

ASSIMILATIVE MORAL REALISM

A CRITIQUE
OF
ASSIMILATIVE MORAL REALISM

By

KEN YASENCHUK, B.A., M.A.

A Thesis

Submitted to the School of Graduate Studies

In Partial Fulfilment of the Requirements

For the Degree

Doctor of Philosophy

McMaster University

(c) Copyright by Ken Yasenchuk, September 1995

DOCTOR OF PHILOSOPHY (1995)
(Philosophy)

McMASTER UNIVERSITY
Hamilton, Ontario

TITLE: A Critique of Assimilative Moral Realism

Author: Ken Yasenchuk, B.A. (University of Calgary)

M.A. (University of Calgary)

SUPERVISOR: Professor Mark Vorobej

NUMBER OF PAGES: vi, 180

ABSTRACT

David Brink, in his book Moral Realism and the Foundations of Ethics, and other writers, have recently offered powerful new arguments for a form of moral realism that sees moral inquiry as being "on par" with scientific inquiry in many important epistemological and metaphysical respects. I call this theory "Assimilative Moral Realism" (AMR). AMR is marked by naturalism about moral facts, and by empiricism about moral knowledge. Moral facts are held to be facts about properties that are constituted by, and supervene upon, complexes of ordinary natural and social properties, just as certain scientific facts are facts about the macrophysical properties that are realized by certain microphysical bases. Moral beliefs are held to be justifiable in the way certain scientific beliefs ostensibly are: via certain explanatory inferences and indirect empirical tests of moral principles. Further, traditional semantical objections to naturalistic moral realism can be avoided, it is supposed, because semantical tests for the existence of moral facts can plausibly be rejected, just as they can be in the scientific case. In this thesis, I make a case for rejecting this assimilative theory. I provide a number of technical reasons to suppose that moral principles cannot be tested empirically using the Duhemian method advocated by AMR theorists. I also give a critique of their main arguments for the relevance and indispensability of assumptions about moral facts to explanations of moral judgments and other items, and conclude that their arguments establish neither the explanatory relevance nor necessity of moral assumptions. Finally, in place of the AMR theory of supervening moral properties, I offer an account of moral supervenience claims - an account that is itself compatible with AMR only if there are sound moral explanations of certain nonmoral facts.

ACKNOWLEDGMENTS

I thank my supervisor Professor Mark Vorobej for his extremely helpful comments on numerous drafts of this thesis, and also for his words of encouragement. I also thank Professor Douglas Odegard and Professor Evan Simpson for providing very useful remarks on earlier drafts of the initial thesis proposal, on earlier drafts of individual chapters, and on the final draft of the entire thesis. I am also indebted to certain anonymous referees at Dialogue and The Southern Journal of Philosophy. Their input on earlier versions of some of the arguments here resulted in revisions to Chapters Four, Five and Six. Finally, I am indebted to Janice L. Redlin, whose editing skills helped me to submit to my committee a very clean copy, very quickly. I also thank her for her understanding, patience, and for encouraging me to persevere during the darkest moments of the thesis-writing process.

CONTENTS

	PAGE
Abstract.....	iii
Acknowledgments.....	iv
Chapter One: Introduction.....	1
1.1. Introduction.....	1
1.2. Background.....	3
1.3. What is AMR?.....	7
1.4. Problems with the Account.....	16
Chapter Two: The Burden of Argument.....	23
2.1. Introduction.....	23
2.2. What is Brink's Argument for the Onus Shift?.....	23
2.3. The PC Argument.....	31
2.4. The PC* Argument.....	41
2.5. Conclusion.....	43
Chapter Three: Are Moral Theories Empirically Testable?.....	45
3.1. Introduction.....	45
3.2. The Case for Testability.....	45
3.3. MT and the Independence Requirement.....	50
3.4. Can <u>Any</u> Moral Principles be Tested Using Duhem's Method?.....	61
3.5. Conclusion.....	64
Chapter Four: Explanatory Relevance.....	65
4.1. Introduction.....	65
4.2. Harman's Case for Irrelevance.....	67
4.3. The Counterfactual Argument.....	70
4.4. Jobe's Application of the Test.....	80
Chapter Five: Explanatory Necessity.....	83
5.1. Introduction.....	83
5.2. The Generality Argument.....	83
5.3. The Argument from Moral Conversions.....	99
5.4. Problems with Werner's Argument.....	104
5.5. Boyd's Response to the Moral Input Objection.....	112
Chapter Six: AMR and Supervenience.....	120
6.1. Introduction.....	120
6.2. Brink on Moral Constitution and Supervenience.....	122
6.3. Necessity, Supervenience, Constitution: Some Distinctions.....	126
6.4. Modal Diversity and Brink.....	135
6.5. Nonmoral Macrophysical Properties.....	136
6.6. Moral Properties.....	141
6.7. The Psychological Functionalist Option.....	154

