one conception of properties, and claims that the one conception can fill all the property roles. On this view, there is a single collection of entities, whose criteria of individuation are determined, jointly, by all the property roles that I have suggested. In my opinion, such a view is implausible until such criteria of individuation can be spelled out in a more precise manner. The extreme diversity of the property roles provides defeasible reason to reject such a view.

Finally, there are views that qualify as genuine forms of pluralism. A weak pluralist will accept at least one conception that satisfies a sparse property role, and at least one conception that satisfies an abundant property role. This philosopher will claim that the conceptions he

¹⁷See, e.g., Quine, 1980 and 1948: 30-1.

accepts come with different criteria of individuation, so they are different conceptions of properties. A strong pluralist, on the other hand, will accept conceptions enough for every property role, sparse and abundant. Depending upon the criteria of individuation he ends up with, he may accept more than two conceptions of property. In any case, there will be one conception for each set of criteria. Thus, while the ontology of the conceptions may overlap, the conceptions are different, so long as their criteria are different.

I'll now register an opinion about these various kinds of nihilism, monism, and pluralism about properties. I believe that a philosopher ought to accept, for each coherent property role, a conception that fills that role. This is not to say that there must be a distinct conception for each role. But it is to say that an adequate theory of properties must accommodate every property role in some way.¹⁸ Therefore, I'll presuppose the strong version of pluralism in my resolution of the problem of entailment. This form of pluralism opens the door for a clean solution to the puzzle.

3.6 A Pluralist Solution to the Problem of Entailment

In this section, I'll distinguish the coherent versions of the metaphysical thesis from its incoherent versions. These distinctions make essential use of the distinction between sparse and abundant properties, though they are not formulated in terms of specific sparse or abundant conceptions. I'll show that, once we acknowledge the many versions of the metaphysical thesis, the problem of entailment dissolves. Finally, I'll argue that the error theorist ought to impose a certain ordering of epistemological priority on the metaphysical theses he accepts.

3.6.1 Four Versions of the Metaphysical Thesis

Recall the "ontological" and "instantiation" versions of the metaphysical thesis. According to the first, there are no moral properties; on the second, although there are moral

¹⁸Compare Lewis, 1986 and Bricker, 1996.

properties, none are ever instantiated. The main problem with these formulations is that they do not suggest which conceptions of properties are at issue. They could be understood as being about properties on certain abundant conceptions ('abundant properties'), or about properties on sparse conceptions ('sparse properties').

I have argued that, for different conceptions of properties, we get different interpretations of the metaphysical thesis. This is a result of the fact that different conceptions have different analytic criteria of individuation: these criteria, in a sense, "give the meaning" of 'property' on their respective conceptions. Since I have not provided us with precise criteria of individuation for the properties of various conceptions, I'll make the simplifying assumption that there is just one abundant conception and just one sparse conception. Then we get two versions of the ontological form of Thesis 1:

Thesis 1OA: The Abundant Ontological Metaphysical Thesis. There are no abundant moral properties; in specific, there are no abundant objective values.

Thesis 1OS: The Sparse Ontological Metaphysical Thesis. There are no sparse moral properties; in specific, there are no sparse objective values.

These theses invoke distinct conceptions of properties, and so, in effect, they presuppose different things about what roles "properties" are required to fill. The abundant version presupposes abundant roles, and the sparse version presupposes sparse roles.

In addition, while it is not blatantly incoherent for an error theorist to accept Thesis 1OS, Thesis 1OA is clearly not something an error theorist will like. On a standard abundant conception of properties, properties are to serve as the contents of predicate expressions: for any contentful predicate expression, there is the abundant property that is its content. Since our error theorist allows that basic moral predicates have content, he thereby accepts abundant moral properties. Importantly, any error theorist who accepts that moral statements have entailments will want to hold that moral predicates have content. So this error theorist will want to reject Thesis 1OA in favor of the claim that there are abundant moral properties. Thesis 1OS is another matter, since, as I have noted, it is coherent to accept a certain

abundant property but reject any sparse correlate of it. Indeed, I will argue below that our error theorist ought to accept Thesis 1OS.

There are also two versions of Thesis 1I, according to which moral properties are not instantiated. On the assumption that there are just one sparse and one abundant conception of properties, we get two versions of this thesis:

Thesis 1IA: The Abundant Instantiation Metaphysical Thesis. There are abundant moral properties, but they are not instantiated; in specific, although there are abundant objective values, they are not instantiated.

Thesis 1IS: The Sparse Instantiation Metaphysical Thesis. There are sparse moral properties, but they are not instantiated; in specific, although there are sparse objective values, they are not instantiated.

Like Theses 1OA and 1OS, these theses invoke different roles for properties. Since I believe that the typical error theorist will accept Thesis 1OS, it would be incoherent for him to accept Thesis 1IS. Instead, such an error theorist ought to be interpreted as accepting Thesis 1IA. If the error theorist claimed that abundant moral properties were instantiated, then it seems he would be committed to the truth of some moral statements. Error theorists clearly want to avoid this result.

This discussion suggests that an error theorist has some reason to accept two versions of the metaphysical thesis (Thesis 1OS and Thesis 1IA), as well as some reason to reject the other two versions (Thesis 1IS and Thesis 1OA).

3.6.2 The Solution

Just as there are four versions of the metaphysical thesis, given pluralism, there are two versions of the second semantic constraint:

SC2A: No Abundant Properties, No Propositions. If a predicate doesn't express an abundant property, then no statement containing it expresses a proposition.

SC2S: No Sparse Properties, No Propositions. If there is no sparse property corresponding to a predicate, then no statement containing the predicate expresses a proposition.

And just as Thesis 10A should be rejected out of hand by the clear-headed error theorist, SC2S should be rejected out of hand by any clear-headed philosopher. In general, from the fact that a certain predicate doesn't correspond to a sparse property, nothing follows about whether statements containing the predicate express propositions. This is clear, for example, if we suppose that the sparse conception involves the notion of objective resemblance. For any sufficiently disjunctive predicate, there are typically (possible) individuals that satisfy the predicate and that objectively resemble each other in no respect whatever. In such a case, there is no sparse property corresponding to the predicate, even though statements that contain the predicate will express propositions.

Perhaps it will be helpful to give an example, in order to explain this point. Consider the predicate: 'is either composed wholly of silly-putty or of aether'. First, note that there are actual individuals that satisfy the predicate: there are plenty of little balls of silly-putty, and so it's true that they are either composed wholly of silly-putty or of aether. In addition, statements containing the predicate do express propositions. Nonetheless, some things made from aether would seem not to bear any objective resemblance to things made from silly-putty. Intuitively, this is because aether is supposed to be a kind of substance that no material objects could be composed of, whereas things that are made from silly-putty must be material. So statements containing this predicate express propositions even while the predicate cannot be associated with any sparse property. Therefore, if the problem of entailment is even to get off the ground, it must invoke SC2A instead of SC2S.

But even if we assume that SC2A is the correct version of the second semantic constraint, the problem doesn't get very far from the floor. For there there seems not to be any way to show that Thesis 10S, Thesis 11A, and Thesis 2E* are jointly inconsistent.

Suppose first that we try to run the problem with Thesis 11A in mind. If Thesis 11A is true, then moral predicates do express abundant properties, though they aren't instantiated. So on this way of reinterpreting the problem, we cannot use SC2A to show that the error theorist is committed to the nonexistence of moral propositions. We could derive that result

only if our form of the error theory implied that there are no abundant moral properties. This is no way to show that moral statements do not have entailments.

Suppose, on the other hand, that we try to run the problem with Thesis 1OS in mind. On this supposition, although we can get the result that moral predicates don't correspond to sparse moral properties, we cannot use this in conjunction with SC2A to argue that the error theorist must reject moral propositions. Thesis 1OS and SC2A have different subject matters: the one is about sparse properties, and the other is about abundant ones. I therefore conclude that SC2A cannot be used to show that the error theorist's metaphysical thesis is incompatible with the entailment discourse thesis.

At least, there is no truth-preserving way to do this: if we equivocate on 'property', and thus ignore the distinctions that I have drawn in this chapter, then we end up with a *prima facie* valid proof that the error theory is incoherent. This is what occurs in the original statement of the problem. But we call the phenomenon of equivocation a 'fallacy' for good reason. We do not want to commit fallacies in any of our alleged derivations. Part of the reason why logicians invented logic was to avoid equivocation.

3.6.3 The Epistemological Ordering of the Metaphysical Theses

In this subsection, my main goal is to clarify what I have tried to show thus far about how to formulate the error theory. I take for granted that most error theories purport to show that there is some inconsistency or incoherence in our moral concepts. This provides a basis for three further ideas. First, that an error theorist ought to accept Thesis 1IA. Second, that an error theorist who accepts Thesis 1IA ought also to accept Thesis 1OS. Third, that an error theorist ought to accept at least one version of the discourse thesis.

All thoroughgoing moral error theorists purport to find some incoherence in our moral concepts. For Mackie, the incoherence lies either in an internal connection to motivation, or the objective prescriptivity of morality, or the idea that morality is brutally supervenient on the natural world. In the next chapter, I'll discuss Richard Joyce's view, on which the

incoherence is in the objective normativity of morality. Others purport to find further respects of incoherence in our moral concepts. In a defense of Mackie's view, Richard Garner writes,

It is hard to believe in objective prescriptivity because it is hard to make sense of a demand without a demander, and hard to find a place for demands or demanders apart from human interests and conventions. We know what it is for our friends, our jobs, and our projects to make demands on us, but we do not know what it is for *reality* to do so. A black hole swallows everything, but it demands nothing. (1990: 143)

I take this passage to express a basic sentiment of the thoroughgoing error theorist: morality conceptually requires of the world things that the world could not provide. Thus, I take it that such an error theorist will accept not just Thesis 1IA, but a stronger version of it, according to which abundant moral properties cannot, on pain of some kind of metaphysical incoherence, be instantiated.¹⁹

If this is the basic thought of the error theory, then such an error theorist ought to accept Thesis 1OS as well. For, if a certain abundant property is necessarily uninstantiated, then any sparse correlate of it would also necessarily be uninstantiated. In general, an abundant property with a sparse correlate is instantiated if, and only if, the sparse property is instantiated. Unfortunately, no necessarily uninstantiated property could play any of the sparse property roles.

Consider the four roles for sparse properties that I briefly outlined above. The first role has it that properties are to explain the internal relations of objective resemblance between individuals. Objective resemblance, as a matter of conceptual necessity, is a relation that bears only between possible individuals. Thus entities that fill this role must be exemplified

¹⁹Of course, there are weaker versions of the error theory in logical space, on which morality does not require anything metaphysically incoherent. On such a view, morality requires certain conditions to obtain in the world, and while these conditions could obtain, they do not in fact hold true (think of an error theory about witch discourse). I do not discuss an error theory of this kind in this work because I know of no philosopher who has defended it, and I know of no arguments for such a view.

in some possible worlds. But these sparse moral properties would be exemplified in no worlds whatever. So no sparse moral property could fill this role.

Suppose that a necessarily uninstantiated property was alleged to play the role of serving as a foundation for natural laws. Then the laws of nature would hold in virtue of properties that couldn't possibly be instantiated. I take it to be part of our concept of a law of nature, though, that the only sparse properties that figure in the laws at a world are ones that are instantiated at some world. So, necessarily uninstantiated properties couldn't satisfy this function, either.

Suppose that a necessarily uninstantiated moral property was expected to play the role of being part of a qualitative supervenience base. Unfortunately, the only properties that supervene on impossible properties are necessary properties. For all other properties, there can be a difference in their instantiation without a difference in the instantiation of impossible properties. Thus, it is hopeless to suppose that impossible properties can fill this role: no properties properly called 'qualitative' supervene on impossible properties.

Finally, consider the role of providing an objective constraint on the interpretation of language and thought. Intuitively, on this role, the properties are supposed to correspond with the eligible contents: the ones that are "objectively favored".²⁰ But it is implausible to suppose that impossible properties are eligible contents. A fundamental idea of interpretation is, very roughly, that there is a presumption in favor of interpreting speakers as uttering truths. To allow that we may easily refer to impossible properties would be unnecessarily to allow violations of this idea. Intuitively, impossible properties are about as ineligible as any of the most ineligible contents.

Therefore, Thesis 1IA requires Thesis 1OS. If abundant moral properties are necessarily uninstantiated, they can have no sparse correlates. So since an error theorist views moral concepts as incoherent, he must reject the existence of any sparse moral properties.

²⁰On the notion of eligibility, see Lewis, 1983: 45-55, 1984, and 1986:104-8.

Finally, because the error theorist ought to accept Thesis 1IA, he ought at least to accept one version of Thesis 2. First, he ought to accept a version of the discourse thesis that is formulated in terms of analytic entailment and abundant properties:

Thesis 2EA: The Abundant Entailment Discourse Thesis. Every moral judgment analytically entails that there are instantiated abundant moral properties.

If our error theorist accepts any plausible semantic framework at all, then he will find that he cannot reject Thesis 2EA, if he accepts Thesis 1IA. For example, on a simple possible worlds framework, roughly, an utterance of ‘x is good’ is true just in case x is in the extension of ‘good’ at the world of the utterance. This requires that ‘good’ has an intension that assigns it an extension at the world of the utterance, so ‘good’ expresses a property in this possible worlds framework. Other standard semantic frameworks, in order to be plausible, should assign semantic values to ‘good’. Thus ‘good’ will express a property relative to those frameworks.²¹

Importantly, a predicate can have an intension even if its intension is empty: it will express a property on the possible worlds framework, though not a very interesting one. Then to say that ‘x is good’ analytically entails the instantiation of moral properties, on this semantics, is to say that at any world where the proposition expressed by ‘x is good’ is true, there is something at that world that falls in the extension of ‘good’ at that world. Clearly, this is trivially satisfied on any view that says moral properties have empty extensions at every world: the proposition expressed by this statement will be true at no worlds whatever. Any other semantics that an error theorist might adopt should provide a similar explanation of the notion of entailment.

We have seen that our error theorist who accepts Thesis 1IA is naturally led to accept Thesis 2EA. I know of no argument to the effect that our error theorist ought to accept the following:

²¹Even a sophisticated expressivist semantics assigns semantic values to moral predicates, and so they express properties relative to the framework. See Gibbard, 2006 and 2003; and Schroeder, 2010 and 2008.

Thesis 2ES: The Sparse Entailment Discourse Thesis. Every moral judgment analytically entails that there are sparse moral properties.

It seems to me that this extra version of the discourse thesis could be affirmed by an error theorist, though it need not be. In the next chapter, I'll argue that Richard Joyce ought to be interpreted as accepting Thesis 2ES.

3.7 Conclusion

In this chapter, I have argued that the problem of entailment does not prove that the error theory is internally incoherent. My argument relied on a controversial view about how to theorize about properties. The main upshot of this view is that we should not expect 'property' to be univocal throughout philosophical discourse. As a result, it is plausible to expect that the problem of entailment contains an equivocation on 'property', and so we cannot trust that the problem shows any real incoherence in the error theory. In the next chapter, I'll apply the lessons of this and the preceding chapters to a case study: Richard Joyce's error theory.

CHAPTER 4

JOYCE'S ERROR THEORY

4.1 Introduction

The preceding chapters have set the stage for a case study: Richard Joyce's error theory. He takes his view to be a direct descendent of Mackie's, and so it is useful to view it as a development of the views expressed by Mackie in *EIRW*. Joyce's is the most developed version of the error theory in the contemporary literature, so it seems a good place to search for the most plausible form of error theory.

In this chapter, I begin by stating the disambiguations of Theses 1-3 that I believe Joyce accepts, or at least nearly accepts. He adopts versions of the theses that, I believe, jointly result in a coherent picture of our moral discourse and the kind of reality it requires. Then I extract two main arguments that Joyce gives in favor of his version of the error theory.

In section 3 I outline his first argument, which appears mainly in his books, *The Myth of Morality* (2001) and *The Evolution of Morality* (2006). Roughly, the argument is that, since morality requires the existence of objective reasons, but there are none, his version of the error theory is true.

In section 4 I criticize Joyce's argument: while both premises seem initially plausible, their plausibility doesn't result in a plausible argument, since it contains an equivocation.

Interestingly, in his more recent work (2011 and 2008), Joyce more-or-less revokes this main argument for his error theory. He replaces the argument with one that purposefully uses vaguer terminology. Roughly, the argument is that since morality requires the existence of "practical oomph", but there is no such thing, his version of the error theory is true. I explicate this argument in section 5.

Finally, in section 6, I criticize this final argument: because it involves such imprecise terminology, it cannot motivate an error theory, especially one as nuanced as Joyce's. In the concluding section, I try to draw a general lesson from my discussion of Joyce's view.

4.2 Some Joycean Theses

With respect to Thesis 1, the metaphysical element of Joyce's view, I interpret him as accepting the existence of abundant moral properties, but rejecting the existence of sparse moral properties. Such sparse moral properties would be defectively objective, in his view. Although he accepts the existence of abundant moral properties, he obviously doesn't believe they are instantiated. I believe he even rejects the idea that they *could* be instantiated. As for Thesis 2, the discourse thesis, I take Joyce as agnostic between a version that uses the notion of logical presupposition and one that is phrased in terms of analytic entailment. But this undecidedness on Joyce's part need not keep us from evaluating his discourse thesis, since logical presupposition and entailment are relevantly similar kinds of phenomena.

4.2.1 Joyce's Metaphysical Theses

I understand Joyce as accepting a formulation of the metaphysical thesis that involves abundant moral properties:

Thesis 1IA: The Abundant Instantiation Metaphysical Thesis. There are abundant moral properties, but they are not instantiated; in specific, although there are abundant objective values, they are not instantiated.

This version of the metaphysical thesis, as we have noted, implies that there are abundantly conceived moral properties. I will assume that Joyce believes that these abundant moral properties are necessarily uninstantiated: they would require the world to be a way that it cannot be.

The evidence that Joyce accepts these entities is derived from his cognitivism and his willingness to assert that no one is ever morally obligated to do anything.¹ If a philosopher sincerely asserts a sentence that involves a certain predicate, and he does not offer some nonstandard semantical explanation of that predicate, then there is a presumption in favor of interpreting him as accepting an abundant property corresponding to the predicate. Joyce satisfies these conditions, and so I must take him as affirming Thesis 1IA, so long as there is no reason to interpret him otherwise.

As I see it, Joyce also accepts a formulation of the metaphysical thesis that involves sparse moral properties:

Thesis 1OS: The Sparse Ontological Metaphysical Thesis. There are no sparse moral properties; in specific, there are no sparse objective values.

As I argued at the end of chapter 3, Thesis 1OS can be thought derivative of Thesis 1IA. The error theorist who accepts Thesis 1IA will have a quick argument for the non-existence of sparse objective values: if there were such sparse properties, then some abundant objective values would be instantiated; but since that's impossible, so is the existence of sparse objective values.

The evidence that Joyce accepts Thesis 1OS derives from his willingness to assert that there are no objective reasons.² If we are to interpret this claim charitably, then we must interpret it as the denial of an existence claim. Since, presumably, a reason would be a relational property, Joyce therefore denies the existence of such properties. But since it would be incoherent for him to accept Thesis 1IA while also denying the existence of abundant moral properties, I cannot interpret him as rejecting the existence of abundant objective reasons. Thus the most charitable way to interpret his denial of objective reasons is as the denial of the existence of certain sparse moral properties.

¹See, e.g., Joyce, 2001: 77.

²See, e.g., Joyce, 2001: 134.

It is instructive to note that Joyce claims officially to be agnostic about properties.³ One may object, on these grounds, to my interpretation of Joyce's comments: if he explicitly says he's agnostic, then we shouldn't interpret him otherwise. I reject the thought behind this objection, since if it were correct, I could not understand what Joyce is to mean when he claims that no one could ever be obligated to do anything, and that there could be no objective reasons. As I interpret these claims, the first is true just in case there are abundant objective values that cannot be instantiated, and the second is true just if there are no sparse objective values. I can only understand Joyce's agnosticism about properties as agnosticism about the nature of properties on various conceptions. I cannot understand it to be a general agnosticism about the existence of properties.

4.2.2 Joyce's Discourse and Error Theses

Joyce is somewhat undecided about whether the discourse thesis should involve analytic entailment or logical presupposition.⁴ Thus, as I understand him, he accepts either Thesis 2FA or Thesis 2EA:

Thesis 2FA: The Abundant Fregean Discourse Thesis. A Fregean presupposition of each moral judgment is that there are instantiated abundant moral properties.

Thesis 2EA: The Abundant Entailment Discourse Thesis. Every moral judgment analytically entails that there are instantiated abundant moral properties.

However, since the differences between Thesis 2FA and Thesis 2EA are minute, I'll assume that the latter is Joyce's primary discourse thesis. The truth of none of the content of this chapter depends on which of these principles Joyce really accepts.

It is important to note that Joyce also accepts a version of the discourse thesis that involves sparse objective values. Thus he accepts either the following thesis, or its Fregean counterpart:

³Joyce, 2001: 9; 2007a

⁴Joyce, 2007a,b; 2001.

Thesis 2ES: The Sparse Entailment Discourse Thesis. Every moral judgment analytically entails that there are sparse moral properties.

The reason why we should interpret Joyce in this way is that he believes our moral judgments about wrongdoing and obligation entail the existence of “objective reasons”. I take it that these reasons would be sparse moral properties.

Joyce’s metaphysical and discourse theses commit him to the following error thesis (or its Fregean counterpart):

Thesis 3EAS: The Both-Abundant-and-Sparse Entailment Error Thesis. Since abundant objective values cannot be instantiated (and thus there are no sparse objective values), but our moral judgments analytically entail that they are instantiated (and that sparse objective values exist), all our moral judgments are false.

If Joyce accepts Theses 1IA, 1OS, 2EA, and 2ES, then he must accept Thesis 3EAS. Indeed, an even stronger conclusion follows from those prior theses, since he thinks moral properties are necessarily defective. That is, as a result of his acceptance of Thesis 1IA, he ought to believe that none of our moral judgments *could* be true.

A question that arises, however, is whether Joyce accepts such a global error theory as is suggested by Thesis 3EAS. Does he really believe that all positive basic moral judgments are false? Or should we interpret his error thesis as restricted in some way to a subset of the basic moral judgments? There is some *prima facie* evidence that he need not be committed to a global version of the error thesis. For example, when he summarizes his argument for his error theory, he writes,

I consider the argument for a moral error theory now to be complete. In short, when we say that a person *morally* ought to act in a certain manner, we imply something about what she would have reason to do regardless of her desires and interests, regardless of whether she cares about her victim, and regardless of whether she can be sure of avoiding any penalties. And yet after careful investigation we have found no defensible grounds for thinking such reasons exist. (2001: 134)

The conclusion of *this* argument, clearly, cannot be that every moral judgment is false. Instead, the proper conclusion is that each moral judgment to the effect that someone ought

to do something is false. This would be a much more localized error thesis. It would seem to leave alone all judgments involving moral goodness and badness, judgments involving the concepts of virtue and vice, as well as judgments involving notions of justice, fairness, and desert.

The *prima facie* evidence is defeated, however, by Joyce's claim that judgments involving moral obligation and wrongness are somehow foundational to our moral discourse. He writes,

[My] arguments have primarily targeted deontological notions like *obligation* and *prohibition*. One might object that even if these arguments were faultless, all they show is that a certain proper subset of our moral discourse is flawed, but there is a rich and robust moral language that remains untouched. However, it is my contention that moral concepts are to a large degree holistically connected, such that a persuasive attack on categorical imperatives will, one way or another, count as a persuasive attack on a great deal more besides. (2001: 175; cf. 2006: 61)

He goes on to suggest that aretaic, axiological, and justice-related basic moral judgments will all be false as a result of these sorts of "holistic connections". In my opinion, Joyce's suggestion here is not plausible. Nonetheless, I'll assume that it is correct, since what I have to say about his arguments is orthogonal to the question of whether they rationalize such a globalized error thesis. That is, I'll suppose that *if* his arguments were sound, they would result in a global error thesis. I merely note in passing that, in my opinion, I am supposing a big 'if': nothing that Joyce has said persuasively shows that deontological notions are at the foundation of our moral conceptual apparatus.

4.3 Joyce's First Argument

In this section, I concentrate on concisely formulating Joyce's main argument for his error thesis. I begin, in subsection 3.1, by citing Joyce's way of formalizing his argument, and I object to it on the basis that it is invalid. Nonetheless, I believe there is another way of putting Joyce's point that results in a *prima facie* valid argument for the correct conclusion.

In subsection 3.2, I explain Joyce's rationales for this cleaner version of his argument. But, in the next section, I'll argue that his argument is unsound, due to an equivocation on a crucial phrase.

4.3.1 Formulation of the First Argument

The following is a direct quotation of Joyce's official presentation of his argument in *The Myth of Morality*:

1. If x morally ought to ϕ , then x ought to ϕ regardless of what his desires and interests are.
2. If x morally ought to ϕ , then x has a reason for ϕ ing.
3. Therefore, if x morally ought to ϕ , then x can have a reason for ϕ ing regardless of what his desires and interests are.
4. But there is no sense to be made of such reasons.
5. Therefore, x is never under a moral obligation. (2001: 77)

Without commenting on the meaning of the various technical terms that appear in this argument, we can see that it is not valid. Although I am willing to accept the inference from (1) and (2) to (3), the second inference of the argument is problematic.

Let us charitably suppose that line (4) asserts that there is no possible analysis, or account, or explanation, of the kinds of reasons Joyce rejects. I.e., this is what it means to say that "we can make no sense of" some alleged phenomenon. Unfortunately, from the fact that we cannot make sense of an alleged phenomenon, it does not follow that the phenomenon never occurs. For example, although we might be unable to make sense of certain microphysical quantum phenomena (there is no deeper explanation of their nature), it does not follow that the phenomena never occur. At most, all that follows is that we are unable to formulate a plausible scientific explanation of why they occur, and what are their natures. Thus the inference from (3) and (4) to (5) does not preserve truth.

I take this to be a relatively uncontroversial point, and I therefore do not think it is damning to Joyce's argument. I mention the point only to set it aside: Joyce's argument for his error theory does not turn on a triviality such as this. He has presented the argument in

other work in such a way that it does not rely on this fallacious sort of inference.⁵ There is a good reinterpretation of the argument that is more compelling than the above interpretation. Before I present my official formulation of Joyce's first argument, however, there are some important preliminaries to explain.

First, it is important to clarify the conclusion of the argument. Which of Joyce's theses is it to be? I will understand Joyce to be providing an argument for his version of the error thesis, Thesis 3EAS. Since he explicitly affirms that this is his argument for his error theory, we may as well have the conclusion of the argument be the official formulation of the error thesis that I have attributed to him.⁶ Second, it is important to note that because Thesis 3EAS presupposes the truth of a metaphysical thesis and a discourse thesis, it cannot be established without first establishing Theses 1IA and 2EA. I therefore understand these prior theses as premises in Joyce's argument. As I will formulate the argument, Thesis 1IA corresponds roughly with line (4) of the above quote, and Thesis 2EA corresponds to line (3). Third, since I believe that Joyce accepts Thesis 1OS if he accepts Thesis 1IA, I will assume that these two theses stand or fall together in the context of this argument. The plausibility of one of them transfers to the other. Therefore, in the context of this argument, I'll act as if they are equivalent.

Finally, a proper understanding of Joyce's argument requires a few bits of technical terminology. Let us say that a reason to act is *categorical* if and only if it does not depend upon the psychological state of its agent.^{7,8} This definition also takes for granted that every reason to act is a reason for someone to perform a certain action. Call a reason *non-institutional* if

⁵E.g., Joyce, 2011: 523.

⁶See, e.g., Joyce, 2001: 134.

⁷See, e.g., Joyce, 2011: 522-4; 2006: Chapters 2 and 6; and 2001: Chapter 2.

⁸I note that, in Joyce's terminology, reasons are associated with imperatives, and it is imperatives that are categorical, whereas reasons are "desire-transcendent" (see Joyce, 2011; 2006). I gloss over this detail because I can find no reason to admit more technical terminology into my presentation of the argument: there would be no gain of content, and there would be an increase in complexity.

and only if it does not depend on the existence or decrees of any institution.⁹ I will sometimes use the phrase ‘objective reason’ to mean the same as ‘categorical, non-institutional reason’, simply to eliminate unnecessary words from my presentation. Finally, following Joyce, I’ll let the notion of an *institution* be relatively vague. We can get an intuitive grasp on it as follows. There is an institution that created the rules of etiquette, one that created the rules of chess, and one that created the rules of Roman gladiatorial combat. Perhaps these are, respectively, the elite of some society, the World Chess Federation, and the Roman emperor.¹⁰

We are now in a position to state a cleaner version of Joyce’s argument:

P1. For any agent x and action ϕ , if x morally ought to do ϕ , then x has an objective reason to do ϕ .¹¹

P2. For any agent x and action ϕ , x does not have an objective reason to do ϕ .¹²

C. Therefore, for any agent x and action ϕ , x is not morally obligated to do ϕ .

The argument is valid, and I am without a doubt that it is Joyce’s argument. Premise (P1) is Joyce’s more precise version of Thesis 2EA, and (P2) is Joyce’s more precise version of Thesis 1IA. We therefore must consider what Joyce has to say in explanation of the premises, as well as how (C) relates to Joyce’s ultimate conclusion, which I have presumed is Thesis 3EAS.

4.3.2 Explanation of the First Argument

I have already noted that Joyce is aware that (C) is not identical with the global error thesis that he would like to be correct. In order to derive such a conclusion, a further premise is required: a premise to the effect that if (C) is true, then so is Thesis 3EAS. Joyce accepts

⁹See, e.g., Joyce, 2011: 522-4; 2006: Chapters 2 and 6; and 2001: 37.

¹⁰See, e.g., Joyce, 2001: 35-7; cf. Mackie, 1977: 80-1.

¹¹See, e.g., 2001: 42; 2011: 523-4.

¹²See, e.g., 2001: 68-9; 2011: 526; cf. Mackie, 1977: 77-9.

such a premise, and I have supposed that it is true, even though I have my doubts. I continue to suppose that it is true, and my criticism of Joyce's argument will be orthogonal to this supposition.

An important feature of the premises of the argument is that Joyce believes each of them to be analytic. I.e., Joyce believes that it is part of the concept of *moral obligation* that whenever someone is obligated to do something, he has a special reason to act; and it is part of our concept of a *reason to act* that no one ever has one of these special reasons. This is a feature that Joyce's argument does not obviously share with Mackie's original arguments for the error theory. As I explicated Mackie's arguments, each was supposed to involve a "non-substantive" and a "substantive" premise. The non-substantive premise was supposed to be a claim merely about our moral concepts. The substantive premise was supposed to be a claim about what the world is really like; it was not supposed to be a conceptual claim.

As Joyce would have it, though, both parts of the argument are claims about our moral concepts. On his view, then, both premises may be prefixed with some such operator as, 'it is conceptually necessary that'. As a result of this fact, Joyce's premises seem to imply that there is some kind of an inconsistency between our concepts of *moral obligation* and *reason to act*. I point this out not as a complaint, but rather as a preliminary to a proper understanding the sorts of rationales we might expect for his premises. We should not expect Joyce to provide some kind of empirically-based rationale for them. Joyce's rationales are supposed to be *a priori*; they only involve appeals to intuitions about possible cases.

With this in mind, we are now in a position to consider the rationales that Joyce provides for his premises. I begin with the first premise, according to which moral obligation requires objective reasons.¹³ Joyce asks us to imagine a circumstance in which an individual has done something morally wrong by our own lights. Joyce's claim is that, provided that we

¹³Compare Garner (1990), who writes, "It is the peculiar combination of objectivity and prescriptivity, rather than any intrinsic motivational power, that makes moral facts and properties queer..." (143).

are willing to say that the individual has done something wrong, we must go on to admit that he has an objective reason not to perform that action. It would be incoherent for us to say that he has no objective reason, provided that we were willing to say that he morally ought not to perform the action.

That is the rationale in the abstract. It may be helpful to consider Joyce's favorite concrete example. He writes,

[Consider] Plato's example of Gyges with his ring of invisibility, who sought to satisfy his basest passions by stealing and raping at will. Of course, there are all sorts of prudential reasons that one might give a person as to why he shouldn't act like this: he'll end up feeling disconnected from his community, people won't truly love him, he'll feel guilty, he'll get bored, etc., etc. But we are allowed to stipulate the example such that none of these speak to Gyges' true ends – he is simply depraved and enjoys the fruits of his degenerate behaviour. Because our moral framework is categorical we can carry on legitimately saying 'Gyges you ought not to do that!' But if our utterances are merely a verbal output that has been validated by an institution of our own creation, then it all begins to sound rather shrill. We can picture Gyges pausing in some sadistic undertaking, acknowledging to us that, yes, we are speaking legitimately when we assert that he morally ought not to act this way, and then shrugging and carrying on. And if our categorical imperatives are nothing more than an institutionally-backed way of speaking, why should he do otherwise? (2011: 524)

The question at the end of this passage is just a bit of rhetoric. Joyce believes that if there are no categorical, non-institutional reasons for Gyges not to do what he's doing, then Gyges can legitimately ignore our requests. But since, intuitively, Gyges could not ignore us like this, we must be committed to an objective reason for Gyges to stop. Thus Joyce continues,

In [*The Myth of Morality*] I tried to pump the intuition that our unease with this scenario leaves us groping for *non-institutional* categorical...reasons. (2011: 524)

Thus Joyce infers that (P1) is true: moral obligations analytically entail the existence of objective reasons.¹⁴

¹⁴Compare Joyce, 2008: 257; 2001: 32.

Joyce's second premise says that there are no categorical, non-institutional reasons to act. What can be said in favor of this claim? Joyce accepts a general analysis of the concept of a *reason to act* that makes reasons dependent upon psychology or on the existence and decrees of institutions. If such a general analysis is correct, then there cannot be any categorical, non-institutional reasons to act. Thus Joyce's evidence for this claim, if he has any, is parasitic on whatever may be said in support of his general account of reasons to act.

Although nowhere in his work does Joyce clearly formulate some particular analysis of reasons, I can confidently assert that he accepts what I have called 'Humeanism about reasons':¹⁵

Thesis 5: Humeanism about Reasons. Whenever a person has a reason to do something, there are, or he believes there are, some desires, interests, or ends that would be promoted by his doing it.¹⁶

Of course, Thesis 5 is not an analysis of the notion of a reason to act. In order to count as an analysis, it would have to state necessary and sufficient conditions for the application of the concept. Nonetheless, I'll treat it as a partial analysis of the notion.

As I see it, then, Joyce's rationalization of (P2) involves an appeal to Humeanism about reasons. If we suppose, toward a *reductio*, that someone has an objective reason to do something, then we are committed to a possible counterexample to Thesis 5. For example, if we accept that Gyges has some categorical, non-institutional reason to stop stealing, we accept that he has a reason to stop. But since this reason won't depend upon Gyges' psychology (it is a categorical reason), and it won't depend on any institution (it is non-institutional), it is possible that no desires, interests, or ends would be promoted by his stopping. Since Thesis 5 entails that such a situation is impossible, our initial *reductio* supposition must be incoherent.

¹⁵See, e.g., Joyce, 2001: Chapters 2 and 3, which would most likely contain a formulation of his analysis of reasons.

¹⁶I believe that Joyce accepts a stronger version of Humeanism about reasons than does Mackie. On this stronger version, when someone has a reason to do something, *he has*, or he believes there are, some desires, interests, or ends *of his* that would be promoted by his doing it.

4.4 Criticism of Joyce's First Argument

In this section, I begin by suggesting two ways to criticize Joyce's first argument. According to the first, (P1) is false; and the second criticism has it that (P2) is false. Although I think these routes press Joyce on the correct issues, I do not think they get to the heart of the problem with Joyce's first argument. In subsection 4.2, I state what I take to be a devastating worry for Joyce's argument.

4.4.1 Two Incomplete Objections

If you assume that there is no equivocation in Joyce's argument, and you do not have an interest in quibbling about whether his rationales really show what he wants them to show, then there are two *prima facie* promising ways to criticize the argument.

First, you might reject the intuition behind the rationale for the first premise. I.e., you might have no intuition that Gyges has a reason not to rape or steal, simply *because* none of his desires would be promoted by doing so. Second, you might reject Joyce's (partial) analysis of reasons. I.e., perhaps it is unmotivated as a general analysis of the notion of a reason, simply *because* it is unable to account for Gyges' reason to stop. I'll comment on both of these routes before I suggest, in subsection 4.2, what I take to be a devastating objection to the argument. Both will play a crucial role in my objection.¹⁷

The first kind of criticism can be motivated by imagining a philosopher who, prior to considering Joyce's argument, has been convinced by some general Humean analysis of reasons to act. Such a philosopher, as I am imagining him, will engage with Joyce's argument as follows:

“When I consider the case of Gyges, I *am* inclined to say that he *morally ought not* to steal. How could I not have this intuition about these detestable actions?! Nevertheless, since all reasons to act must be dependent either on psychological or institutional facts, I have *no* inclination to accept that Gyges

¹⁷Shafer-Landau (2005) rejects both premises of Joyce argument, though his attack on Joyce is quite different from mine. It is hard to see how Shafer-Landau's criticisms would be compelling against Joyce (see below). I believe that my criticisms would be compelling, and so I prefer mine to his.

has a reason not to steal – so long as none of his desires would be furthered by doing so, and he doesn't accept any institutional rules against stealing and raping. Moral obligation simply doesn't have this kind of internal connection to reasons. I therefore deny Joyce's first premise."

Such a philosopher is able to imagine the case of Gyges, but is not bothered by Gyges' ignoring our requests to stop raping and stealing: it's perfectly reasonable for Gyges to do so. In my opinion, there is no obvious *prima facie* conceptual incoherence in such a view. Moreover, it is unclear what Joyce could say in reply: if a philosopher doesn't have the intuition, it's hard to get him to accept it without a certain amount of browbeating. Call this the *Reject the Intuition* response.

A second kind of objection targets (P2). This objection can be motivated by imagining the philosopher who, prior to considering Joyce's argument, has been convinced that moral obligations always entail reasons to act. This philosopher will engage with Joyce's argument as follows:

"It's intuitively obvious that Gyges morally ought not to steal. He therefore has a reason not to steal: that's just part of what we mean when we say that he has such an obligation. But I have no inclination to accept a Humean analysis of reasons – at best, any Humean set of necessary and sufficient conditions for reasons to act should be considered an *incomplete* analysis. It only tells us about those reasons that depend on psychology or institutional decrees. It doesn't tell us about the kinds of reasons that are grounded in moral obligation. I therefore reject Joyce's second premise: as an analysis of all reasons to act, it is false."