	PAGE
6.8. Why is the Absence of Modal Parity Important?.....	159
6.9. Conclusion.....	163
Chapter Seven: Conclusion.....	165
Bibliography.....	173

CHAPTER ONE

INTRODUCTION

1.1: Introduction

David Brink and other philosophers have recently presented powerful new arguments for a thoroughly naturalistic form of moral realism.¹ These writers reject traditional objections to moral realism which regard the facts that that theory affirms as ontologically "queer" and epistemically inaccessible; on their view, moral facts are no more queer or inaccessible to cognizers than those that are affirmed by other "recognizably realist" disciplines, such as physics or empirical psychology. Moral facts are not ontologically queer because moral properties can be construed as akin to ordinary nonmoral macrophysical (and psychological functional) properties (both kinds of properties can be regarded as supervening upon, and constituted by, specific constellations of underlying properties). They are not epistemically inaccessible since knowledge about them can be obtained a posteriori via the same cognitive and perceptual means employed in science. Thus moral facts are epistemologically and ontologically "on par" with the facts and properties affirmed by scientific theories. Given this parity, we can reject realism about morals only if we are willing to give up realism about science: the two stand or fall together.² Brink calls his theory "moral realism," but it has earned a number of different labels in the literature, depending on which of its salient features writers have wanted to stress: e.g., "new wave moral realism," "synthetic ethical naturalism," and simply "moral

¹ Brink 1989; Sturgeon 1984; 1986; also Boyd 1988.

² David McNaughton calls this the "companions in guilt" strategy (1988: 63).

realism."³ I shall call it "assimilative moral realism" (AMR) because it is ostensibly based on the "assimilation" of moral and empirical inquiry. Other forms of moral realism might be based on different assimilative considerations. For example, one might assimilate, say, moral truths to truths in mathematics. I shall reserve the "assimilative" label for accounts that assimilate moral and empirical (esp., scientific) inquiry.

AMR is a metaphysical thesis (as opposed to, say, a purely semantic one); it holds that "there are moral facts and true moral claims whose existence and nature are independent of our beliefs about what is right and wrong" (Brink 1989: 7). A number of features make the theory worth examining from a metaethical point of view. First, although neither general realism, nor moral realism, needs to be construed as a metaphysical thesis, metaphysical versions of extra-moral realism are prominent enough in the literature on that topic to make the extension of that theory - which is ostensibly what AMR represents - to additional areas worthy of examination.⁴ Second, AMR theorists appear to have plausible answers to many familiar objections to metaethical naturalism, such as those that cite failures of semantic reduction, and the existence of an "is/ought gap." Third, the theory promises to provide for the idea that moral demands impose on agents (and communities of agents) from without (thus refuting relativism). If AMR is true, then moral facts and properties could plausibly be regarded, pace John Mackie (1977), as forming part of the fabric of the world. It thus promises to be a plausible and metaethically satisfying account of the objectivity of morals.

³ See Horgan and Timmons 1992; 1990-1991.

⁴ Realism about an area might be construed as a semantic claim, or a claim about the concept of truth for a class of propositions, or as a claim about their truth values (esp. that some of a class of propositions are true). But see, for example, Devitt (1984: 34), who defines common sense and scientific realism as the thesis that "Tokens of most current common sense, and scientific, physical types objectively exist independently of the mental."

This work is a critique of AMR. Since the most sustained and systematic articulation and defence of AMR appears in Brink's Moral Realism and the Foundations of Ethics, the arguments of that book will be my main target, but I will also target the work of theorists whose work bears on the viability of AMR. I will argue that, notwithstanding the profundity of recent arguments for AMR, Brink and other AMR theorists have not provided sufficient reason to accept that theory.