This philosopher rejects the Humean analysis of reasons because it isn't able to account for the alleged reason that Gyges has in Joyce's example. Of course, it seems cheap in the context of Joyce's argument to perform such a G.E. Moore shift.¹⁸ Joyce could accuse such a philosopher of "begging the question". But it is crucial to note, at least, that cases like

¹⁸Suppose that someone presents a *modus ponens* inference to you: $p \rightarrow q, p \vdash q$. Then one way to perform the G.E. Moore shift is to "tollens their ponens": $p \rightarrow q, \neg q \vdash \neg p$. Obviously, another way to perform the G.E. Moore shift is to "ponens someones tollens".

Gyges are commonly thought to be the main problem for Humean analyses of reasons.¹⁹ Call this the *Reject the Analysis* response.

I take it that both of these objections would lead to a sort of deadlock in a debate with Joyce. If someone took the first option, Joyce could have no alternative but to agree to disagree. You can't force someone to have an intuition that he doesn't have. If someone took the second option, Joyce could accuse him of begging the question. But that philosopher could equally accuse Joyce of having the burden of proof, since the argument in question is Joyce's.

My reaction to this is that neither of these routes, in the end, seems promising. I have no interest in playing burden-of-proof-tennis, and I am hopeful that there is a way to avoid deadlock. The way forward, in my opinion, is to argue that Joyce's argument equivocates on 'has a reason'. Once we distinguish two natural senses of the phrase, we will find that on any uniform interpretation of Joyce's argument, one of the premises is false.

4.4.2 A Complete Objection

I am not the first to point out an ambiguity in Joyce's argument. In the context of discussing possible versions of the moral error theory, David Lewis writes,

[M]orality has been accused of presupposing quite a variety of errors. ... Perhaps there is the error of supposing that the dictates of morality give us reasons for acting regardless of our actual or potential desires. [Here Lewis cites Joyce, 2001.] But again I am not sure that someone who sees fit to use 'reason' in such an expansive way has exceeded his linguistic rights. I do agree, however, that it is an error to suppose that the dictates of morality give us reasons that are reasons in just the same non-disjunctive sense in which reasons based on serving our actual and potential desires are reasons, so that moral and desire-based reasons can compete for the status of reasons *simpliciter*. (Lewis, 2005: 316; cf. Balaguer, 2011: 382)

In this passage, Lewis seems to accept that there may be a sense of 'reason' on which Joyce's first premise is true. He also seems to accept that there is a sense of 'reason' that makes

¹⁹See, e.g., Schroeder, 2007: Chapter 6.

Joyce's second premise is true. What Lewis would deny is that there is one sense of 'reason' on which both of Joyce's premises are true. That is what Lewis implies when he suggests that moral and desire-based reasons should not be supposed to be in competition for the status of reasons *simpliciter*.

Although I agree with Lewis in the basic form of this objection, I would not put it in quite the way that he does.²⁰ In my opinion, we need not say that Joyce's equivocation occurs in the phrase 'reason', but rather, in the phrase 'has a reason'. That is, I believe that the crucial ambiguity occurs in 'has', instead of in 'reason'. In order to see the ambiguity, first consider the following sentences:

1. My body has some physical parts: my hands, legs, and so on.
2. There are some physical things that are parts of my body: my hands, legs, and so on.

There is a usage of 'has' (and its syntactic variants) where all that is required for you to "have" something is for that thing to exist. In our mereological discourse, for example, we readily paraphrase (1) with (2): there is no significant difference in content between these sentences. This holds for a large portion of our discourse in other areas, too: to say that I have parents just is to say that there are some parents of me; to say that I have a bike just is to say that there is a bike that I own; and so on.

In some of our discourse involving 'has', however, there is an extra requirement. In order to "have" something in this sense, it not only must exist: you must also bear some intentional relation to it. I.e., you must either accept it, or you must at least be aware of its existence, or maybe you must have some kind of psychological access to it. As an illustration of this stronger sense of 'has', consider an example from epistemology. Suppose that I see Fred in his office, but after a while I temporarily forget this fact: it has skipped my mind. I could perhaps recall that I saw him earlier if I were prompted and I thought about

²⁰But I do accept that 'reason to act' is not univocal in our discourse. Our usage has not determined a single unique sense of the term.

it for a bit, though I may initially deny this fact. The following sentences seem true when evaluated relative to such a context:

3. There is some reason for me to believe that Fred is in his office.
4. I don't have reason to believe that Fred is in his office.

On the most natural reading of (3), it should come out true. Since I saw Fred in his office earlier, there is some reason for me to believe that he's in his office. In general, our perceivings give us reasons to believe what they are perceivings of, at least so long as we haven't taken hallucinogens of some kind. Nonetheless, on at least one natural reading of (4), it should be true. Since, in this context, I do not presently accept that I saw Fred in his office earlier, I don't presently *have* reason to believe that he's in his office. Thus it seems that in these epistemic contexts, 'has' does not have the same meaning as does 'has' in mereological contexts. If it did, then (3) and (4) couldn't both be true of me in this context.

Before I move on, it is important to note that there are at least two more-or-less natural readings of (4). First, there is the reading that I relied on above to make (4) come out true. We can paraphrase it roughly as follows:

- 4'. Although there is a reason for me to believe that Fred is in his office, I *don't presently accept it*, and it is *not accessible by me* right now; so I don't *have it*.

This corresponds to the stronger sense of 'has', where in order to have something you must bear some sort of intentional relation to the object. But there is also a reading that corresponds to the weaker sense that we readily invoke in mereological contexts:

- 4''. There is no reason *for me* to believe that Fred is in his office; so I don't *have* such a reason.

On this reading, in order to have a reason to believe something, it merely needs to be a reason *for me*. This parallels the mereological example, where in order to have a part, it merely needs to be a part *of me*. I don't have to bear any crucial intentional relations of acceptance of it or accessibility to it.

How should this ambiguity in ‘has’ affect our evaluation of Joyce’s argument? I believe that in the argument, both premises come out true only if we do not interpret ‘has’ uniformly throughout. However, on such an understanding, the argument commits the fallacy of equivocation, and so it is invalid. Let me explain.

For simplicity, let us abstract away from the technical terms in Joyce’s argument, and instantiate his premises to the example involving Gyges. Thus, in the following sentences, let ‘reason’ mean the same as ‘categorical, non-institutional reason’. Joyce’s argument then looks like this:

5. If Gyges morally ought not to steal, then he has a reason not to steal.
6. But Gyges does not have a reason not to steal.
7. Therefore, it is not the case that Gyges morally ought not to steal.

If we apply the above lesson about our talk of “having” to (5)-(6), then we find that there are two uniform ways to interpret these premises.

On the first way, we may interpret ‘have’ with its weaker meaning:

- 5’. If Gyges morally ought not to steal, then there is a reason for him not to steal.
- 6’. But there is no reason for Gyges not to steal.

On this understanding of the premises of Joyce’s argument, I think Joyce would be unable to provide a coherent reply to an analogue of the Reject the Analysis response. That is, the Humean analysis of reasons that Joyce relies on is not plausible (in the least) when it is understood as an analysis of the reasons *there are* for people to perform actions. It doesn’t even come close to capturing all of our talk of such reasons, since a large part of that talk is grounded in our considered convictions involving moral obligation.²¹

²¹This objection mirrors the classical objection to Humeanism about reasons to the effect that it implies the existence of “too few reasons”. In my opinion, no attempt in the literature to reply to this objection on behalf of the Humean has been successful. See Schroeder, 2007: Chapter 6, which is titled ‘Too Few Reasons’. The only possibility available to the Humean, in my opinion, is to bite the bullet. But, unless some explanation of our intuition is provided, that is just to acknowledge that the Humean analysis simply isn’t an adequate account of our ordinary concept of a reason to act. (And I remain puzzled at the dogged persistence some still have in accepting some kind of Humean analysis of reasons, in spite of its inability to explain away our

Notice that this criticism of this disambiguation of Joyce's argument is very similar to what Lewis writes in the above passage. According to Lewis, someone who uses 'reason' in such a way that there may be reasons that do not promote our desires "hasn't exceeded his linguistic rights". I accept this claim, but also a somewhat stronger claim: someone who refuses to use 'reason' in such a way that it is impossible for there to be reasons that do not promote our desires has ignored his linguistic obligations.

The second way of interpreting Joyce's argument injects our stronger meaning into 'has':

5". If Gyges morally ought not to steal, then there is a reason for him not to steal, and he is aware of and accepts this reason.

6". But there is no reason for Gyges not to steal that Gyges is aware of and accepts.

On this alternative understanding of the argument, the second premise, (6"), seems obviously true. Gyges is a despicable human being, and although he may be aware that plenty of people disapprove of raping and stealing, he simply doesn't care. He doesn't accept those reasons, and so in our stronger sense of 'have', he doesn't have them. Unfortunately for Joyce, however, there seems not to be any way for him to reply to the analogue of the Reject the Intuition response in this case. That is, Joyce's first premise, (5"), is now obviously false. Although it is intuitive that Gyges morally ought not to rape and steal, Gyges accepts no reason not to rape and steal.

So I accuse Joyce of equivocating. When I consider (P1) and (P2) in their non-disambiguated forms, both seem to have some intuitive appeal. But once I recognize the crucial ambiguity in the phrase 'has a reason', I realize that the initial appeal does not count for much at all, so long as we can provide no uniform disambiguation of the argument on which each of the premises seems true. All the intuitive appeal is explained by the clear plausibility of

intuitions. (Of course, I am willing to accept that there is some form of Humeanism that is able to provide a plausible account of some *subset* of our discourse involving "reasons to act"; but as a general account of reasons, Humeanism is hopeless.)

(5') and (6''), but there is no valid argument for Joyce's conclusion that arises out of these premises. I therefore reject this argument as unsound.

4.5 Joyce's Second Argument

I do not know whether Joyce would accept my criticism of his argument. In his more recent work, however, he seems to retract the argument. Joyce writes,

We want something that *really* binds Gyges, such that he is making a mistake *by his own standards*, even if he cannot be brought to recognize this. I have come to suspect that my inadequacy in articulating this idea is not because of any failure of imagination or eloquence. Rather, morality may be imbued with a deeply mysterious kind of force – a kind of primitive feeling of 'being bound by rules and ends' that resists explication. ... Perhaps Mackie and I fumble to dissect something that by its very nature cannot be brought into the light to be picked over by philosophical scrutiny. (2011: 524-5)

In this passage, I understand Joyce to accept that his argument is something of a failure. He suggests that it was a mistake to formulate it in terms of 'categorical, non-institutional reasons'.²² It seems, then, that he wants to replace his reference to reasons with reference to some "deeply mysterious kind of force". In *The Evolution of Morality*, Joyce coins a term for this force: 'practical oomph'. He writes,

I claim that moral normativity has a distinctive kind of practical "oomph". Despite appearances, I choose the word "oomph" carefully, since it is indeterminate, non-theoretical, and metaphorical – and thus, I maintain, does a decent job of capturing certain aspects of the phenomenology of moral judgments made by ordinary thinkers in everyday contexts. Ordinary thinkers probably have a thoroughly inchoate idea of what the "must-be-doneness" of moral rules consists in, but this is not to say that it is a peripheral or negotiable aspect of morality. On the contrary, I argue that it is of utmost importance. (2008: 257; cf. 2006: 57-64.)

As I interpret Joyce in these passages, although he retracts the argument that consists in (P1), (P2), and (C), he does not give up on his version of the error theory. A new argument for

²²See also 2008: 260.

Thesis 3EAS is supposed to arise out of Joyce's alleged insight about the practical oomph of morality.

Before I explicate what I take to be his second argument for the error theory, it is important to comment on this novel idea in Joyce's writings. It represents a significant departure from his earlier thinking. In *The Myth of Morality*, Joyce criticized Mackie for using a term ('objective value') that was "too blunt for a proper argument to be conducted".²³ Joyce's choice to use the allegedly more-precise terminology of 'reasons', instead of 'practical oomph' or 'objective value', was a product of the fact that he "was not content to conduct the argument in these mysterious terms".²⁴ Joyce's newer argument for the error theory explicitly uses the term 'practical oomph', and thus represents his relaxation about using mysterious terms in philosophical argumentation. Although I do not yet raise an objection to Joyce's new argument on the basis of this fact, I'll eventually appeal to it in my criticism of the argument.

What can be said to clarify this new term, 'practical oomph'? The above passage from Joyce suggests that he intends to use the term with whatever meaning it has in our public usage. This is an interesting thought, but I doubt the *OED* explicates a meaning of 'oomph' that is appropriate for Joyce's purposes: there it is defined as "the quality of being exciting, energetic, or sexually attractive".²⁵ Clearly, this will not do for Joyce's purposes: first, because his argument has nothing much to do with sexual attractiveness; and second, because he explicitly says that he suspects there is no interesting analysis of the concept that is expressed by 'practical oomph'.

It would help to have some explanation of the categories of entities to which 'oomph' is supposed to apply. What is it that has practical oomph? In the first of the above passages, Joyce speaks as if *morality* has practical oomph. Unfortunately, it is unclear what sort of

²³Joyce, 2001: 17.

²⁴Joyce, 2008: 259.

²⁵*Concise Oxford English Dictionary*, 2008, online edition.

object morality is. As I typically read the term, ‘morality’ is supposed to refer to a subject matter. But it is unclear how to formulate an argument for the error theory on such an understanding. So, I reject this interpretation. In the second of the above passages, Joyce suggests that *moral normativity* has practical oomph. Unfortunately, I cannot understand what Joyce uses ‘moral normativity’ to refer to in this context. As I understand the term ‘moral normativity’, it is used to talk about moral reasons. But since Joyce intends his new argument not to involve reference to reasons, it seems that would be uncharitable to interpret ‘oomph’ this way.

Elsewhere, Joyce suggests that it is our *moral judgments* that have practical oomph. He writes,

When it is judged that someone is under a moral obligation, that judgment is imbued with a distinctive kind of practical oomph, but we theorists can make no sense of this quality, therefore nobody is ever really under a moral obligation. (2008: 259)

I believe that this leads to the most charitable understanding of Joyce’s new terminology. I’ll understand ‘practical oomph’ to be a predicate of judgments. I’ll also understand ‘practical oomph’ to be connected with some notion expressed by the phrase ‘must-be-done’, which is another term that Joyce uses in one of the above passages.²⁶

With this in mind, we may define ‘practical oomph’ as follows: a judgment that x ought to do ϕ has practical oomph if, and only if, in order for that judgment to be true, it has to be the case that x “simply must do” ϕ . I should note that this definition is not supposed to be explanatory of the content of ‘practical oomph’. Joyce appears to assert that there could be no such explanation. I merely intend the definition to help clarify the syntactic category of the term, so that we may at least formulate Joyce’s second argument. In addition, while it is tempting, we should not take ‘ x simply must do ϕ ’ merely to mean the same as ‘ x morally ought to do ϕ ’. For then the argument that I present below on behalf of Joyce surely would not be persuasive.

²⁶See the passage from 2008: 257.

Like his first argument, this argument focusses only on judgments involving moral obligation. It thus concludes in the same way as the first argument: that no one ever has a moral obligation.

P1'. For any agent x and action ϕ , if J is a judgment to the effect that x morally ought to do ϕ , then J has practical oomph.

P2'. For any judgment J , if J has practical oomph, then J is not true.

C. Therefore, for any agent x and action ϕ , x is not morally obligated to do ϕ .

As before, as far as I can tell, Joyce takes the premises of this first argument to be analytic. Consider, for example, his explanation of (P1'). He writes,

I maintain that *some* kind of special practical oomph is a necessary feature of moral judgment – though whether this quality can be given any clear articulation is open to question. ... The issue is whether this quality of “practical oomph” is an expendable aspect of morality – whether a normative system stripped of this quality would warrant the description “moral.” I claim that it is not and would not (respectively). Of course I realize how difficult it seems to assess the claim when it uses this purposely blurred term “oomph.” Yet it may be that this is the best that we can do, and it would be a mistake to try any better. (2008: 258-9)

In this passage, I take him to assert that the existence of practical oomph is a presupposition of our moral discourse about obligation. In chapter 2, I formulated a definition of ‘discourse presupposition’ according to which, when a proposition is a presupposition of our discourse, it is analytically entailed by judgments within the discourse. Since discourse presuppositions indicate the existence of analyticities, I interpret him as asserting that (P1') is analytic.

It seems, however, that Joyce has nothing new to offer as a rationale for (P1'). Any rationale that Joyce might offer in favor of (P1') would look very similar to his initial rationale for (P1). Recall the rationale for (P1): it would be incoherent for us to say, all in one breath, that Gyges morally ought not to steal but that he has no objective reason not to steal. If I understand Joyce correctly, this idea should transfer to his new argument: if we were willing to say that Gyges morally ought not to steal, then we would have to accept that he

“simply must not” steal. In order for our judgment to be a moral judgment, it must have the oomphiness that Joyce claims is essential to morality.

As for the second premise, I believe that Joyce expresses his rationale in one of the above passages when he says that “we theorists can make no sense of” practical oomph.²⁷

He also writes,

It seems to me not implausible that the sense of “practical requirement” with which natural selection may have endowed us ... is a primitive sort of feeling/thought which resists analysis, decomposition, explication, or naturalistic demystification. (2008: 259)

That is, he appeals to the idea that there can be no informative analysis of practical oomph, and so no judgment that has this feature could possibly be true. That Joyce appeals to this inference is further evidenced by the above passage where he writes, “...we theorists can make no sense of this quality, therefore nobody is ever really under a moral obligation.”

As far as I know, Joyce has expressed nothing in his writings that suggests any further justification for these premises. And I have no further charitable explication of his writings to offer in favor of them. We must now turn to the question whether the argument is persuasive.

In next section, I present what for me is the most compelling criticism of Joyce’s argument. Before I present that criticism, though, I’ll comment on two other worries one might have for the argument.

First, I have already noted that it is a mistake to infer that a putative quality doesn’t exist merely from the fact that we have no analysis of the quality. So, insofar as Joyce’s second argument relies on such an inference, it is unsound. There is, however, a related form of inference, which might seem to be what Joyce is appealing to, viz., ‘the predicate ‘F’ is unintelligible, therefore any statement that requires that ‘F’ apply to something is untrue’. Unfortunately, I have a hard time understanding how falsity could result from unintelligibility, and so I do not think that Joyce finds ‘practical oomph’ to be literally unintelligible.

²⁷See the passage from 2008: 259.

(When we say that a term is unintelligible, don't we mean that it is mere gibberish, something that doesn't really mean anything?) In addition, I do not believe that Joyce's choice to interpret speakers in this way is up to standards of charity. In my opinion, we should interpret speakers as being unintelligible only if all other reasonable interpretive options have been exhausted. I believe there are other options.

My second comment involves the question of whether we can 'make sense of' the notion of practical oomph. As a piece of biographical information, I must admit that I cannot make sense of what Joyce means by 'practical oomph'. As another piece of biographical information, I must note that, when I cannot make sense of a term, I have a hard time understanding arguments that are formulated with the term. It is hard to evaluate the premises of an argument when you cannot understand a central term that occurs in it! Thus, my main criticism of Joyce's argument will not focus on a unique premise of the argument. It is more of a big picture objection.

4.6 Criticism of Joyce's Second Argument

In general, a certain amount of caution is required in the evaluation of arguments that use vague or otherwise indeterminate terminology. They can implicitly equivocate (and, as a result, they are invalid), and their premises can end up not being determinately true (and, as a result, such arguments are unsound for want of true premises). In my opinion, such due caution will point us toward the main problem for Joyce's argument.

I begin this section by briefly explicating an approach to the treatment of vague language. Then, in subsection 6.2, I show how this treatment of vague language provides the resources for a convincing criticism of Joyce's argument.

4.6.1 Vagueness and the Sorites Paradox

Consider a long horizontal sequence of colored panels on a white wall. At the far left, there is a blue panel, and at the far right a yellow panel. There are one-hundred ninety-eight

panels in between, one-hundred of which are determinate shades of red. Thirty of the panels (further to the left) are determinate shades of blue-or-purple, and thirty of them (further to the right the right) are determinate shades of orange-or-yellow. What are we to say about the forty leftovers, twenty of which are to the left of the determinate reds, and twenty of which are to the right of the determinate reds?

Blue — — Purple — — [??] — — Red — — [??] — — Orange — — Yellow

I take it as a datum that we cannot assert that they are determinately red, nor that they are determinately not red. The ones on the left are on the borderline between purple and red, and the ones to the right are on the borderline between orange and red. But there is no antecedently justified method of resolving the question of their color (not even “in principle”).

How might we explain this sort of irresolvability? I will presuppose the following rough sketch of a proposal. ‘Red’, as we use the term, doesn’t determinately express a particular range of precise shades of color. Our patterns of usage of a term help to specify our tacit, conventional decisions about what the term is to mean. Since we haven’t decided in exactly what circumstances the predicate ‘red’ is supposed to apply, we therefore haven’t given it a determinate meaning. There are many slightly different determinate ranges of precise shades of color that are admissible given everything we’ve decided upon; at least, we haven’t ruled them out. There are also quite a few that are absolutely inadmissible; e.g., ranges that include determinate shades of blue or yellow. Call these, respectively, ‘admissible precisifications’ and ‘inadmissible precisifications’ of ‘red’. Intuitively, we may think of the admissible precisifications as possible ways of deciding, by fiat, which shades are to be red and which are not to be red.

This approach provides a compelling way to deal with what is known as the ‘Sorites paradox’, which can be expressed by the following argument:

8. If a panel is not red, then neither is the next panel to the right.
9. The panel furthest to the left is not red.
10. So, each panel is not red.

Since (10) is unacceptable, either the inference from (8) and (9) to (10) is invalid, or one of (8) and (9) is false. However, *prima facie*, both premises seem true. Clearly, the panel furthest to the left is not red, since it is blue. And how could such a tiny change in color result in a difference in redness?

Provided with this picture of how to deal with indeterminacy, we are in a position to say what is wrong with the argument in (8)-(10). First, let us say that one of the premises is *super-true* only if it is true on all of its admissible precisifications; *super-false* if false on all admissible precisifications. Thus (9) is super-true only if the panel furthest to the left is not included in the range of shades expressed by 'red' on each of its admissible precisifications. And (8) is super-true only if, given any admissible precisification of 'red', if a particular panel is not included in the range of 'red', then neither is the one to its immediate right.

But it then seems clear that (8) is, in fact, super-false. Provided with any determinate range of precise shades of red, some panel in the borderline area will count as not red, but the one immediately to the right will count as red. Intuitively, on any way of deciding by fiat which shades are red and which aren't, there will be a sharp cutoff between the red and non-red. (It doesn't really matter where, so long as the cutoff occurs somewhere in the borderline area.) And if (8) is super-false, then it cannot be used in a rationally compelling argument for (10).

4.6.2 An Application of this Approach to Joyce's Argument

I believe that considerations along the same lines show that Joyce's argument is not rationally compelling. Let me explain. First, let us instantiate (P1'), (P2'), and (C) to the case of Gyges once again. I do this only to make my discussion less complex. We get the following argument:

11. If someone judges that Gyges morally ought not to steal, then his judgment has practical oomph.
12. If a judgment has practical oomph, then it is not true.
13. So, if someone judges that Gyges morally ought not to steal, then his judgment is not true.

Joyce has explicitly asserted that ‘practical oomph’ is indeterminate.²⁸ As I understand this claim, it implies that our usage of ‘oomph’ does not decide on a precise class of judgments that are to be the ones that have oomph. I agree with this claim. (Indeed, it seems to me that our usage of ‘oomph’ may not even decide whether *any* moral judgments are to have oomph; but I’ll not press this point.) There are many precise ways of deciding for us, by fiat, which judgments are to have oomph. Similarly to the case of ‘red’, I’ll call each way of making precise our term an ‘admissible articulation’ of ‘oomph’.²⁹ And, similarly to sentences (8)-(10), we can say that (11) is super-true only if, given any admissible articulation of ‘oomph’, our judgment that Gyges morally ought not to steal is included in the class of judgments that have practical oomph. Likewise for (12): it is super-true only if on every admissible articulation of ‘oomph’, there are no true judgments included in the class of judgments that have oomph.

With this setup in place, we are now in a position to see whether Joyce has provided us with a convincing rationale for either premise. I believe that he has failed in this task. Moreover, I believe that it is impossible for him to succeed. Let me explain.

As I see it, in order to rationalize either (11) or (12), Joyce would have to give us reason to think that both premises come out true on all admissible articulations. But he has not done this. At best, he has only given us an argument that they come out true on one admissible articulation: where we take ‘practical oomph’ to be articulated in terms of categorical, non-institutional reasons. Joyce writes,

[I]n [*The Evolution of Morality*], I was not content to conduct the argument [for the error theory] in these mysterious terms [i.e., in terms of ‘oomph’]. (Perhaps I ought to have been.) Rather, I endeavored to give some concrete artic-

²⁸See above; Joyce, 2008: 257.

²⁹I do not use the terminology of “precisifications”, and have switched to talk of “articulations”. I do this for two reasons. First, because Joyce uses the term ‘articulation’ in one of the above passages. Second, because the notion of a precisification has become commonplace in the literature on vagueness, and it seems to me that, although ‘oomph’ is indeterminate, it is not literally vague (it doesn’t admit of borderline cases). Sometimes Lewis talks of ‘reasonable resolutions of semantic indecision’ (Lewis, 1993; 1989; 1982). I take it that this phrase stands for the same idea as my ‘admissible articulation’.

ulation of *oomph* So understood, the battle lines get drawn at the dispute over agents' *reasons*...

Of course, in this passage, Joyce is referring to his first argument for the error theory, which I criticized in the preceding sections of this chapter. Unfortunately for Joyce, however, we have seen that one or the other of his premises comes out false when 'oomph' is precisified in this way. This I have already argued, and I have nothing further to say about Joyce's first argument: it is unsound, no matter how you look at it. The point that I am now making is that the unsoundness of Joyce's first argument infects his second argument. If the first argument ends up with a premise that isn't true, then the second argument likewise ends up with an unacceptable premise.

We can motivate this point by extending our analogy to the Sorites paradox. Intuitively, the soundness of the argument that consists in (8)-(10) is parasitic on the soundness of its precisifications. (A 'precisification' of an argument is just a way of uniformly precisifying each of the premises.) If there is a precisification where the argument is unsound, then the argument from (8) and (9) to (10) is not *super-sound*. Since there is at least one precisification of (8) that is untrue, the Sorites reasoning based on that precisification is unsound. A parallel result holds for Joyce's argument: it is super-sound only if it is sound on each admissible articulation. Since it has a false premise on at least two articulations (I argued for this in the preceding sections of this chapter), the argument itself is not super-sound, and so it is not rationally compelling.³⁰

³⁰One might object to this criticism, claiming that an argument need not be sound on all articulations in order to be rationally compelling. On such a view arguments may be rationally compelling when they are merely sound on *most* articulations. Thus, Joyce might claim that his argument is "close enough" to being super-sound. In order for Joyce to make this claim, however, he must provide some reason to believe that the argument is sound on most articulations. Until he does this, it does not provide us any reason to accept his error theory.

4.7 Conclusion

I have criticized Joyce's first argument for his version of the error theory. A crucial phrase in his argument is ambiguous, and the ambiguity affects the truth of the premises of the argument: on no way of disambiguating 'has a reason' does the argument end up sound.

I have claimed that my objection to Joyce's first argument provides a compelling ground for criticizing his second argument. Since the first argument is supposed to be a "concrete articulation" of the second argument, the second argument would be rationally compelling only if the first argument is sound. But since the first argument is not sound, we must conclude that the second argument is no good, either.

My criticism of Joyce's arguments rested on some controversial claims. First, my objection to his first argument required that we recognize an ambiguity in our talk of "having reasons". Joyce may deny this, but I do not see how he could do so plausibly: the ambiguity is so manifest that any adequate account of our talk of "having reasons" would have to accommodate it. Second, my objection to his second argument presupposed that indeterminacy is a result of semantic indecision, and that Joyce treated his first argument as a more precise way of articulating his second argument. Joyce could deny either of these claims, but, again, I do not see how this would make his arguments any more plausible. I therefore conclude that both of Joyce's arguments fail to motivate his error theory. If there is a good argument for the error theory, it is not to be found in Joyce's writings.

I believe that this discussion of Joyce's view, especially of his second argument, uncovers an idea that would undermine any error theory like Joyce's. It is that moral discourse is in a considerable disarray – enough to make plausible the idea that there might be no determinately true version of the error theory. Since we have left so much undecided as to what we mean with our moral terminology, it is impossible to find an analytic commitment of our moral discourse that is a defective commitment on every way of precisifying the terminology. This makes it plausible that there might be compelling considerations in favor of the idea that no version of the error theory can determinately be the case. In the

next chapter, I develop this thought into an argument against any error theory according to which an entailment version of the error thesis is true.

CHAPTER 5

AGAINST THE ERROR THEORY

5.1 Introduction

At the end of chapter 4, I concluded that Joyce's arguments for his error theory do not work because of his failure to find a defective commitment of moral discourse that is sufficiently determinate. I claimed that the fallacy of equivocation infects his first argument, and that his second argument only brings with it more indeterminacy than the first. In this chapter, I argue that this worry does not merely plague arguments for error theories like Joyce's. It is a problem that cuts to the very heart of his error theory, and it therefore can be used to argue against all views like his.

So I aim to argue against any version of the error theory that holds Joyce's principal theses. Here are the relevant parts of my interpretation of his view:

Thesis 1IA: The Abundant Instantiation Metaphysical Thesis. There are abundant moral properties, but they are not instantiated; in specific, although there are abundant objective values, they are not instantiated.

Thesis 2EA: The Abundant Entailment Discourse Thesis. Every moral judgment analytically entails that some abundant moral properties are instantiated.

Thesis 3EA: The Abundant Entailment Error Thesis. Since abundant objective values cannot be instantiated, but our moral judgments analytically entail that they are instantiated, all our moral judgments are false.

While Joyce seems also to accept versions of these claims that invoke sparse conceptions of properties, those versions will not directly be relevant to the argument of this chapter. Thus my argument primarily invokes semantical considerations that tell against these error theories; metaphysical considerations are only indirectly relevant. In what follows, my talk of properties should be read as talk of abundant properties, or properties conceived as

meanings. Furthermore, my argument requires a strong reading of Thesis 1IA, on which not only are no moral properties actually instantiated: they cannot be instantiated, on pain of some kind of incoherence. In what follows, I'll use 'the error theory' to refer to any view that accepts these theses. Finally, as I see it, my argument ought to be effective against an error theory that includes Theses 2FA and 3FA, since such a view would be similar in the relevant respects to this version of the error theory.

I choose to focus my attention against this sort of error theory for roughly three reasons. First, because I believe that this is the kind of view that most philosophers have in mind when they comment upon error theories in morality. Second, because I believe that this is the kind of view that, deep down inside, most self-proclaimed error theorists would like to accept. Finally, because Joyce's view seems to be an error theory of this kind, and his is the most well-developed and nuanced view deserving of the name 'moral error theory'. It is this view that, in my opinion, deserves the name 'the most plausible form of error theory'.

My argument against the error theory goes roughly as follows. The view requires that moral discourse is determinate enough that our moral judgments all definitely contain some semantic incoherence. But moral discourse is in significant disarray, and so at worst our moral judgments only indeterminately contain such an incoherence. So, the error theory is untrue. In section 2, I outline some recent literature that seems to agree with the main thrust of my argument. In section 3, I explicate a framework in which to present my argument. In section 4, I present the premises of the argument in full detail and I begin to rationalize them. Sections 5 and 6 contain further rationalization of the main premise of the argument.

In addition to the main aim, this chapter has two subsidiary goals. The first is to sketch some possible reasons to believe an anti-realist view that may very well deserve to be called an 'error theory', but that has not been extensively discussed in the literature (sections 5 and 6). I call the view 'moral indecisionism', and I am justified in including discussion of it in this chapter because the considerations that support it are also the considerations that I use to argue against the error theory. Second, I suggest that moral indecisionism can be developed

in a variety of interesting ways depending on the motivations that one has to accept it, and thus it ought to be given further careful consideration (sections 3 and 7). Unfortunately, the chapter is often somewhat sketchy and speculative in the respects having to do with these secondary goals, but I see no other option. A full discussion of moral indecisionism could easily fill another dissertation-sized work.

5.2 Pieces of My Argument in the Literature

The goal of this section is to summarize parts of the meta-ethical literature that are most similar to the line of thought that I pursue against the error theory. There are plenty of similarities between these views and mine, but there are also some differences. It will be important to consider both the similarities and the differences. Subsection 2.1 discusses Kirchin's (2010) argument against the error theory, subsection 2.2 contains a discussion of John Burgess's (2007) and Stephen Schiffer's (2003 and 1990) views, and in subsection 2.3 I discuss Lewis's (2005 and 1989) meta-ethical views.

5.2.1 Kirchin's Tension

In a recent discussion of moral error theories, Simon Kirchin argues that, at present, there is no error theory that is sufficiently motivated, and that it is unlikely that anyone will ever generate a plausible motivation for moral error theories. He writes,

There exists a tension at work in the moral error theory that has received little if any attention. In order to make their position convincing moral error theorists must alight on a particular idea and argue that it is a crucial commitment of everyday thought and language. Further, in order to be sure that they can convict a commitment of error, error theorists might need to specify that commitment so as to rule out alternative defensible understandings. The danger is that the more one specifies the formulation of a commitment, the more it is likely that the commitment is less crucial to many people's everyday moral thought and language. ... Hence, it is questionable whether error theorists can plausibly convict everyday moral thought and language overall of error. (2010: 167)

Kirchin does not aim to show that all moral error theories are untrue. Instead, he wants to show that it is difficult to argue persuasively for an error theory. So his argument seems to be this: if an error theory were justifiable, then there would have to be some defective commitment of moral discourse that is not open to many “alternative understandings”. Since it is unlikely that there is any such commitment, it is unlikely that any error theorist will ever be able to argue compellingly for his view.

Neither does Kirchin seem to have explicitly in mind any particular kind of error theory. He makes it clear that any error theory that roughly satisfies my imprecise renderings of Theses 1-3 is to be included in his target.¹ However, I am inclined to think that, at least implicitly, his main target is Joyce’s form of error theory.

The most controversial part of Kirchin’s argument is the claim that there will always be “alternative understandings” of defective commitments in moral discourse. In favor of this, he writes,

What makes the moral stand out ... is the fact that moral thought and language is messy, disorganized, cluttered, and lacks a unique and specific identifying point. There might be nothing that is fundamental [to moral discourse] in the necessary way, or if there is it is so general that one has to specify it so as to argue that it is erroneous and in doing so one invites other plausible specifications that are equally legitimate. (168)

We get the following picture. Suppose an error theorist is conversing with a philosopher who is not an error theorist. The non-error theorist will want to know why the error theorist holds that view: what are his arguments for his error theory? The error theorist must then explain what he thinks is defective about morality. In order to do this properly, he must be very precise; otherwise, the non-error theorist will not be able to come to grips with the explanation. Once the error theorist has precisely explained what he thinks is the defective commitment, it will then be open to the non-error theorist to question why the error theorist accepts *that* as a commitment of moral discourse, instead of *this*: some other similar precise

¹See Kirchin, 2010: 168

commitment that is not defective. It is alleged that there is nothing an error theorist could do to alleviate this worry, and so the argument contained in his explanation should not compel the non-error theorist.

While I tend to agree with the main point of Kirchin's argument, I am unsatisfied with it for three reasons. First, I believe that Kirchin's rough idea can be used only to argue against a certain subset of error theories, viz., ones relevantly like Joyce's. So the target of Kirchin's worry should be much more focussed. Second, I believe that his idea can be used to show not only that the error theory is unmotivated, but also that it is untrue. Third, and most importantly, I believe that Kirchin's idea can be made more rigorous. In my opinion, the fundamental problem with the error theory has to do with indeterminacy in moral discourse. The reason why "other plausible specifications" of an alleged defective commitment are "equally legitimate" is that any ordinary commitment of moral discourse that is *prima facie* defective is also semantically undecided, and there are many reasonable resolutions of this indecision that are not *ultima facie* defective. A goal of section 3 of this chapter is to set up part of a semantic framework on which this claim of mine can be understood.

5.2.2 Burgess and Schiffer on Widespread Moral Indeterminacy

John Burgess (2007) and Stephen Schiffer (2003 and 1990) both hold that moral language is radically indeterminate, and so (virtually) all of our moral judgments lack truth-values. Their views, however, are not entirely similar, nor are they equally developed (Schiffer's is much more developed), so I will discuss them somewhat separately.

It is instructive to locate Schiffer's view within the rubric of the cognitivist/non-cognitivist divide and the realist/anti-realist divide. First, Schiffer is explicitly permissive about the existence of properties and propositions: he accepts roughly that for any meaningful predicate or sentence of natural language, there is a property or proposition expressed by that linguis-

tic item.² Further, he accepts a general thesis to the effect that “belief is the backbone of cognitive meaning”: sentences derive their meanings from our use of them to express our beliefs.³ As a result of this, Schiffer is a cognitivist about morality. But he also believes that the properties expressed by moral predicates do not have instantiation conditions: conditions under which their predicates would apply or fail to apply to things in the world. As a result of this, he qualifies as an anti-realist about morality: while there are moral properties, propositions, and so on, they do not successfully describe any parts of the world.⁴

Schiffer recognizes that he has a burden of explaining why moral properties are defective in this way. He writes,

My explanation of why moral properties fail to have instantiation conditions is as follows. Typically, we judge two people to mean the same thing by a predicate only if they use, or are disposed to defer to, the same criteria for applying the term, but this is not so in the case of moral predicates like ‘wrong’. The “criteria” of application a person associates with ‘wrong’ are determined by her ultimate moral views as to what is wrong, and *there may be great interpersonal differences among these views, these criteria of application*. At the same time, important interpersonal similarities of use ... account for our judging two people to mean the same by ‘wrong’, notwithstanding the different criteria of application they attach to the term. Here there is nothing to distinguish any one person’s criteria as the correct criteria of application; this is the result of there being sameness of meaning in the absence of conditions of application that are criterial for everyone. (1990: 611; emphasis my own)

Some commentary on this passage is in order. Schiffer refers to “interpersonal similarities of use” of moral terms between individuals who have different “criteria of application” for those terms. As I understand it, a person’s criteria of application of a term are some set of conditions that the person would accept as being roughly necessary and sufficient for the application of the term. In addition, according to Schiffer, interpersonal similarities of use of moral terms include the fact that we typically have motivational reactions to our

²Schiffer, 2003: 61-71; 1990: 603-5.