In this Chapter, I begin by situating AMR in its proper context as a theory of the objectivity of moral values by giving a brief description of previous work in the area (Section 1.2). I then give a brief overview of AMR and of my objections to it (in Sections 1.3 and 1.4 respectively). Subsequent chapters contain more detailed analyses of particular aspects of AMR.

1.2: Background

Although moral realism enjoyed a brief period of respectability under G. E. Moore's (1903) tutelage during the very early part of this century, the metaethical landscape since that period has been marked by a prevailing moral antirealism. In fact, Moore himself is at least partly responsible for this: his "open question" argument convincingly ruled out naturalistic moral realism, and his own nonnaturalist moral realism seemed to imply an "occult," and hence unconvincing, view of moral facts. Thus, in 1946, John Mackie could write that elements of his skeptical antirealism about moral value were "the stock instruments of all modern discussions of morals" (1946: 77), and this was still largely true in 1977, when his Ethics: Inventing Right and Wrong appeared. Moral philosophy remained dominated by an empiricist noncognitivism inspired by Hume (more recently by logical positivism), and by a prevailing concern with the analysis of meaning. Accepting Hume's denial of moral facts, and inspired by a general philosophical concern with language, philosophers had concentrated on developing a theory of moral language.

These theories had become increasingly sophisticated. A. J. Ayer's suggestion (1952: 107) that moral utterances were mere expressions of emotion was rejected in favour of accounts that attempted to explain conflict between judgments, apparent reasoning about ethics, and prescriptivity (e.g., Stevenson 1944; Hare 1963, 1981). Meanwhile, efforts were made to avoid the metaethical skepticism implicit in Hume, and thought to be implicit in empiricism generally. These roughly followed Kant, and included attempts to ground morality and establish the prescriptivity of at least some moral claims on a particular conception of ethical rationality (Baier 1958), commitment (Searle 1964), analyses of action (e.g., Gewirth 1978), or coherence of conviction (Rawls 1973; Daniels 1979). These accounts may make certain sorts of behaviour rational, and other sorts irrational (assessed against certain normative criteria or canons of rationality), but they all share a received skepticism about the idea that any complete description of the world must include the postulation of moral facts.

It was in concert with this prevailing metaethical view that Mackie denied that moral facts are part "of the fabric of the world," and Gilbert Harman insisted that the observational untestability of moral judgments marks them as targets for skepticism.⁵ Both writers ostensibly attempted to summarize, in a concise way for those new to metaethics, a rationale for the prevailing view. However, since 1977, the philosophical landscape in the area has changed considerably. A flood of work on, and often advocating, moral realism, has appeared. Predating Brink are papers by David Wiggins (1976), Mark Platts (1980), John McDowell (1979, 1981, 1985), and Sabina Lovibond (1983). There have also been two major conferences that have resulted in published collections of papers on moral realism. One, the 1986 Spindel Conference on "Moral Realism," produced several papers, almost all of which argue for some version of realism

⁵ Mackie 1977; Harman 1977.

about moral value (Gillespie 1986). More recently, the 1990 Spindel Conference on "Moral Epistemology" resulted in papers offering further, largely sympathetic, refinements (Timmons 1990a). Other collections that deserve mention are anthologies: David Copp (ed.) Morality, Reason and Truth (1984), Ted Honderich (ed.) Morality and Objectivity (1985), and Geoffrey Sayre-McCord (ed.) Essays on Moral Realism (1988). The first contains Sturgeon's landmark "Moral Explanations," the second contains discussions of Mackie's Ethics, and the third includes some of the most important papers to date (Wiggins' 1976 paper, some advancements on earlier work by John McDowell, and a lengthy paper by Richard Boyd that marshals his research in the philosophy of science in support of moral realism).⁶ As well, in addition to Brink's book, several books have appeared that systematically defend versions of moral realism: David McNaughton's Moral Vision (1988), Panayot Butchvarov's Skepticism in Ethics (1989), and Alan Gilbert's Democratic Individuality (1990). Among these, Brink's is the most thorough, but McNaughton's book is significant because it provides important alternatives to Brink's account, particularly on the issue of moral motivation. Butchvarov offers a careful and extensive rehabilitation of G. E. Moore's ethical theory. Gilbert's work is a sweeping, conversational, defence of moral progress.

The realist writers broadly divide into two camps. The first camp can be described as "nonassimilative" because it rejects the idea that objectivity of morality must be assessed via a comparison of moral and scientific inquiry. For example, Nagel (1980; 1986), Butchvarov (1989), and Lemos (1994), evidently think that moral facts can meet acceptable standards of objectivity without being "real" or "objective" in the way scientific facts ostensibly are. Others, notably Wiggins (1976), Platts (1980), McDowell (1979, 1981, 1985), and Lovibond (1983), seem to make a kind of "transcendental" case for moral

⁶ Boyd, incidentally, promises to expand his arguments in a forthcoming book.

objectivity. In their view, scrutinizing moral commitments requires that they be adequately understood, and this in turn requires that we adopt the participant's point of view rather than the scientific point of view; but from there, our commitments seem like rational, perspicuous responses to the world; any adequate understanding of morality must therefore include the recognition of their objectivity.

By contrast, members of the second camp - that represented by AMR theorists such as Brink (1989), Sturgeon (1984; 1986), and Boyd (1988) - regard science as the proper object of comparison, but contend that moral inquiry compares favourably against that area of inquiry. It advocates an empiricist, naturalist, "scientific" moral realism according to which moral facts and knowledge are thoroughly unexceptional. As such, AMR contrasts also with Moore's position in Principia Ethica (1903), and also with epistemologically "rationalist" versions of moral realism.⁷ It accepts the core idea on which early antinaturalist arguments were based (viz., that moral claims are not equivalent in meaning to naturalistic claims, and that it is impossible to derive moral statements via an "extension-fixing" analysis of the meaning of moral property ascriptions), but it refuses to regard this as a sufficient reason to reject naturalistic moral realism. On the AMR account, moral property ascriptions do not depend for their justification entirely on semantic analyses any more than scientific and perceptual claims do; rather, they are justified in the same way scientific and perceptual claims are (on what is assumed to be a plausible account of the justification of such claims). It is because AMR eschews the antinaturalism and anti-empiricism of Moore's account, while also advancing novel arguments for moral realism (arguments that are thought by AMR theorists to be capable of avoiding the

⁷ See, e.g., Audi 1990. Audi contends that there are objectively existing moral properties, but denies that they are apprehended by empirical means.

problems that plagued earlier forms of ethical naturalism), that AMR is sometimes called "New Wave Moral Realism" (Horgan and Timmons 1992).

1.3: What is AMR?

AMR implies both an ontological or metaphysical thesis (about the existence and nature of moral facts or properties), and an epistemological thesis (about how those facts are discerned by moral inquirers, and about how statements affirming the existence of moral facts are justified). In this section, I provide a brief description of these theses.

In order to appreciate Brink's account of the nature of moral facts, we need to understand certain of the technical terms that Brink employs in giving his account. On Brink's view, if a certain basic property (or set of properties) G actually realizes a property F and F could not have been realized by any other property(ies), then F is identical to G; but if G realizes F even though F could have been realized in another way, then F is constituted by, but not identical to, G; and if G produces F as a matter of nomological necessity, then F supervenes on G (1989: 160-162). Brink claims that moral properties are akin to the ordinary upper-level properties affirmed by science: just as macrophysical properties (such as the property of being water or of possessing water's empirically-discernible qualities) can often be said to be constituted by, and to supervene upon (but not be identical to), complexes of more basic (e.g., microphysical) properties, so moral properties can be said to be constituted by, and to supervene upon (but not be identical to), complexes of natural and social properties. Brink presents this account of the nature of moral properties in the course of defending his theory against traditional 'semantical' objections to ethical naturalism. According to those objections, naturalism is untenable because there is no naturalistic reduction for moral terms.⁸ According to Brink, such objections depend on two assumptions: (a) that synonymy is a test of property identity, and

⁸ The source of the view is thought to be Hume 1739: III, i, 1/469-470. For a representative set of discussions, see Hudson 1969.