³Schiffer, 1990: 602.

⁴Schiffer, 2003: chapter 6.

moral beliefs; that there is some significant connection for us between moral judgments and moral emotions like shame, regret, and guilt; and that pieces of moral terminology come in families (for example, there are value terms, there is deontic terminology, and so on).⁵

The emphasized part of the passage also requires some explanation. In later work, he more clearly states roughly the same thought that is expressed in that part:

...as skeptics of moral realism have long claimed, whatever moral principle one believes a priori [i.e., whatever ultimate moral views one has], it would seem to be possible that there is someone else who doesn't believe that principle [who does not have all those views] but who does not differ in any relevant way as regards one's epistemic situation vis-a-vis the principle. ... This seems entirely conceivable to me, and if the example is adequately described, there need be no flaw of reason or concept possession or imagination to explain this person's not believing the principle; the conceptual strain is in imagining someone with such an unusual psychology. (2003: 247-8)

As I see it, these passages, taken together, contain the materials for an argument for the indeterminacy of moral language. It goes like this: if moral judgments had truth-values, then moral properties would have instantiation conditions. For, in order to ascribe truth or falsity to a judgment, one must first determine whether the property ascribed by the judgment is instantiated by the subject of the judgment. And if moral properties had instantiation conditions, then rationally irresolvable disagreement with respect to ultimate moral views would not be possible. For, if either party to the "disagreement" fundamentally associated application conditions that differed from the instantiations conditions of the relevant moral properties, then it wouldn't really be a disagreement. Instead, it would be some kind of linguistic or terminological confusion. But it is clearly possible to have this kind of disagreement with respect to morals; we can conceive of it if we think hard enough. So, moral language is radically indeterminate.⁶

Burgess expresses a similar view in the following passages:

⁵Schiffer, 1990: 609-11; compare the discussion in Hare, 1952: 146-9.

⁶There are important similarities between this argument and the main argument contained in Horgan and Timmons, 1991; cf. Balaguer, 2011: Section 1.1.

I...maintain that moral judgments are without truth-value.... My reason is this: It is part of the normal understanding of moral judgments that what truth-value they have is supposed to be independent of the person by whom, and the circumstances in which, they are uttered. Yet though each speaker might have quite definite (albeit unconscious) criteria for applying the term 'moral', there is not enough common to all speakers' criteria to provide 'Abortion is immoral' with a speaker-independent truth-value. (2007: 433)

He continues,

The spectacle of the diverse moral codes left us by stern Moses, gentle Buddha, and a thousand others, makes it apparent that there is not enough common to different people's criteria for what counts as right and wrong to make any answer to the question, 'Is abortion immoral?' the right one. (2007: 434)

As I understand Burgess here, to say that the truth-value of a statement is supposed to be "speaker independent" is to suggest that correct criteria of application for the predicate contained in the statement, if indeed there are any, are not determined by the speaker's own criteria that he associates with it, nor by the speaker's own psychological states. Thus he seems to reject most kinds of subjectivism about moral judgments.⁷ Instead, he seems to think that the correct criteria, if there were any, would be determined by the way people use, and are disposed to use, the predicate in public moral discourse.

Thus his argument seems to be this: if moral judgments had truth-values, then their truth-values would not be determined by speakers's psychological states. And so, the criteria of application for moral predicates would have to be shared widely between users of moral language. This is supposed to be a result of the idea that the correct criteria of application would be determined by public usage. But there seems not to be any such agreement; indeed there is a huge amount of diversity in the criteria of application that people associate with moral predicates (witness stern Moses and gentle Buddha). So, moral language is radically indeterminate.⁸

⁷Cf. Stevenson, 1944 and Harman, 1975. For some relevant discussion about subjectivism, see Schroeder, 2010: 65-9.

⁸Compare the similar but contrary argument in Balaguer, 2011: 384-6.

This argument bears some similarity with Schiffer's. Both Burgess and Schiffer suggest that considerations having to do with "moral diversity" undermine the thought that moral judgments have truth-values. On Burgess's view, this diversity is actual: for him, the gap-piness in moral judgments is a result of the widespread actual existence of disagreement in norms. Schiffer, on the other hand, merely cites an idealized sort of diversity: for him, the reason why moral judgments lack truth-values is that it is possible for there to be in-principle rationally irresolvable moral disagreement.

I am inclined to take seriously the considerations that Burgess and Schiffer use to argue for their views; the arguments in sections 5 and 6 partially rest on similar considerations. There are, however, important differences between their lines of thought and mine. First, both Schiffer and Burgess view their arguments as arguments *for* error theories. Burgess, for example, suggests that he and Mackie hold the roughly the same meta-ethical view; and Schiffer favorably compares his view with Mackie's.⁹ Indeed, the conclusions of their arguments are compatible with a version of the error theory on which moral judgments suffer from Fregean presupposition failure. On such a view, moral judgments will all be gappy in truth-value. It seems to me, however, that the premises that Burgess and Schiffer rely on to derive their results are incompatible with the error theory. For example, in Schiffer's view, moral properties entirely lack instantiation conditions; but the error theory has it that moral properties have instantiation conditions that are impossible to meet. Similarly, in Burgess's view, public usage has not determined any particular application conditions for moral predicates; but the error theory would have it that moral predicates have impossible conditions of application.¹⁰

⁹Burgess, 2007: 427; Schiffer, 2003: 259.

¹⁰I acknowledge, however, that in some general sense of 'error theory', both Burgess and Schiffer accept error theories about moral discourse.

Second, both Burgess and Schiffer accept that (virtually) all moral judgments are without truth-value.¹¹ But it will become apparent in what follows that one might accept the premises of my argument against the error theory without having to accept this sweeping result. As I see it, moral language may be semantically undecided enough so as to undermine the idea that moral predicates are incoherent, without also implying that (virtually) all moral judgments are without truth-value.

So while I tend to agree with many of the ideas in Burgess's and Schiffer's arguments, I do not believe that I have to accept conclusions as radical as theirs. In addition, I aim to use their premises as means to slightly different ends: I'll use their ideas to argue against the error theory.

5.2.3 Lewis: Earlier and Later

In this subsection, I outline two important thoughts from the work of David Lewis, one earlier (1989) and one later (2005). I begin with the earlier thought, and contrast it with the later one, as well as with the premises of my argument against the error theory.

Lewis (1989) presents and defends an analysis of moral values: to be a value is to be that which we are necessarily disposed, under ideal conditions, to desire to desire. His defense of the analysis includes a controversial theoretical understanding of it, on which many of our value claims would be indeterminate, wavering between incoherence and possible truth. Since it is the way he interprets his analysis with which I am concerned, I'll not comment on many of the implications of this analysis.

In Lewis's view, when we decide to search for analyses of our terms, we ought first to look to the theory that is implicit in our discourse involving the term. Thus, in order

¹¹Schiffer suggests that he is willing to countenance certain exceptions to this rule; see 1990: 612, where he writes, "Perhaps it will be agreed that there maybe huge interpersonal differences among ultimate moral principles but insisted that there are limits: no one who has our moral concepts can deny that it is wrong to torture people for smoking in public places. But little need change even if this were conceded. For while this would give moral properties some instantiation conditions, it would still leave a great deal that matters to us uninstantiated by moral properties. Morality, the sum of moral truths, would be the feeblest of guides." It is interesting to compare Schiffer's view with Balaguer (2011).

to determine the proper analysis of ‘value’, we must consider the theory about values that is carried about in our “folk morality”. Our folk theory of value is some collection of shared opinions about “values”. According to this methodology, then, to be a value is to be whatever is the most natural candidate that would make those shared opinions true.¹² There may, however, be problems in this process of philosophical analysis. Lewis writes,

One way to create indeterminacy ... is to define names implicitly in terms of a theory (folk or scientific), and later find out that the theory is wrong enough that nothing perfectly deserves the names so introduced, but right enough that some things, perhaps several rival candidates, deserve the names imperfectly. (1989: 92)

In this picture, the degree to which something *deserves a name* is proportional to the extent to which it approximates satisfying the parts of the theory in terms of which the name is analyzed. To deserve the name *perfectly* is exactly to satisfy the theory; to deserve a name *imperfectly* is to come sufficiently close to satisfying it (and it is something of a vague matter of temperament what counts as *sufficiently close*). As a result of this caveat, Lewis introduces a distinction between speaking *strictly* and speaking *loosely* about values.¹³ He writes,

There are no perfect deservers of the name [‘value’] to be had. But there are plenty of imperfect deservers of the name.... Strictly speaking, nothing shall get the name without deserving it perfectly. Strictly speaking, Mackie is right: genuine values would have to meet an impossible condition, so it is an error to think there are any. Loosely speaking, the name may go to a claimant that deserves it imperfectly. Loosely speaking, common sense is right. There are values, lots of them, and they are what we are disposed *de facto* to value. (1989: 93)

The reason why Lewis thinks there are no perfect deservers of the name is that there is nothing we are necessarily disposed to desire to desire. It is a near platitude that our psychological states are contingent, and so our second-order desires could easily have been

¹²Compare Jackson, 1998: Chapter 2 and Smith, 1994: Chapter 2. This technique is something of a reflection of Lewis, 1970. Obviously, what I have presented here is technically and philosophically much less sophisticated than Lewis’s official view. Be that as it may; it conveys the essential ingredients of the view.

¹³Compare Bedard, 1993; and Lewis, 1972.

different. But there seem to be plenty of imperfect deservers of the name ‘value’: things we are disposed to desire to desire, though perhaps we are not necessarily so-disposed.

There are two important things to note about Lewis’s (1989) view. First, he thinks that at least one of our moral terms, ‘value’, is semantically indeterminate. He thinks also that on at least one way of resolving this indeterminacy, values are impossible; and that on some other way, values are not only possible, but that they are actually instantiated. Similarly, the premises of my argument against the error theory will imply that much of our moral terminology is indeterminate. I, however, need not accept that moral terminology is indeterminate between the impossible and the possible. Only a weaker thesis is required in order to accept the premises of my argument, viz., that on some ways of resolving the indeterminacy in moral language, moral terms may actually apply to something. I hope to leave it open whether also on some other ways of resolving the indeterminacy, moral terms cannot possibly apply.

Second, Lewis does not intend to argue against the error theory. Indeed, in a certain light, Lewis may seem to be an error theorist, as he believes that, strictly speaking, there are no values as we ordinarily understand them. But it is also natural to understand Lewis’s view here as something of a rejection of the error theory. After all, it seems that the error theory would have it that on *the one-and-only* way of speaking, moral properties cannot be instantiated. Since Lewis rejects the idea that there is just one appropriate way of speaking of values, it therefore seems that he rejects the error theory to a certain extent.

In a posthumously published piece, Lewis seems to accept a meta-ethical view that is incompatible with his earlier one. He writes,

Morality...could presuppose an error. It consists in part of a (rather ill-defined) system of alleged truths; in part of a practice of appealing to those alleged truths in order to guide one’s own conduct or that of others; and maybe in part of other things.... Whatever may be said about the rest of morality, at least the alleged truths might carry presuppositions, and those presuppositions might include one or more of the [putative] errors.... It could be so – but I don’t think it is. The system of alleged truths is just too ill-defined. (2005: 317)

Lewis here suggests that our folk theory of morality is in considerable disarray. If this is correct, then it seems that there isn't really any way at all of speaking strictly about values. This is because, if folk morality is sufficiently ill-defined, there is no definite collection of opinions that constitute *the* folk theory of value. As I understand the picture we get from Lewis, we will as a result only be able to speak loosely about values: there will only be imperfect deservers of the name 'value'. For our folk morality, whereby we determine whether anything is a value, is "just too ill-defined".

Therefore, it seems to me that, unlike the earlier passages, in this passage Lewis is indeed arguing against the error theory. In addition, his argument seems roughly to be the same as my argument: the error theory requires that moral discourse be relatively determinate; but morality contains enough disarray to undermine such determinacy; so, the error theory is not true. (I further believe that this represents an important and heretofore unrecognized change in Lewis's considered meta-ethical views.)

There are two important upshots of this section. First, some philosophers (Kirchin and later Lewis) have already pursued arguments against the error theory that are very similar to my own. But Kirchin's doesn't use exactly the same set of ideas, and Lewis's argument is not very developed. Second, some philosophers (Burgess, Schiffer, and earlier Lewis) already accept ideas that, in my view, can be used to argue against the error theory. They tend, however, to accept slightly stronger claims than is required in order for my argument to go through. In the remainder of this chapter, I'll outline my argument against the error theory and defend its premises.

5.3 A Framework for Indeterminacy

In this section, I begin by extending the linguistic account of indeterminacy that I introduced in chapter 4 in order to apply it to more radical sorts of semantic indeterminacy. The definitions that I introduce in subsection 3.1 will be important for the argument that I present in the next section. In addition, in subsection 3.2, I discuss a variety of ways of

understanding the approach to moral discourse that naturally arises out of this framework. I begin the section with some broad, impressionistic brushstrokes in order to paint a picture of the account.

In this section and the next, I help myself to an understanding of ‘property’ on which every property determines a precise class of possible individuals, the possible things that instantiate the property. However, I do not presuppose a conception of properties on which they can strictly be individuated in this way: sameness in intension will not suffice for property-identity. (Similar stipulations hold for talk of propositions in these sections.) In addition, talk of properties is to be understood in an abundant sense; properties are still to be conceived as meanings. These are meant merely to be terminological stipulations, and I hope it is clear what they amount to, given the discussion of properties in chapter 3.

5.3.1 The Linguistic Account of Indeterminacy

Color terminology is not special in being semantically undecided. Just as we have not decided which precise range of the color spectrum ‘red’ is to apply to, for almost all of our terms, there are far more equally eligible candidate meanings than we ever acknowledge. Indeed, we don’t have any reason to acknowledge most possible meanings, nor could we acknowledge many of them if we tried. However, it is natural to think that some parts of our language are more semantically undecided, and are in more disarray, than other parts. First, there is the case of perfect collective decision of a linguistic community on a perfectly precise meaning. Terms that are most plausibly understood in this way probably include the logical connectives and (unrestricted) quantifiers.^{14,15} Then, with vague predicates, it seems that we come to a limited consensus on a “core” meaning, but leave open “borderline”

¹⁴Obviously, the idea that linguistic communities come to a consensus on meanings is an extreme idealization of the way in which our terms get their meanings; cf. Quine’s preface to Lewis, 1969. I take part of this metaphor to be cashed-out in the main text of Lewis, 1969.

¹⁵There may be other factors that “help us to decide” or “decide for us” on precise meanings, e.g., facts involving which meanings are “intrinsically eligible”. See Lewis, 1983 and 1981.

cases. A slightly more extreme sort of indecision occurs when there is something close enough to consensus on each of many possible “core” meanings. As I see it, this is true for ambiguous and polysemic terminology. Chapter 3 of this work should make it clear that I believe that ‘property’ is undecided in roughly one of these ways. Finally, at the far end of semantic indecision, our discourse involving a term may be in such disarray as not to make plausible the existence of any core denotations for the term. It is difficult to come up with uncontroversial ordinary examples of this kind, but I am inclined to think that ‘love’ and ‘hate’ are semantically undecided in this way. I believe that we could, with a certain amount of effort, construct artificial examples of terms that are radically semantically undecided, though I will not attempt to perform this exercise here.¹⁶ Clearly, there is a whole spectrum of more and less extreme sorts of semantic indecision in between these natural points.

I assume as my starting point that this is an accurate sketch of some of the crucial relations between our language and the world. We need some way of modeling this picture; the following is most convenient for me. If our terminology were perfectly semantically decided, then every name would refer to a single individual, every predicate would express a unique property, and every statement would express a distinct proposition. But since our language is far from being perfectly semantically decided, the best we can say is that each name refers to a *range* of individuals, each predicate expresses a *range* of properties, and each statement expresses a *range* of propositions.^{17,18} Intuitively, these ranges correspond to the various meanings that we could have meant, had our linguistic community come to

¹⁶Feldman 2010: 135 seems to suggest that ‘happy’ in English may be a case of this sort. Compare the ‘smidget’ example in Soames, 1999: chapter 6.

¹⁷As in chapter 4, I still ignore any problems of higher-order indeterminacy, such as the problem of higher-order vagueness. As I see it, these problems are orthogonal to the present discussion. See Williamson, 1994: 156-61.

¹⁸It may be more natural to say that names, predicates, and statements indeterminately express or refer to precise individuals, properties and meanings. I believe that this would result in a more-or-less equivalent model of the relation between our language and the world, and that any differences would be orthogonal to the purposes of this discussion.

more of a consensus. Each member of a range is a *reasonable resolution* of the semantic indecision of the term.

As an upshot of this sort of model, there is no straightforward way of assigning semantic properties, such as truth and falsity, to pieces of our language. If the language were perfectly decided, we could simply say that our terminology directly inherits the semantic properties of what it expresses: a statement would be *true* if the proposition it expressed were true; a predicate would *apply* to something if the property it expressed were instantiated by that thing; and so on. In chapter 4, I explained a natural way of associating the semantic properties of truth and falsity with statements that contain vague terminology. I now aim to generalize this in order to deal not only with vagueness, but with other more extreme kinds of semantic indecision, and also in order to cover other crucial semantic properties. First, we may say that a statement is *super-true* if every proposition that is a reasonable resolution of it is true; *super-false* if each of its reasonable resolutions is false; and *indeterminate* otherwise. Obviously, this exactly parallels the definitions that I provided for these terms in chapter 4. The only difference is that ‘reasonable resolution’ occurs wherever ‘precisification’ occurred before.

Truth and falsity are not the only semantic properties that are important for my discussion. Statements also have analytic entailments, where analytic entailment is conceived as a semantic relation. Indeed, a crucial part of the error theory involves a claim about the analytic entailments of moral statements. If our language were perfectly precise, then we could say that a statement entails a proposition if and only if the proposition it expresses entails that proposition. But since statements typically cannot be said to express unique propositions, this notion must be redefined as follows: a statement *analytically entails* a proposition *P* just when each of its reasonable resolutions entails *P*; that is, for any proposition in the range expressed by the statement, that proposition is true at a possible world only if *P* is also true at that world.

In addition, we need some way of talking of conceptual analysis, since the error theorist holds beliefs about the correct analyses of moral predicates. If our predicates expressed unique properties, then we could say that the analysis of a predicate ‘F’ would be the one true necessary *a priori* way of filling in the blank in the following schema:

The predicate ‘F’ applies to something if and only if that thing has the property _____.¹⁹

But since predicates express ranges of properties instead of unique properties, there will be no unique necessary *a priori* way of filling in the blank. For each reasonable resolution of ‘F’, there will be a different *admissible analysis*, where what goes in the blank is the precise, fully-semantically-decided name of the property that is the salient reasonable resolution of ‘F’.