(b) that meaning implication is a test of constitution or necessitation. Brink represents these claims as "the semantic test of properties" (Brink 1989: 162). Motivating the semantic test of properties is a particular view of necessity and analyticity: viz., that only analytic claims can be necessary (1989: 162). According to Brink, we should reject the semantic test of properties for two reasons. First, its underlying motive is just untenable: analyticity and necessity are plausibly regarded as distinct properties (1989: 165). Analyticity is semantical; necessity is metaphysical; so H₂O might necessarily realize water even though 'x is H₂O' does not mean 'x is water'. Second, accepting the semantic test would require rejecting many otherwise acceptable scientific and commonsense claims (e.g., that water is composed of H₂O, that World War I was comprised of "events x, y, z," etc.) (Brink 1989: 164). But if we reject the constellation of views represented by the semantic test of properties, then the lack of entailment relations between moral and naturalistic property ascriptions is not itself a reason to reject any particular assignment of moral properties. A given moral property might be constituted or identical to, and necessitated by, a complex of natural properties even if there are no semantic assurances that this is so. For instance, setting a cat on fire for entertainment purposes might be wrong even though 'x is wrong' is not extensionally equivalent to 'x is an instance of setting a cat on fire for entertainment purposes'.⁹

Of course, a realism that postulated a class of inquiry-independent facts but denied inquirers 'epistemic access' to them would be untenable,¹⁰ so AMR would be implausible if it included no account of how inquirers can know, and be justified in believing, that moral facts of the sort described above exist. And Brink's argument against the semantic test of properties, if sound, establishes only that the existence of moral facts cannot be denied on

⁹ See Brink 1989: 156-163.

¹⁰ Devitt (1984: 15) calls this "anti-realism with a fig-leaf."

purely semantic grounds; it does not provide a reason to suppose that there are such facts. This raises the epistemological question: how do cognizers know that any particular natural property complex G produces some moral property F, and not moral property J (or no moral property at all)?

The answer, given by Brink and other AMR theorists, is that cognizers obtain knowledge of moral facts in the same way they discern scientific facts. Moral property ascriptions, like their scientific counterparts, are synthetic; moral knowledge is not gained merely by analyzing the meaning of moral property-ascribing terms; rather, it is gained a posteriori via a broadly empirical investigation into the nature of (moral) reality. Now, at first glance this suggestion seems untenable: how could the fact that a given cluster of natural and social properties G bears some moral property F be discerned empirically? The claim that it could be suggests that cases in which we gain moral knowledge must be like those scientific cases in which we know what the underlying properties of a thing are, but still need to discern its observable macroproperties. Scientists may know, for instance, that something is methane but not know what methane tastes like, or about its behaviour under certain conditions. Such knowledge is gained empirically: we just taste the stuff, and so forth. But knowledge of moral facts is not obviously gained like this. While one can no doubt empirically discern the presence of good things once one is given a certain conception of what good things are, discovering whether anything really is good would not appear to be an empirical enterprise. We may know, for instance, that an action is altruistic, but it is not clear that we could empirically discern the fact that such actions are good. And the challenge facing moral epistemologists is not just to show that we can know of the presence of things that are good, bad, and so forth, given a conception of what these values are; it is to show that we can know that certain things are indeed good, bad, and so forth.

The AMR claim that moral properties, in a perfectly ordinary sense, are natural properties, seems unhelpful in this regard. Natural properties may be epistemically accessible, and moral properties may be identical to or constituted by them, but it is not obvious that we are, in virtue of these facts, in a position to know that the natural properties we apprehend bear any particular moral properties. What is required is not merely a constitution or identification of moral properties in terms of ontologically unremarkable and empirically discernible natural properties; what is required is an account of how cognizers can apprehend the fact that moral properties distribute over natural properties in some particular way, as opposed to distributing differently, or being entirely absent. And because it seeks to assimilate the moral and scientific cases, the AMR theory must employ an epistemological strategy that is plausible for empirical knowledge. Brink's own strategy is inspired largely by the coherence theory of empirical justification recently formulated by BonJour (1986). What follows is a general outline of Brink's theory.

On a coherence theory of epistemic justification, beliefs are justified by their coherence with other beliefs. In this respect coherentism contrasts with foundationalism, which holds that beliefs are justified either because they are (i) self-justifying or (ii) entailed by self-justifying beliefs. According to Brink, coherentism is both (a) a better account of justification than foundationalism¹¹ and (b) compatible with realism.¹²

Foundationalism fails largely because there are no self-justifying beliefs. Rather, the justification of any belief *p* depends on further beliefs about what type of belief *p* is and why such beliefs should be true. Coherentism is advantaged because it does not postulate self-justifying beliefs. And it involves neither vicious regress, nor vicious circular reasoning. It avoids regress because it allows justificatory chains to, as Brink puts it, "loop

¹¹ Brink 1989: 116-122; cf. BonJour 1986: 32.