Finally, because the error theorist believes our moral predicates are inconsistent, we require a way of talking of predicates that could not possibly apply to anything. If our language were perfectly decided, we could say that a predicate is inconsistent when the property it expresses is not instantiated by any possible individual. Now, however, it is most natural to say that predicates may only be either “unequivocally” or “equivocally” inconsistent. A predicate is *unequivocally inconsistent* if, and only if, for each of its admissible analyses, what goes in the blank is a precise name for a property that is instantiated by no individual at any possible world. A predicate is *equivocally inconsistent* just in case on some but not all of its admissible analyses, what goes in the blank is a name for a property that is uninstantiable. Otherwise, the predicate is *unequivocally consistent*.

¹⁹Necessity and *a prioricity* are, strictly speaking, not the only constraints. Intuitively, if I were to fill the blank with ‘of being F’, then I wouldn’t be giving an analysis of the predicate, for, while the result would be necessarily true and *a priori*, it would also be *trivial*. Also, if we suppose that ‘of being G’ gave us the correct analysis of ‘F’, then ‘of being G and either red or not-red’ would be another necessary *a priori* way of filling the blank, though it wouldn’t be another analysis. It is unclear to me how to solve these problems. I hope, however, that I have conveyed the crucial idea.

5.3.2 Some Brands of Moral Indecisionism

With these notions, we are almost in a position to state my argument against the error theory. A main premise in my argument will involve an appeal to the existence of widespread semantic indecision in moral discourse. If one thinks this is a feature of moral discourse, then one apparently has many choices about how to formulate his overall approach to moral language. In the remainder of this section, I outline some important “choice points”: questions that ought to be answered by anyone who accepts a view of this kind. As noted in the introductory section of this chapter, these views are kinds of ‘moral indecisionism’.

The reason why I am outlining these sorts of moral indecisionism is that different forms of the view may be more relevant to the argument to come than others. It will be important to keep these different forms in mind. I also want to classify the views expressed by Burgess, Schiffer, and Lewis, since their views seem to be kinds of moral indecisionism.

The first question that ought to be answered by the philosopher interested in moral indecisionism is: how extreme is the semantic indecision in moral discourse? Does the view hold that pieces of moral terminology are vague, that they are ambiguous or polysemic, or that they are infected with some kind of more radical sort of semantic indecision? As I see it, in order for such a view to qualify as a form of anti-realism, and in order to be an interesting alternative to most other views that have been explicitly accepted in the literature, it ought to imply that the indeterminacy in morality is one of the more radical kinds.²⁰ Most of our terminology is vague in some way, and so it is not terribly controversial to believe that moral terminology is included in this class. Finally, it seems that both Burgess and Lewis would hold views according to which the indeterminacy is relatively radical.

Second, there is the question of how wide-reaching the indecision is supposed to be. Are merely some of our moral predicates semantically undecided, or are they all undecided in this way? The sorts of moral indecisionism that I am presently concerned with need merely

²⁰See, for example, Shafer-Landau, 1994 and Balaguer, 2011.

hold that the moral predicates that the error theorist targets are undecided. Thus, if the error theorist claims that all of our moral predicates are inconsistent, then the moral indecisionist that is most relevant will hold that all moral predicates are semantically undecided. Since Joyce merely concentrates on deontic predicates like ‘obligatory’ and ‘wrong’, a weaker form of indecisionism may be all we need to countenance in order to undermine the heart of his view. In addition, it is interesting to note that Schiffer and Burgess hold that the indeterminacy is close to universal.²¹ Lewis, on the other hand, seems to hold a restricted version of the view, on which the semantic indecision in moral discourse may be located in our statements about values.^{22,23}

The third and fourth questions are intimately related, but are worth separating nonetheless. We must ask: are our moral predicates equivocally inconsistent? As I see it, a brand of indecisionism may answer ‘yes’ to this question and remain incompatible with the error theory. I presuppose this as part of my argument against the error theory. It is important to note that Lewis seems to answer ‘yes’ to the question, but both Burgess and Schiffer remain silent about it.

Finally, it is important for a moral indecisionist to take a stand on how the indeterminacy affects the truth-values of our moral statements. Are all the infected statements gappy, or may some of them be super-true or super-false? Clearly, semantic indecision need not result in uniform gappiness, as vague predicates are semantically undecided but some statements that involve vague predicates are super-true, and some are super-false. Burgess says that each statement that is infected with semantic indecision is gappy, whereas Schiffer suggests

²¹It is important to note, however, that Schiffer holds a different view about the nature of indeterminacy than I do here. See Schiffer, 2003: chapter 5.

²²But see Lewis, 2005: n2.

²³There are actually a couple of importantly different questions here. We might ask what is the *horizontal* extent of indecision, thus asking about which moral predicates are undecided. We might also ask, with respect to those predicates, what is the *vertical* extent of indecision. This second question may involve such issues as whether it is indeterminate what is the correct analysis of the predicate, or whether it is indeterminate whether we ought to build implicit relativization to moral standards into our semantic treatment of the predicate.

that he is willing to allow that some small minority of moral statements have determinate truth-values. The earlier Lewis, on the other hand, apparently would be allowed to say that some of our moral statements about value are super-false and some indeterminate. But he could not say that any of them are super-true, since he thinks statements about value are equivocally inconsistent.

Moral indecisionism therefore comes in many varieties. The most radical sorts of moral indecisionism would hold one of these views:

Radical semantic indecision infects every moral statement, and each one is equivocally inconsistent.

Radical semantic indecision infects all moral statements, but only some are equivocally inconsistent.

Radical semantic indecision infects all moral statements, and none of them is equivocally inconsistent.

Slightly less-radical versions of the view may weaken either of these claims, resulting, for example, in one of these views:

Radical semantic indecision infects some moral statements, and all of those infected are indeterminate.

Radical semantic indecision infects some moral statements, and only some of those infected are indeterminate, while some are super-true or super-false.

Thus it is consistent for a form of moral indecisionism to hold that some of our moral statements are super-true. Again, I note that it will be interesting to keep in mind whether the considerations that I advance in favor of my premises will commit me to particular answers to these questions.

5.4 Against the Error Theory

The main goal of this section is to lay out, in as clear and rigorous a manner as possible, my argument against the error theory. The argument crucially presupposes a framework on which semantic indecision is the primary cause of indeterminacy. I do not attempt to rationalize this assumption. The argument also crucially requires, as part of its main premise,

that moral discourse contains widespread semantic indecision. In the subsequent sections, I provide a barrage of considerations in favor of this hypothesis.

My argument involves three steps. The first step is to notice a trivial implication of the error theory: according to Thesis 2EA, every moral statement analytically entails a necessary falsehood, viz., that abundant moral properties are instantiated. This is trivially implied by the error theory because, on this view, the defect in morality is located in our moral predicates. Further, it holds that when we go to analyze the meanings of moral terms, we find that if they were to apply to something, the world would have to meet an impossible condition.

The second step in the argument is to draw out a corollary of this implication: if it's impossible for abundant moral properties to be instantiated, then the predicate of every moral statement must be unequivocally inconsistent. So, the error theory implies that the predicate in any moral statement is unequivocally inconsistent. Recall that in order to analyze the meaning of a predicate, we must fill in the blank for each of its admissible analyses. If there are many reasonable resolutions of semantic indecision for the predicate, then there will be many admissible analyses; many ways of filling in the blank. Thus, in order for the error theorist to claim that abundant moral properties are necessarily uninstantiated, he must also accept that, for each admissible analysis of a moral predicate, what goes in the blank is an impossible condition. That is, the predicate must be unequivocally inconsistent.

The final step in the argument is the main premise: at worst, moral predicates are merely equivocally inconsistent. It would follow from this premise and the corollary in the second step that the error theory is not true. So clearly this premise requires some rationalization.

The reason why I accept the main premise comes in two parts. First, like Kirchin and Lewis, I believe that moral discourse is “messy” and “ill-defined”. As a result of this, upon reflection, we ought to allow that there will be no unique correct interpretation of the language contained in moral discourse. There will be a multitude of admissible interpretations. Second, while it may be that on some of these interpretations of moral language, the pred-

icates therein will express impossible conditions, charity considerations dictate that on at least some of the interpretations, moral predicates will express possible conditions.

The next section contains further justification for the first part of this rationale for the main premise. I'll now spend some time to explain what I mean by "charity considerations". I here mean to invoke a version of *the principle of charity*, according to which "we should work on the general presumption that the folk are not badly confused".²⁴ This is obviously very vague, and there are at least two crucially different ways of making it more precise: one stronger and less plausible; and one weaker and more plausible.

On a stronger version of the principle of charity, it might be demanded that we interpret ordinary speakers as failing to mean truths, or as having irrational beliefs, only if there is some kind of overwhelming evidence for this. I do not intend to invoke the principle of charity in this strong form. Instead, I find it somewhat more plausible that, as a constraint on interpretation, we give an initial presumption to the hypothesis that ordinary speakers mean truths and don't hold irrational beliefs. This is the weaker version of the principle of charity.

Notice that an error theorist would apparently have to violate this principle at least some of the time. For the error theorist interprets people as being in incoherent belief states: belief in a moral proposition would be belief in a necessary falsehood. However, belief in a necessary falsehood constitutes a strong form of irrationality: if there is anything that a rational agent will not believe, it is a necessary falsehood. Now, it may be that on some admissible interpretations of moral language, the presumption grounded in the principle of charity is overridden by other considerations. But given that our moral language is "messy and ill-defined", we ought to expect that there will be interpretations of moral discourse where moral thinkers do not hold incoherent belief states.

²⁴Jackson, 1998: 103. For references, see also Lewis, 1974: 112.

In summary, we ought to expect that at least sometimes the presumption of charity wins out over opposing considerations. On these admissible interpretations, moral predicates will not be inconsistent. As a result of all this, we should conclude that moral predicates are not unequivocally inconsistent; at worst, they are merely equivocally inconsistent. And so the error theory is not true.

5.5 A First-Person Consideration

If the main premise of my argument is to be justified, then there must be some good reason to think that moral language is semantically undecided. A crucial part of the rationale for the main premise involved my belief that moral discourse is radically semantically undecided. In this section and the next, my goal is to provide some reasons to believe this hypothesis.

I want to be very upfront about one major caveat: I believe that these considerations, taken individually, provide only scanty support for this hypothesis. I further believe, however, that the arguments I advance in these sections, taken jointly, provide at least some non-negligible reason to accept the hypothesis. I am inclined to think that, unfortunately, only a substantial amount of empirical linguistics could ever give us very powerful reasons to accept moral indecisionism.²⁵

In this section, I outline the first consideration in favor of my main premise. I do not expect it to compel most others, but I take it to be a reason for me to accept the premise. It is something of a first-person consideration, and so I would expect it only to carry over to others who have a similar frame of mind. Very roughly, it is this: since I have some access to what words in English mean, and it seems to me that the meanings of moral terms are not very determinate, I have some reason to think that moral terminology is quite semantically undecided.

²⁵Compare, e.g., Balaguer, 2011: 378-84.

I am a reasonably competent speaker of English. I grew up speaking the language. A consequence of this is that I have some access to what our terms mean. Of course, my ability to know the meanings of English terms may be limited in a variety of ways, and it may not be totally clear how I have it. But these are not reasons to reject the claim that, in general, I have an understanding of what words of English mean.

In virtue of my possession of this capacity, I am able to do a certain amount of conceptual analysis. I can think about possibilities, both ordinary and outlandish, and then consult my linguistic intuition as to whether a certain word can correctly be said to apply in those possibilities.²⁶ I further believe that my understanding of English words allows me to hold their meanings “before my mind”. When I use this ability, I can come to reasonable conclusions about the parts of a concept expressed by a word, and I can also come to know about connections between the concept expressed by a word and other concepts. For example, I know that if you’re a bachelor, then you must be an unmarried adult male.

But when I reflect on the meanings of moral terms, my mind’s eye sees only a blur. I can think about possibilities where I know I am *supposed* to have inclinations to make moral claims: the little boys who burn the cat, the older sibling who gets his little brother to steal, and so on. I know that some people would expect me to judge that ‘wrong’ applies in either of these situations. But I typically don’t have any inclinations to make these judgments.²⁷ In addition, when I hold my moral concepts before my mind, and inquire about their components and their connections to other concepts, I am unable to see anything but blurry components and connections.

One might suggest that I investigate the way others use moral language in order to gain a more determinate set of moral concepts. (Indeed, I think it’s plausible that I came to my

²⁶Cf. Jackson, 1998: Ch. 2.

²⁷Of course, sometimes I can bring myself to feel strong moral emotions about situations. If I concentrate on certain features of the burning of the cat, I begin to feel moral repugnance. But this is not a judgment to the effect that ‘wrong’ applies in the situation. Sometimes I also am inclined to judge that the burning of the cat is “immoral”. But, again, this is not a judgment to the effect that ‘wrong’ applies. Rather, I believe it is a judgment to the effect that some negative moral term or other applies in the situation.

understanding of moral terminology through this sort of process.) But I don't trust that I can look to the way other people use moral language in order to help me get any clearer on the criteria of application of moral terms. Usually, if you're uncertain about when a word can correctly be said to apply, you can look into a dictionary, ask someone who is older than you, or watch the way others use the word. In this case, however, the existence of widespread disagreement about how to use moral terms prevents me from taking the usage of others too seriously: most other speakers seem just as confused about moral terminology as I am.²⁸ In my opinion, any evidence about the criteria of application for moral terms that I might gain from others is undermined by the huge amount of actual moral diversity.

So I conclude (to myself): there must not be any clear, distinct, decided criteria of application there in the first place. I don't have any clear moral concepts. As far as I can tell, the majority of language users don't either. Indeed, this disarray is exactly what I would expect of a discourse that contains a radical amount of semantic indecision.

5.6 Some Inferences to Explanations

In this section, I argue that the truth of moral indecisionism would help to explain some phenomena underlying a variety of meta-ethical puzzles. The puzzles all have to do with the idea that the moral analytically supervenes on the descriptive. The phenomena that I'll discuss are: (i) the attraction of Hume's "Law", (ii) the open questions of Moore's famous argument, and (iii) the conceivability of rationally irresolvable moral disagreement. In each case, my claim is that, given the uncontroversial existence of analytic supervenience, and the truth of the error theory, these phenomena are unexpected. But on moral indecisionism, the phenomena are not so surprising. So, that is some reason to believe moral indecisionism instead of the error theory, thus some reason to accept the crucial part of my main premise.

²⁸This is where I use the main point that Burgess, 2007 makes in his argument. But whereas he uses this point positively to show that moral language is undecided, I use it to reply to a possible objection to the consideration that I am presenting.

5.6.1 The Supervenience of the Moral on the Descriptive

It is nearly universally accepted in the philosophical community that moral predication supervenes on descriptive predication. That is, there could be no difference or change in the world with respect to morals without a difference with respect to descriptive features of the world. Further, it is relatively uncontroversial that this supervenience is the broadest possible: it is not contingent or restricted in any way. For those who accept the notion of analyticity, it's natural to say that the moral analytically supervenes on the descriptive. Apparently with this in mind, Simon Blackburn writes,

It seems to be a conceptual matter that moral claims supervene upon natural ones. Anyone failing to realize this, or to obey the constraint, would indeed lack something constitutive of competence in the moral practice. And there is good reason for this: it would betray the whole purpose for which we moralize, which is to choose, commend, rank, approve, or forbid things on the basis of their natural properties. (1984: 137)

It is easy to find echoes of this passage from Blackburn.²⁹ As I see it, this is because it is not only analytic that the moral supervenes on the descriptive. It is also *obviously* analytic: even for those of us who haven't explicitly thought about the matter, if we were to think about it, we'd come to this conclusion.

It is not entirely clear how to formulate the thesis, even though the underlying thought may seem obvious. It is standard to formulate supervenience claims in terms of properties: if different moral properties hold of two situations, then different descriptive properties must also hold of them. This, however, would beg the question against the moral indecisionist. The reason why is that the most natural way to define 'moral property' is to take moral properties to be the properties expressed by moral predicates. But for the indecisionist, moral predicates express no unique properties; at best, the indecisionist can only say that moral predicates express ranges of properties. Thus, to accept this definition, in the present context, is to presuppose that moral indecisionism is false.

²⁹See, e.g., Balaguer, 2011: 370; Schroeder, 2007: 70-1; Jackson, 1998: 118-9; Smith, 1994: 40; Hare, 1952: 80, 145

In my opinion, it is better in the present context to formulate the supervenience principle directly in terms of the linguistic items in question. Thus, the supervenience of the moral on the descriptive becomes this:

For any possible situations x and y , if it is correct to say that x and y satisfy different moral predicates, then it must also be correct to say that x and y satisfy different descriptive predicates.

Or, bit more informally: if you encounter two situations that are exactly alike with respect to which descriptive predicates you would apply to them, then you must also ascribe exactly the same moral predicates to them. Alternatively, if you're willing to say that different moral conditions hold of the situations, then you must also be willing to say that different descriptive conditions hold of them.

There are three further issues, which I note but do not intend to discuss in depth. First, there is the question of how we are to individuate the class of moral terms from the class of descriptive terms. As in the rest of this work, I presuppose merely an intuitive understanding of this distinction. Second, since I have formulated the supervenience principle directly in terms of predicates, it is important to be clear about what language the predicates are in. Am I merely talking about some natural language, some artificial language, or some hybrid thereof? There is a worry that, if the supervenience principle is to be relativized to a natural language, then the language may not have enough expressive power within its descriptive predicates to make true the supervenience principle (for example, its descriptive predicates may not be able to distinguish between possibilities that its moral predicates are able to distinguish). To allay this worry, I allow that the supervenience principle may be relativized to an extension of a natural language such as English, which contains not only the predicates of the natural language, but also predicates that denote perfectly natural properties.^{30,31} Third, there is the further question of what variety of supervenience this is

³⁰On “perfectly natural” properties, see Lewis, 1983 and 1986.

³¹I believe that this does not make the supervenience principle trivial, since it seems to be uncontroversial that moral predicates do not express perfectly natural properties.

supposed to be.³² For my purposes, all that matters is that, when a supervenience principle holds true, there are corresponding principles of the same logical strength that say which subvening conditions are sufficient for any supervening condition.³³ In the present context, this amounts to the truth of a principle of the following form, for any moral predicate ‘M’:

If it is correct to ascribe [insert descriptive predicates here] to *x*, then it must be correct to ascribe ‘M’ to *x*.

That is, there must be some descriptive basis for any correct ascription of a moral predicate. This idea is an important part of the intuition that the moral depends upon the descriptive.

5.6.2 Interest in Hume’s Law

I intend ‘Hume’s Law’ to refer to whatever vague thesis was apparently affirmed by Hume in the following passage and extensively discussed during the twentieth century under a variety of other names.³⁴

“In every system of morality, ...the author proceeds for some time in the ordinary way of reasoning, and establishes the being of a God, or makes observations concerning human affairs; when of a sudden I am surpriz’d to find, that instead of the usual copulations of propositions, *is* and *is not*, I meet with no proposition that is not connected with an *ought* or an *ought not*. This change is imperceptible; but is, however, of the last consequence. For as this *ought*, or *ought not*, expresses some new relation or affirmation, ‘tis necessary that it shou’d be observ’d and explain’d; and at the same time that a reason should be given, for what seems altogether inconceivable, how this new relation can be a deduction from others, which are entirely different from it. (Hume, 1740: 469)

I find this idea initially compelling: prior to any systematic investigation into the meanings of the technical terms involved, and without couching it in a sophisticated philosophical

³²On the distinctions between strong and weak supervenience, and local and global supervenience, see Kim 1993; Sider, 1999; and Bennett, 2004.