¹² Brink 1989: 125-130; cf. BonJour 1986: 169-179.

back upon themselves."¹³ It avoids vicious circularity because it is only incomplete, "contextual," justifications that form simple linear loops; complete or "systematic" justification is, according to Brink, nonlinear and is comprised by a belief's coherence with the totality of beliefs one does or might hold.¹⁴

Finally, Brink argues that coherentism and realism are compatible. On his account, truth is correspondence with inquiry-independent facts, and the coherence of a belief with all or most of our other beliefs is evidence of its truth so construed. In order to fully appreciate Brink's reasons for thinking that coherence is evidential, it will help to consider how he resists two standard objections to coherentist realism.¹⁵ The "input objection" relies on the apparent fact that belief sets do not require input from an inquiry-independent reality in order to be coherent. There seem to be many examples of coherent, but nonetheless false, bodies of belief. But if so, the objection runs, there seems to be no reason to suppose that as belief sets approach maximal coherence, they become more probably true. Second, the "alternative coherent systems" objection claims that different cognizers, each employing a coherentist method of justification, could come to hold belief sets that are largely incompatible with each other while enjoying the same degree of coherence; but it seems that, if so, the coherence of a cognizer's beliefs provides no reason to think that any of them are true.

Brink concedes that maximally coherent beliefs might be wildly false, but insists that the possibility of global or massive error need not be ruled out by coherentist realism

¹³ Brink 1986: 123.

¹⁴ Brink 1986: 123. See also BonJour, who insists that beliefs approaching maximal coherence are not related to each other in a circular way; rather, "the justification of a particular belief finally depends, not on other particular beliefs as the linear conception would have it, but instead on the overall system and its coherence" (BonJour 1986: 91-92).

¹⁵ BonJour 1986: 25, 107-108.

(1989: 128, 129). The coherentist realist need claim only that maximal coherence is truth-conducive; not that coherence guarantees truth. Why is coherence truth-conducive? Essentially, because of the special nature and role that certain second-order beliefs play in the justification of empirical beliefs. Recall that the critique of foundationalism depended on the denial that there are any self-justifying beliefs, which in turn depended on the claim that second-order beliefs (beliefs about what kind of belief p is and why such beliefs should be true) are essential to justification. According to Brink, the second-order beliefs we in fact hold are realist beliefs (1989: 127). They include the convictions that are germane to a commonsense account of the world and our place in it, an account that is developed, extended, and often revised, by scientific investigation. This is where explanatory power and testability enter into the account. Many of our beliefs are coherent in the sense that they serve as explanations of the observations we make in confronting the world; they cite facts about the world, our perceptual capabilities, the biological and other causal mechanisms that underlie our ability to discern features of reality, the conditions under which those capabilities are reliable, and so on (1989: 127-128). These, together with further beliefs about the role and importance of such explanations in justification, combine to produce a belief set that is essentially a comprehensive explanatory account of our experiences, and since it turns out that this account includes realist assumptions, we may regard the overall coherence of our beliefs as evidence of their truth (1989: 128). Coherence provides no guarantees (even maximally coherent beliefs might be "wildly false"); but it provides no reason to accept skepticism, and many reasons to reject it (1989: 128-129).

The theoretical beliefs that comprise this coherent explanatory account of our observations are also testable against those observations - testable in a way that is amenable to coherentism. It is widely recognized that our scientific theories have observational

consequences, and are testable, only when conjoined with "auxiliary hypotheses."

Predictive failure is reason to reject either the theory or one or more of the auxiliary claims.