³³Terminology: a subvening condition is a condition that makes it into the base of the supervenience principle; in this case, a descriptive condition. A supervening condition is a condition that makes it into that which is said to supervene; in this case, a moral term.

³⁴See, e.g., Searle, 1964; Hare, 1967; Jackson, 1974; Prior, 1960.

theory, I am inclined to agree with Hume that an ‘ought’ cannot be derived from an ‘is’. Further, I find the idea initially to be a deep and interesting truth about the independence of moral language from descriptive language.

I do not think this is idiosyncratic. I think that most ordinary people, if you put it to them, will, after a bit of reflection, be inclined to agree with me about this. In addition, the fact that this has been a main point of contention in meta-ethics also suggests that many philosophers have found the idea *prima facie* deep, interesting, and true. Why would so many philosophers take so much time to discuss it if they didn’t think this?³⁵

This should not generally be expected of subject matters that are intimately related in the way the moral and the descriptive are supposed to be. We should not find Hume’s Law a compelling truth. Provided with the fact that it seems obvious that the moral analytically supervenes on the descriptive, we ought to expect there to be, for any moral claim, a descriptive claim that is analytically sufficient for it.^{36,37} And if we ought to expect this to be the case, then it would be more natural for us to find Hume’s Law false – and therefore *not* deep and interesting.

That is, we should expect there to be analytic principles linking the descriptive with the moral, so long as we presuppose that moral language is relatively semantically decided. But given the hypothesis of moral indecisionism, we should expect these principles to be indeterminate (neither super-true nor super-false). For example, on some reasonable resolutions of ‘wrong’, stealing may be analytically sufficient for wrongness, while on others, it may not be analytically sufficient. Further, the moral indecisionist claims that this is a deep and interesting fact about the moral, since this is not true of our discourse in general: not everything we utter is radically semantically undecided like moral language.

³⁵I am further inclined to think that the literature on the subject tends to show that, given standard semantic and logical assumptions, either Hume’s Law is not deep and interesting, or else it is false.

³⁶This is more or less the main point of Blackburn (1984).

³⁷Terminology: a predicate ‘F’ is analytically sufficient for ‘G’ just in case ‘if something is F, then it is G’ expresses an analytic entailment.

Finally, in my opinion, the error theorist must say that it is a confusion on our part to find Hume's Law compelling. Indeed, this seems to be the move that Joyce favors.³⁸ According to him, either Hume's Law is trivially true, or (nearly) trivially false, but never is it deep, interesting, and true. So, it seems to me that the moral indecisionist has a better explanation of this phenomenon, since it is able to preserve the idea that Hume's Law is an important fact about moral language.

5.6.3 Open and Closed Questions

In his *Principia Ethica*, Moore allegedly uses the famous open question argument to undermine various forms of "naturalism" in meta-ethics. It is a matter of some controversy how to interpret Moore's arguments.³⁹ As I see it, any such Moorean argument turns on the existence of what I'll call 'open question phenomena'.

Consider any moral value predicate. I'll use 'good', though analogous phenomena occur with respect to other predicates of value or disvalue, e.g., 'bad', 'desirable', 'evil', or 'preferable'.⁴⁰ Now imagine any plausible attempt to provide a descriptive analytically sufficient condition, 'D', for the correct application of that predicate. This will yield a sentence of the form, 'If something is D, then it is good'. We can conceive of a conversational context in which someone who is a competent speaker of English, and who understands the meanings of the terms involved, may sensibly and genuinely utter the question, 'Is everything that is D also good?' This is the *open question phenomenon* with respect to 'good' and the descriptive predicate 'D'.

³⁸Joyce, 2006: 152-6

³⁹See Feldman, 2005 for explication of the arguments that Moore actually uses, and for further references.

⁴⁰In fact, it seems to me that open question phenomena occur with respect to most moral predicates. But Moore only used the open question argument to show that naturalism about evaluative predicates is incorrect.

It is puzzling that open question phenomena exist.⁴¹ In my opinion, given standard assumptions, we should expect there not to be any open question phenomena. This is because of the fact that the moral analytically supervenes on the descriptive. Given supervenience, we should expect there to be, for any evaluative predicate, a descriptive predicate that is analytically sufficient for its application. But the open question phenomena undercut this idea: if indeed it is always conceivable to sensibly ask the open question, then there are possible discourse contexts in which there are some possible worlds where ‘D’ is not materially sufficient for ‘good’. (This is the natural explanation of the openness of this question on the Stalnakerian model of conversation that I introduced in chapter 2. I continue to presuppose that model in this section.) But if there is such a possible world in the discourse context, then it seems that ‘D’ is not, after all, analytically sufficient for ‘good’. So, we should expect open question phenomena not to exist.

That is, we should expect open question phenomena not to exist, provided that moral language is not radically semantically undecided. But on the hypothesis of moral indecisionism, there is a good general explanation of these phenomena.⁴² For the moral indecisionist about moral value predicates, there are no descriptive predicates that unequivocally entail any moral predicates. Instead, on some reasonable resolutions of the semantic indecision in ‘good’, one descriptive predicate (say, ‘pleasant’) may be sufficient for ‘good’, while on others, a different descriptive predicate (say, ‘what we desire to desire’) may be sufficient for ‘good’. When a speaker asks, ‘Is everything that is pleasant also good?’, he is inviting his conversational participants to resolve the meaning of ‘good’ in such a way that there are worlds in the discourse context wherein some pleasant things are not also good.

⁴¹As I see it, Moore essentially embraced the puzzle when he concluded that ‘good’ must be a name for a “non-natural” property. It is important to note that Moore’s (1908) distinction between natural and non-natural properties is not the same as Lewis’s (1983) distinction between natural and unnatural properties.

⁴²Compare Lewis, 1989: 85-6.

If the parties to the conversation are agreeable conversational participants, then they will accommodate his invitation.⁴³ Thus it will be sensible for him genuinely to ask the question.

In my opinion, the best explanation of open question phenomena that is available to the error theorist must invoke the idea that some analytic truths may be “unobviously” analytic.⁴⁴ On this view, the reason why the open question phenomena exist is that it is not easy to learn the correct analysis of ‘good’. Joyce, for example, writes, “A priori truths may be less than obvious to us, and may be revealed only after a lot of intellectual labor.” He approvingly cites Michael Smith, who writes,

An analysis of a concept C in terms of another concept C* is correct just in case knowledge of C* gives us knowledge of all and only the platitudes surrounding C: that is, knowledge of all and only the inferential and judgemental dispositions of someone who has mastery of the concept C. ... Why are analyses unobvious and informative? Because even though someone who has mastery of some concept C must have certain inferential and judgemental dispositions, it may not be transparent to her what these inferential and judgemental dispositions are, and so, a fortiori, it need not be transparent to her what the best summary or systematization of the platitudes that describes these dispositions is. (1994: 38)

Thus, according to this explanation, there is some descriptive predicate ‘D’ that is analytically sufficient for ‘good’, and whenever someone can sensibly ask ‘Is everything that is D also good?’, that person hasn’t done the “intellectual labor” required to know this.

It seems to me that this explanation is undermined by a generalization of the original open question phenomena. Consider an individual who, in search of an analysis of ‘good’, has done whatever “intellectual labor” that is in general required to know the analysis of a concept. Perhaps this individual – the *ideal speaker* – has taken into consideration all of his “inferential and judgmental dispositions” having to do with ‘good’. I believe that, for any plausible descriptive predicate ‘D’, we are still able to conceive of a conversational context

⁴³On the concept of accommodation, see Lewis, 1979; and Stalnaker 1973 and 1974.

⁴⁴See Joyce, 2006: 149-50; compare Lewis, 1989: 85-6. On this view, the open question argument and the paradox of analysis involve the same sets of issues.

in which the ideal speaker may sensibly and genuinely utter the question, ‘Is everything that is D also good?’ Further, he may utter this question to an audience full of individuals who are ideal in exactly the same respects: each of them has taken into consideration all the platitudes involving moral terms, and so on. If this is true, then the error theorist’s best explanation of the open question phenomena does not represent a solution to the general puzzle that is posed by their existence.

In addition, the moral indecisionist’s explanation *does* generalize: the ideal speaker’s speech act should be viewed partially as an invitation to resolve the meaning of ‘good’ in such a way that there are worlds in the discourse context at which something is pleasant but not good. So the existence of open question phenomena gives us some reason to accept moral indecisionism over the error theory.

5.6.4 Rationally Irresoluble Fundamental Moral Disagreement

Disagreement in morality is commonplace. Often it seems to stem from differences in opinion about empirical matters of fact. Is it wrong to kill fish? It is natural to think that the answer to this question at least partially depends on whether fish feel pain. We tend to presuppose that the parties to a typical disagreement over this question will share some ultimate moral view to the effect that it’s wrong to inflict pain, but what they will disagree about is whether fish in fact feel pain. This sort of moral disagreement is not particularly puzzling or problematic. It could be resolved (at least “in principle”) were we provided with sufficient descriptive empirical information (which, “in practice”, we may never acquire).

But there is another kind of disagreement that does represent a genuine puzzle: disagreement in ultimate moral views, or “fundamental” moral disagreement, where there is apparently no way rationally to resolve the dispute. There are a couple of ways to explicate

the notion of rationally irresolvable fundamental moral disagreement. It is useful to present each of them.⁴⁵

First way. Imagine that two individuals disagree over the question whether it is wrong to kill fish, even though they possess all descriptive empirical information, they can easily draw any semantically permissible inferences from this information, and they have done as much “intellectual labor” as possible in analyzing their moral concepts. It is hard to conceive of how such a disagreement might rationally be resolved. Of course, we could imagine that one of the parties to the disagreement is an extremely persuasive rhetorician. He might use these abilities to browbeat his opponent into submission (either by forcing her to accept a different set of ultimate moral views, or by forcing her to accept an answer to the question that doesn’t cohere appropriately with her actual views). But this would not result in a *rational* resolution of the disagreement. Further, we might ask the individuals to consult their ultimate views about wrongness once again, just to see whether they made a misstep somewhere in their reasoning in favor of their respective verdicts. But it should not surprise us that each ends up with the same verdict he accepted before; for these individuals, the reasoning will go exactly as it went before. They’re extremely unlikely to change their views.

Second way. Imagine that two ideal speakers of the same language have arrived at different analyses of their moral terminology after going through the requisite tough intellectual labor. As a result of the fact that they accept different analyses, they disagree over which descriptive conditions the correct predication of ‘wrong’ supervenes. Thus they will accept different descriptive conditions as analytically sufficient for the correct application of ‘wrong’. For example, one might accept (1), whereas the other might accept (2):

⁴⁵In my opinion, the notion of fundamental moral disagreement that I attempt to explicate in this section is similar in many important respects to the sort of disagreement that is supposed to hold between the Earthlings and the Moral Twin Earthlings of Horgan and Timmons, 1991. So, I think that moral indecisionism provides a nice explanation of the intuitions that we have about that sort of disagreement, as well. For a level-headed discussion of the Moral Twin-Earth argument, see Balaguer, 2011.

1. If it is correct to ascribe ‘...maximizes utility’ to x , then it must be correct to ascribe ‘...is not wrong’ to x .
2. If it is correct to ascribe ‘...treats someone as a mere means’ to x , then it must be correct to ascribe ‘...is wrong’ to x .

Since it is possible for something to satisfy the antecedents of both (1) and (2), this would result in disagreements over everyday moral questions. These disagreements, however, would stem from a more fundamental disagreement as to the correct criteria of application for moral predicates. And again, it is unclear how the disagreement might rationally be resolved. We can imagine that, in the course of their intellectual labor, each individual has systematized his inferential and judgmental dispositions equally well; each of them has reached an ideal state of “reflective equilibrium”.

I intuit, with Schiffer, that this kind of moral disagreement is possible.⁴⁶ In addition, it seems to me that, given standard assumptions, we should expect such a disagreement not to be possible. Typically, if individuals associate different conditions of application with a term, it is intuitive to think that they aren’t “really” disagreeing on any substantive matter.⁴⁷ If they accept that the correct ascriptions of the salient predicate supervenes on different bases, then they seem to mean different things with the predicate. Instead, it often seems correct to say that they are merely engaging in a disagreement in terms; they are talking past one another. But this is not plausible in the moral case: we may stipulate that the parties to our disagreement have the same (or largely overlapping) inferential and judgmental dispositions with respect to moral terminology. As a result of this, it seems that they cannot merely be talking past one another.⁴⁸

⁴⁶Schiffer, 2003: 244-56. Notice that R.M. Hare seems to have preempted Schiffer’s argument in Hare, 1952: 148.

⁴⁷Or at least one of them is committing some kind of semantic error: he has used a term, thinking that it means one thing, when it determinately means another thing. But in the situation as I have described it, there is good reason to think that neither of them is making any kind of semantic error. (See the long footnote below.)

⁴⁸This is roughly the point that Schiffer makes when he says that there are “important interpersonal similarities of use” between moral speakers; see 1990: 611. Further, I believe that this is what Horgan and

The moral indecisionist has a good explanation of what is happening in a situation like this. According to the indecisionist, there is no unique correct analysis of ‘wrong’ in descriptive terms. Instead, since ‘wrong’ merely expresses a range of properties, there are many admissible analyses: one for each reasonable resolution of the meaning of ‘wrong’. On each way of resolving the indeterminacy in ‘wrong’, it ends up that the disagreeers express incompatible propositions, since (we are imagining) one person utters ‘ x is wrong’ while the other utters ‘ x is not wrong’. So, they genuinely disagree.⁴⁹

In addition, moral indecisionism has a good explanation for why there is no way of rationally resolving the issue: there is no rational way to guarantee that the individuals will accept compatible positions; neither person is rationally forced to accept the other person’s assertion as to the wrongness of x . It is natural for the indecisionist to say that each individual accepts criteria of application for ‘wrong’ that correspond to some but not all of its reasonable resolutions. And there is nothing in their inferential or judgmental dispositions that makes either individual’s criteria preferable to those accepted by the other. That is, in the example, there is nothing about the meaning of ‘wrong’ that suggests ‘wrong’ should mean the same as ‘maximizes utility’ instead of ‘doesn’t treat anyone as a mere means’.⁵⁰

Timmons rely on to elicit the intuition that Earthlings and Moral Twin Earthlings are disagreeing; see Horgan and Timmons, 1991. Compare also the discussion in Hare, 1952: 148.

⁴⁹There may be some intuition to the effect that they are not really disagreeing, or not having a substantive disagreement. The indecisionist can accommodate this intuition by admitting a distinction between what is *semantically expressed* by an utterance and what is *pragmatically asserted* on an occasion of utterance. In virtue of accepting different criteria of application for ‘wrong’, the individuals assert different propositions when they utter sentences involving ‘wrong’, even though what is semantically expressed by their utterances may be shared. In addition, it is plausible that in a conversation between these individuals, their discourse context will be *defective* in Stalnaker’s technical sense of the term. In defective contexts, it is natural to have an intuition that the speakers are merely talking past one another; communication “breaks down”. On defective contexts, see, e.g., Stalnaker, 2002.

⁵⁰*Objection.* One might claim that, if moral indecisionism is true, then agents who are idealized in these ways will not come up with “ultimate moral views” at all. Instead, after reflecting on their inferential and judgmental dispositions, they’ll be *undecided* about the application conditions for ‘wrong’. If so, then either moral indecisionism implies that these situations are impossible (rational agents will always agree in their indecision about ultimate moral views), or else it implies that the disputes are rationally resolvable (they can rationally come to agree to be undecided in exactly the same ways).

Reply. I reject the first claim of this objection. It’s false that, if indecisionism were true, then the agents *would* agree in their indecision. In place of that claim, I accept that if moral indecisionism were true, then

It seems to me that the error theorist must claim that, in any disagreement that nearly satisfies this description, either it *is* rationally resolvable, or else it is not a real disagreement. In order to explain why, I'll show two things: first, on the supposition that the dispute is rationally irresolvable, it becomes a non-substantive disagreement; second, on the supposition that it is a substantive disagreement, the dispute becomes rationally resolvable.

First, suppose that there is no way rationally to resolve the disagreement. We cannot ask our ideal speakers to reflect again on their inferential and judgmental dispositions and expect them to arrive at different analyses of 'wrong'. They've systematized their discourse; each has correctly analyzed 'wrong'; each is in "reflective equilibrium". It's a decided fact that, in one individual's mouth, 'wrong' means 'doesn't maximize', whereas in the other speaker's mouth 'wrong' means 'treats someone as a mere means'. Then it seems that they must be talking past each other; they are using the same word with different meanings (or else one of them is confused about what he actually means).

Second, suppose that the individuals really are engaged in a substantive disagreement: what the one expresses with 'wrong' is the denial of what the other expresses with 'wrong'. Then it seems that, if we were to ask them to reflect again on their inferential and judgmental dispositions, at least one of them would come to realize that he just hadn't quite reached reflective equilibrium the first time. But then whichever individual comes to this realization

the agents *might* agree in their indecision about application conditions for 'wrong', but they might also agree on application conditions, or they might instead disagree in application conditions. And the latter happens to be true in the example that I have described. In order to understand why this is true, consider the following facts. (i) There are two options for any agent who seeks conditions of application for a term: either to accept determinate conditions of application for the term, or not to accept any. (ii) When a term is semantically undecided, each reasonable resolution of the indecision corresponds to a determinate set of application conditions, and thus each of these sets of application conditions is *reasonable*: the individual won't be breaking any "laws of rationality" in accepting them. (iii) As a consequence, neither option that I mentioned in (i) is "rationally forced" upon an agent who seeks to analyze a term: both are "rationally permissible". (iv) Thus a rational agent might take either of these options. And so it's false that if moral indecisionism were true, then the agents *would* agree to be undecided: depending upon which options they choose, the agents might agree in their indecision, or they might agree on ultimate moral views, or they might disagree on ultimate moral views.

ought to agree with his opponent on whether ‘wrong’ applies to x . And this would constitute a way of rationally resolving the issue.

So we have found another reason to believe moral indecisionism instead of the error theory. Roughly, this is because moral indecisionism has a good explanation of how there may be genuine, rationally irresolvable moral disagreements between individuals who are in reflective equilibrium, and who adopt very different fundamental moral views.⁵¹

5.7 Conclusion

I conclude that the error theory, which holds Theses 1IA, 2EA, and 3EA, is not true. It implies that moral discourse is relatively semantically decided, but there is evidence to the effect that moral discourse is radically semantically undecided. This was the primary goal of the chapter.

I now mention a slight weakening of my conclusion. Suppose you were an error theorist before you read this chapter. Suppose that you have not been convinced by the considerations I presented in favor of the main premise of my argument. Suppose further that the reason you are unconvinced is that these considerations are largely *a priori*, and do not involve any sort of systematic empirical investigation into our usage of moral terminology. This, it seems to me, is a natural and reasonable reaction to my arguments. However, even if you satisfy these suppositions, you are not justified in ignoring my objection to the error theory. For it is not obvious that moral language is semantically decided; it is (or, I claim, it ought to be) a controversial hypothesis. In addition, I reject the idea that there is some kind of initial theoretical presumption in favor of the idea that moral language is relatively semantically decided. So if you are justified in rejecting moral indecisionism because of empirical reasons, then you must present the empirical data that would disprove moral in-

⁵¹It is interesting to note that a newfangled relativism, such as the kind argued against in Cappelen and Hawthorne (2009), would also apparently be able to provide a neat explanation of this phenomenon.

decisionism. Therefore, my weakened conclusion is this: the burden of proof is on the error theorist to justify the claim that moral language is sufficiently semantically decided.⁵²

Another main goal of the chapter was to explicate a novel meta-ethical standpoint, which has received very little attention in the literature. I hope to have made it plausible that moral indecisionism is a rival to the major theories about the semantics of moral discourse. Roughly, the reason why it is a competitor is that it is able to explain at least as much of what is puzzling about moral discourse as any other views; further, it is able to explain some phenomena that may be puzzling on other views. Moral indecisionism also admits of precise explication, depending upon how you want to develop it. That is further reason to admit it as a plausible rival view.

I want to close the chapter, and the dissertation, with a few more thoughts about how moral indecisionism may be developed. In what follows, I suggest at some answers to the following questions:

How should the moral indecisionist view the practices of normative ethicists?

How should the moral indecisionist continue to participate in moral discourse?

Must the moral indecisionist accept some kind of “error theory” about ordinary moral discourse?

My thoughts about these questions ought, however, to be viewed as elaborations and extensions of the view. Moral indecisionism *per se* is a claim about the correct way of modeling the semantics of moral discourse: it says that the best semantic model of moral language treats moral terms as radically semantically undecided.

First, it may be natural for the moral indecisionist to view a lot of what normative ethicists do as slightly confused. If normative ethics is characterized as the program of the search for the one true necessary theory of right, wrong, good, bad, justice, and injustice,

⁵²I would reply similarly to meta-ethical realists who are unimpressed because my arguments fail to provide much empirical justification for a hypothesis that requires it.

then the moral indecisionist should view normative ethics as misguided. On moral indecisionism, there will be no complete moral system that can satisfy that demand. Instead, there will be a variety of “equally correct” moral systems, and none of them can be put determinately above the rest.

But even though normative ethics looks confused in one respect, the moral indecisionist does not mean to imply that it is a waste of time to participate in the program. Rather, it is an extremely important part of philosophy, since we want to know, roughly speaking, what are the reasonable resolutions of moral discourse. And we may view the program of normative ethics as the pursuit of these reasonable resolutions. Whenever a philosopher puts forward a complete ethical system as true, the moral indecisionist takes him to be putting it forward as a joint reasonable resolution of our moral concepts.

In addition, the moral indecisionist has a good interpretation of the classical methodology of normative ethics. The “classical methodology” is the method of trying to come into reflective equilibrium: to systematize, *via* various theoretical virtues, our one-off moral intuitions with general *prima facie* plausible moral principles. On moral indecisionism, this is a useful methodology *not* for seeking the truth about morality; instead, it is the way to determine what are the reasonable resolutions of our moral terminology. And so, the research program of normative ethics may be retained, if slightly reunderstood.

Second, the moral indecisionist has to come to a decision on how to continue to participate in moral discourse. The recent literature on “fictionalism” about ethical discourse suggests a variety of options.⁵³ I, however, am interested in a specific question: how should the moral indecisionist use moral language in conversation? In my opinion, in actual moral discourse, we often use moral statements to communicate our subjective attitudes of approval and disapproval. We also often use moral statements in order to goad others into

⁵³Cf. Joyce, 2001 and 2005.

doing what they aren't initially inclined to do. The moral indecisionist clearly can retain both of these kinds of moralizing.

Furthermore, the moral indecisionist may use moral statements to communicate about the external world. For conversational participants may resolve the meanings of moral terms one way or another in context, just as we may shift standards of precision for vague terminology one way or another, depending upon our interests. If it is clear in context just how the semantic indecision is being resolved, then speakers may update their contextual presuppositions about the meanings of moral terms, and thus communicate determinately, using these updated presuppositions.⁵⁴ So it is not like a moral indecisionist must give up on moral discourse in general. Indeed, I think that would be ill-advised.

Finally, I think the moral indecisionist has to accept some sort of error theory about moral discourse. According to the indecisionist, if people utter moral sentences with the belief that what they are uttering is true, then they are committing errors. For none of these statements is super-true; instead, each of them ends up being indeterminate. Of course, how often this error occurs depends on how far-reaching the indeterminacy in moral discourse is: if it is universal or near-universal, then there may be a strong case that it happens very often.

However, by appealing to the distinction between what is asserted on an occasion of utterance and what is semantically expressed on an occasion of utterance, the indecisionist's error theory need not be too extreme. For it is very plausible that we often communicate facts about our subjective states through moral statements, and the indecisionist can allow for this. And I have already noted that it is natural for the indecisionist to view moral speakers as resolving the meanings of moral terms in context in order to make determinate assertions. So, in my opinion, the moral indecisionist can retain much of ordinary moral discourse.

⁵⁴Cf. Lewis, 1979.

BIBLIOGRAPHY

- [1] Ayer, A.J. (1952), *Language, Truth and Logic*, New York: Dover.
- [2] Balaguer, Mark. (2011), “Bare Bones Moral Realism and the Objections from Relativism”, in *A Companion to Relativism* (ed. Hales), pp. 368-390.
- [3] Bedard, Katherine. (1993), “Partial Denotations of Theoretical Terms”, *Nous*, 27:4, pp. 499-513.
- [4] Bennett, Karen. (2004), “Global Supervenience and Dependence”, *Philosophy and Phenomenological Research*, 68: 3, pp. 501-529.
- [5] Blackburn, Simon. (1984), “Supervenience Revisited”, in *Exercises in Analysis: Essays by Students of Casimir Lewy*, Ian Hacking (ed.), Cambridge University Press.
- [6] Boghossian, Paul. (2011), “Three Kinds of Relativism”, in *A Companion to Relativism* (ed. Hales), pp. 53-69.
- [7] — ——. (2006), “What is Relativism?”, in *Truth and Realism*, Greenough and Lynch (eds.), pp. 13-37, OUP.
- [8] Bricker, Phillip. (1996), “Properties”, in *The Encyclopedia of Philosophy*, pp. 469-73.
- [9] Burgess, J.P. (2007), “Against Ethics”, *Ethical Theory and Moral Practice*, 10, pp. 427-439.
- [10] Cappelen, Herman; and John Hawthorne. (2009), *Relativism and Monadic Truth*, OUP.
- [11] Daly, Chris; and David Liggins. (2010), “In defence of error theory”, *Philosophical Studies*, 149, pp. 209-230.
- [12] Dreier, James. (2010), “Mackie’s Realism: Queer Pigs and the Web of Belief”, in Joyce and Kirchin, pp. 71-86.
- [13] — ——. (2005), “Moral Relativism and Moral Nihilism”, in *Oxford Handbook of Ethical Theory*, Copp (ed.), pp. 240-64, OUP.
- [14] — ——. (1990), “Internalism and Speaker Relativism”, *Ethics*, 101:1, pp. 6-26.
- [15] Feldman, Fred. (2010), *What is This Thing Called Happiness?*, Oxford: OUP.

- [16] ——. (2005), “The Open Question Argument: What it Isn’t; and What it Is”, *Philosophical Issues 15, Normativity*, pp. 22-43.
- [17] ——. (1979), “Review of Mackie, *Ethics: Inventing Right and Wrong*, *The Philosophical Review*, pp. 134-7.
- [18] Finlay, Stephen. (2008), “The error in the error theory”, *Australasian Journal of Philosophy*, 86:3, pp. 347-69.
- [19] ——. (2011), “Errors upon errors: a reply to Joyce”, *Australasian Journal of Philosophy*, 89:3, pp. 535-47.
- [20] Frege, Gottlob. (1892), “On Sense and Reference”, in Martinich 2001, pp. 199-211.
- [21] Garner, Richard. (2007), “Abolishing Morality”, *Ethical Theory and Moral Practice*, 10, pp. 499-513.
- [22] ——. (1990), “On the genuine queerness of moral properties and facts”, *Australasian Journal of Philosophy*, 68:2, pp. 137-46.
- [23] Geach, Peter. (1965), “Assertion”, *Philosophical Review*, 74, pp. 449-65.
- [24] ——. (1960), “Ascriptivism”, *Philosophical Review*, 69, pp. 221-5.
- [25] Grice, H.P. (1975), “Logic and Conversation”, in Martinich (2001).
- [26] Hare, R.M. (1999), “Objective Prescriptions”, in his *Objective Prescriptions and Other Essays*, pp. 1-18, Oxford: Clarendon Press.
- [27] ——. (1985), “Ontology in Ethics”, in his *Essays in Ethical Theory*, pp. 82-98, Oxford: Clarendon Press.
- [28] ——. (1981), *Moral Thinking*, Oxford: Clarendon Press.
- [29] ——. (1967), “The Promising Game”, in his *Essays in Ethical Theory*, pp. 131-144, Oxford: Clarendon Press.
- [30] ——. (1957), “ ‘Nothing Matters’ ”, in his *Applications of Moral Philosophy*, pp. 32-47, University of California Press.
- [31] ——. (1952), *The Language of Morals*, Oxford: Clarendon Press.
- [32] Harman, Gilbert. (2000), *Explaining Value and Other Essays in Moral Philosophy*, Oxford: Clarendon Press.
- [33] ——. (1975), “Moral Relativism Defended”, in his *Explaining Value*, pp. 3-19, Oxford: Clarendon Press.
- [34] ——. (1978), “What is Moral Relativism?”, in his *Explaining Value*, pp. 20-38, Oxford: Clarendon Press.

- [35] Harman, Gilbert; and Thomson, Judith Jarvis. (1996), *Moral Relativism and Moral Objectivity*, Blackwell.
- [36] Horgan, Terence and Timmons, Mark. (1991), “New wave moral realism meets moral twin earth”, *Journal of Philosophical Research*, 16, pp. 447-465.
- [37] Huemer, Michael. (2005), *Ethical Intuitionism*, Oxford: OUP.
- [38] Hume, David. (1777), *Enquiries concerning Human Understanding and concerning the Principles of Morals*, Nidditch (ed.), Oxford: Clarendon Press.
- [39] ———. (1740), *A Treatise of Human Nature*, Nidditch and Selby-Bigge (eds.), Oxford: Clarendon Press.
- [40] Jackson, Frank. (1998), *From Metaphysics to Ethics: A Defence of Conceptual Analysis*, Oxford: OUP.
- [41] ———. (1974), “Defining the Autonomy of Ethics”, *Philosophical Review*, 83: 1, pp., 88-96.
- [42] Janis, Allen. (2010), “Conventionality and Simultaneity”, in the *Stanford Encyclopedia of Philosophy*.
- [43] Joyce, Richard. (2011), “The Error in ‘The Error in the Error Theory’ ”, *Australasian Journal of Philosophy*, 89:3, pp. 519-534.
- [44] ———. (2010), “Patterns of Objectification”, in Joyce and Kirchin, pp. 35-53.
- [45] ———. (2008), “Replies”, *Philosophy and Phenomenological Research*, 77:1, pp. 245-267.
- [46] ———. (2007a), “Moral Anti-Realism”, in the *Stanford Encyclopedia of Philosophy*.
- [47] ———. (2007b), “Morality, schmorality”, in *Morality and Self-Interest*, P. Bloomfield (ed.), pp. 51-75, Oxford: OUP.
- [48] ———. (2006), *The Evolution of Morality*, Cambridge: MIT Press.
- [49] ———. (2005), “Moral Fictionalism”, in *Fictionalism in Metaphysics*, Kalderon (ed.), Oxford: OUP.
- [50] ———. (2001), *The Myth of Morality*, Cambridge: CUP.
- [51] Joyce, Richard; and Simon Kirchin (eds.). (2010), *A World Without Values*, *Philosophical Studies* 114.
- [52] Kim, Jaegwon. (1993), “Concepts of supervenience”, in his *Supervenience and Mind*, Cambridge: Cambridge University Press.
- [53] Kirchin, Simon. (2010), “A Tension in the Moral Error Theory”, in Joyce and Kirchin, pp. 167-82.

- [54] Lewis, David. (2005), "Quasi-Realism is Fictionalism", in *Fictionalism in Metaphysics*, Kalderon (ed.), pp. 314-21, Oxford: OUP.
- [55] ———. (1995), "Should a materialist believe in qualia?", in his *Papers in Metaphysics and Epistemology*, Cambridge: CUP.
- [56] ———. (1993), "Many, but almost one", in his *Papers in Metaphysics and Epistemology*, pp. 164-82, Cambridge: CUP.
- [57] ———. (1989), "Dispositional Theories of Value", in his *Papers in Ethics and Social Philosophy*, pp. 68-93, Cambridge: CUP.
- [58] ———. (1986), *On the Plurality of Worlds*, Blackwell.
- [59] ———. (1984), "Putnam's Paradox", in his *Papers in Metaphysics and Epistemology*, pp. 56-77, Cambridge: CUP.
- [60] ———. (1983), "New work for a theory of universals", in his *Papers in Metaphysics and Epistemology*, pp. 8-55, Cambridge: CUP.
- [61] ———. (1982), "Logic for equivocators", in his *Papers in Philosophical Logic*, pp. 97-110, Cambridge: CUP.
- [62] ———. (1980), "Index, Context, and Content", in his *Papers in Philosophical Logic*, pp. 21-44, Cambridge: CUP.
- [63] ———. (1979), "Scorekeeping in a Language Game", in his *Philosophical Papers, Vol. 1*, pp. 233-49.
- [64] ———. (1974), "Radical Interpretation", in his *Philosophical Papers, Vol. 1*, pp. 108-118.
- [65] ———. (1970), "How to Define Theoretical Terms", in his *Philosophical Papers, Vol. 1*, pp. 78-95.
- [66] ———. (1969), *Convention: A Philosophical Study*, Blackwell.
- [67] Mackie, J.L. (1985), *Persons and Values: Selected Papers, Vol. II*, Oxford: Clarendon Press.
- [68] ———. (1980), *Hume's Moral Theory*, Routledge, Kegan, and Paul Ltd.
- [69] ———. (1977), *Ethics: Inventing Right and Wrong*, Penguin Books.
- [70] ———. (1969), "Aesthetic Judgments – A Logical Study", in Mackie (1985), pp. 60-76.
- [71] ———. (1946), "A refutation of morals", *Australasian Journal of Philosophy*, 24:1, pp. 77-90.
- [72] Martinich, A.P. (2001), *The Philosophy of Language, 4th Edition*, Oxford: OUP.

- [73] Nolan, Daniel; Restall, Greg; and West, Caroline. (2005), "Moral fictionalism versus the rest", *Australasian Journal of Philosophy*, 83:3, pp. 307-330.
- [74] Phillips, David. (2010), "Mackie on Practical Reason", in Joyce and Kirchin, pp. 87-99.
- [75] Pigden, Charles R. (2010), "Nihilism, Nietzsche, and the Doppelganger Problem", in Joyce and Kirchin, pp. 17-34.
- [76] Potts, Christopher. (2007), "Into the Conventional-Implicature Dimension", *Philosophy Compass*, 2:4, pp. 665-679.
- [77] Prior, A.N. (1960), "The Autonomy of Ethics", *Australasian Journal of Philosophy*, 38, pp. 199-206.
- [78] Quine, W.V. (1980), "Logic and the Reification of Universals", in *Logico-Philosophical Essays*, pp. 102-29, Cambridge: Harvard University Press.
- [79] ——. (1980), "On What There Is", *The Review of Metaphysics*, 2:5, pp. 21-38.
- [80] Salmon, Nathan. (1998), "Nonexistence", *Nous*, 32: 3, pp. 277-319.
- [81] Schiffer, Stephen. (2003), *The Things We Mean*, Oxford, OUP.
- [82] ——. (1990), "Meaning and Value", *The Journal of Philosophy*, 87: 11, pp. 602-14.
- [83] Schroeder, Mark. (2010), *Noncognitivism in Ethics*, Routledge.
- [84] ——. (2008), *Being For*, Oxford: OUP.
- [85] ——. (2007), *Slaves of the Passions*, Oxford: OUP.
- [86] Searle, John. (1964), "How to Derive 'Ought' from 'Is' ", *Philosophical Review*, 73: 1, pp. 43-58.
- [87] ——. (1962), "Meaning and Speech Acts", *Philosophical Review*, 71, pp. 423-32.
- [88] Shafer-Landau, Russ. (2006), "Ethics as Philosophy: A Defense of Ethical Nonnaturalism", in Horgan and Timmons, pp. 209-32.
- [89] ——. (2005), "Error Theory and the Possibility of Normative Ethics", *Philosophical Issues*, 15, Normativity, pp. 107-20.
- [90] ——. (2003), *Moral Realism: A Defence*, Oxford: OUP.
- [91] Sider, Theodore. (1999), "Global Supervenience and Identity across Times and Worlds", *Philosophy and Phenomenological Research*, 59, pp. 913-937.
- [92] Simons, Mandy. (2003), "Presupposition and Accommodation: Understanding the Stalnakerian Picture", *Philosophical Studies*, 112:3, pp. 251-278.

- [93] Sinnott-Armstrong. (2010), “Mackie’s Internalisms”, in Joyce and Kirchin, pp. 55-70.
- [94] ——. (2006), *Moral Skepticisms*, Oxford: OUP.
- [95] Smith, Michael. (2010), “Beyond the Error Theory”, in Joyce and Kirchin, pp. 119-139.
- [96] ——. (1994), *The Moral Problem*, Blackwell.
- [97] Soames, Scott. (2008), “The Gap Between Meaning and Assertion: Why what we say often differs from what our words literally mean”, in *Asserting, Meaning, and Implying* (eds. Hackl and Thornton).
- [98] ——. (2005), *Reference and Description: The Case against Two-Dimensionalism*, Princeton University Press.
- [99] ——. (2002), *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*, OUP.
- [100] ——. (1999), *Understanding Truth*, Oxford: OUP.
- [101] ——. (1989), “Presupposition”, in *Handbook of Philosophical Logic, Vol. IV*, Gabbay and Guenther (eds.), pp. 553-616.
- [102] Stalnaker, Robert. (2002), “Common Ground”, *Linguistics and Philosophy* 25, pp. 701-721.
- [103] ——. (1999), *Context and Content*, Oxford: OUP.
- [104] ——. (1987a), *Inquiry*, Cambridge: MIT Press.
- [105] ——. (1987b), “Semantics for Belief”, in his 1999, pp. 117-129.
- [106] ——. (1978), “Assertion”, in his 1999, pp. 78-95.
- [107] ——. (1974), “Pragmatic Presuppositions”, in his 1999, pp. 47-62.
- [108] ——. (1973), “Presuppositions”, *Journal of Philosophy*, 2, pp. 447-57.
- [109] ——. (1970), “Pragmatics”, in his 1999, pp. 31-46.
- [110] Strawson, P.F. (1950), “On Referring”, in Martinich 2001, pp. 228-242.
- [111] Stevenson, C.L. (1944), *Ethics and Language*, Oxford: OUP.
- [112] van Inwagen, Peter. (1977), “Creatures of Fiction”, in his *Ontology, Identity, and Modality: Essays in Metaphysics*, pp. 37-56, Cambridge: Cambridge University Press.
- [113] Yablo, Stephen. (2006), “Non-Catastrophic Presupposition Failure”, in *Content and Modality* (eds. Byrne and Thomson).