Blame for predictive failure is cast on whichever conviction is less well supported, i.e., is less coherent with the rest of our beliefs and observational judgments. The coherence of the conjoined claims is assessed on the basis of further tests involving still other

(theoretical and observational) claims. As Brink puts it:

To determine at whose door we should lay the blame for predictive failure, we must ascertain the independent plausibility of each of the contributing hypotheses, and because hypotheses can be tested only in bundles, we must judge the independent support of contributing hypotheses by testing them and varying the bundles in which we test them. (1989: 137)

Making claims testable, and justifiable in virtue of their explanatory power, allows the account to respond to the aforementioned objections. We have reason to suppose that independent features of the world have input into our belief sets because the beliefs that affirm them are testable (albeit in bundles), because they figure in the best explanations of our observational judgments, and because both testability and the capacity to explain can be assessed in terms of coherence. Moreover, solving the concern about input makes it possible to respond to the worry about many incompatible, but equally coherent, belief sets: maximally coherent belief sets are unlikely to radically disagree because they each receive input from the same source.

Now, the same objections can be raised against coherentist moral realism. It may be argued that a cognizer's moral beliefs do not require any input from an inquiry-independent moral reality in order to be coherent. A person's moral beliefs may therefore be internally coherent, but false. Indeed, it may be argued that they are probably false because different cognizers can and do hold moral beliefs that are equally coherent, but incompatible, and this is best explained without mentioning moral facts. The AMR response is this: the objection is supposed to show that moral beliefs are only internally

coherent (i.e., that moral beliefs are coherent with each other but not with nonmoral beliefs); but this is just false.¹⁶ Like ordinary perceptual beliefs, moral beliefs are often formed under conditions that are thought to be generally conducive to the formation of true beliefs, so they can be coherent with certain second-order epistemic convictions about the reliability of beliefs so formed, and can have the same sort of initial credibility that ordinary perceptual judgments possess (Brink 1989: 132, 136). Moreover, moral judgments cohere with many of our other beliefs by contributing to a comprehensive explanatory account of our observations. Particular moral judgments can count as evidence for our moral theories because they are best explained on the assumption that the facts they affirm are present and cause us to judge as we do (1989: 182-197).

The claim that some nonmoral facts can and must be explained on moral assumptions is at least plausible because we commonly do offer moral explanations of much of what we encounter in the world. AMR theorists point to persuasive evidence of this. We may attempt to explain the widespread unrest in South Africa during the apartheid period by citing the injustice of the apartheid system (Brink 1989: 187-188); to explain Hitler's behaviour (and consequently many important historical events) on the assumption that he was morally depraved (Sturgeon 1984: 68-69); or to explain a person's sudden change in moral opinion on the assumption that he has seen a film that has caused him to perceive a hitherto unnoticed moral fact that has disconfirmed his original opinion (Werner 1983).

Brink makes two arguments in defence of moral explanations: the Counterfactual Argument and the Generality Argument.¹⁷ The first is an argument for the explanatory

¹⁶ However, Brink thinks that the internal coherence (of moral beliefs with other moral beliefs) can be evidential if moral disagreement can be explained without assuming nihilism (1989: 182-183).

¹⁷ These labels are mine and not Brink's.

relevance of moral facts to explanations of certain nonmoral facts (e.g., the fact that a person has made a certain moral judgment). Drawing largely on work by Nicholas Sturgeon (1984), he argues that given the principle of moral supervenience, the explanatory relevance of moral facts must be assessed via a counterfactual test that moral facts can often pass (1989: 190-191). The second is an argument for the indispensability or necessity of making moral assumptions in explaining certain nonmoral facts. He argues that sometimes only moral assumptions are sufficiently "general" to explain nonmoral facts (1989: 194-197).

On Brink's account, our moral theoretical convictions are, moreover, testable in the same indirect Duhemian way as are our scientific convictions: that is, in bundles against our observational beliefs.¹⁸ So, the BonJour-Brink account runs, moral beliefs are or can be coherent with the rest of our beliefs, and if coherence is conceded to be evidence of the truth of our nonmoral beliefs, it must be evidence of the truth of our moral beliefs too.

Finally, Brink makes an important dialectical point: i.e., that we can and should assume that moral realism is true unless confronted with "powerful" reasons to think otherwise (1989: 24). His rationale for this dialectical position is unclear. His considered view seems to be that the presumption is warranted by certain second-order beliefs that we hold about the nature of moral inquiry; in particular, by beliefs about those features that are commonly thought to support ethical cognitivism, such as beliefs about the fact-affirming and property-ascribing form of moral judgments (1989: 23-29). But he also seems inclined to justify the position on phenomenological grounds (i.e., on the grounds that the world just seems to us to contain real moral values).

We are now in a position to give a concise summary of AMR's core claims, as they are presented by Brink:

¹⁸ Brink 1989: 136-138; see also Sturgeon 1984, 1986.

1. There exist inquiry-independent moral facts; these facts are ontologically and epistemologically unremarkable when compared to facts affirmed by recognizably realist disciplines such as science; in particular:

(a) like scientific facts, moral facts are facts about properties that supervene upon, and are constituted by (but often not identical to), certain complexes of natural and social properties; and

(b) moral beliefs, such as the belief that (natural and social property complex) G is (moral property) F, can be and often are warranted by their coherence with certain other beliefs (e.g., with other reliably-formed and thus similarly coherent moral beliefs, with second-order beliefs about the reliability of beliefs formed in the way moral beliefs are formed; and with observational beliefs [by explaining certain of our observations, such as what we observe people doing, etc., and by passing Duhemian tests]).

2. We should presume that inquiry-independent moral facts obtain unless we are given "powerful" reasons to think otherwise. We should make this presumption either (a) in light of certain facts about moral inquiry, or (b) because the world seems to contain such facts.

1.4: Problems with the Account

In this thesis I argue that AMR as it is presented and defended by Brink and other theorists (e.g., Sturgeon, Werner) is untenable. My case depends largely on the claim that moral facts are inefficacious; they cannot figure as the causes of anything. This claim is plausible because, as Harman (1977: 7-8; 1986: 62) has pointed out, we have no idea how moral facts could figure as causes. Now, I can think of two reasons to reject Harman's claim, both of which I find untenable. First, it might be supposed that moral facts can figure as causes because the natural property complexes upon which they are supposed to supervene can be causes. For instance, it might be contended that Hitler's moral depravity can be regarded as a cause of certain historical events just because the complex of natural properties that comprises that moral property can be regarded as the cause of those historical events. But the efficacy of moral facts cannot be established just by pointing to the efficacy of the bases upon which they putatively supervene. We need, in addition, some account of how the fact that certain complexes of natural and social properties

realizes a moral property could be the cause of something. We have such an understanding in the case of many nonmoral facts involving supervening properties. For instance, there is a perfectly acceptable scientific explanation of how the fact that certain things are red can cause people observing those things to have a certain visual experience. That is, all red things have a particular arrangement of microphysical particles that, under certain (standard) conditions, cause suitably-constructed creatures to have a certain visual experience. But we have no comparable understanding of how the fact that certain things are, say, good, could affect cognizers so as to cause them to detect the goodness of those things. Indeed, we have no comparable explanation of how such a fact could serve as the cause of anything. Now, it might be contended that moral facts are in this regard unremarkable when compared to many similarly construed nonmoral facts that we have reason to affirm. Consider, for instance, a parallel claim about tables, viz., that items possessing certain sets of underlying properties, such as being of sundry size, having a flat, hard surface elevated to a certain height from the ground, "really are" tables. Of course, we do have reason to suppose that such items are tables. But our reasons are conceptual, not abductive: such items are tables just because part of our concept of a table holds that to be a table is to possess the property of having a hard, flat surface, and so forth, and not because the assumption that they are tables services a causal explanation. Indeed, just like in the moral case, we have no understanding of how the fact that items possessing hard, flat surfaces are tables could figure as the cause of anything. So, the AMR theorist cannot explain the efficacy of moral facts by pointing to table facts. And since AMR theorists insist that our knowledge of moral facts is not gained via conceptual analysis but via empirical investigation, they would not, I think, be satisfied to know that our knowledge of moral facts is gained in the way our knowledge of table facts is gained.

The other possible reason to reject Harman's claim cites the fact that we commonly ascribe causal status to facts even though we do not understand the precise causal mechanisms by which two items are connected. For instance, if there were a statistical correlation between the presence of chinook winds and migraine headaches, we might (plausibly) regard the presence of chinook winds as a cause of migraine headaches, even if we have no understanding of the exact mechanisms involved. And if the lack of such an understanding does not undermine our ascriptions of causal status there, then it might be thought that it should not do so in the moral case. But there is, I think, a significant difference between the two cases: even though we may not fully understand how chinooks cause migraines, we have no difficulty in understanding how chinooks can be causes of some things (e.g., certain meteorological phenomena, perhaps); by contrast, we have no understanding of how the fact that something is good could be the cause of anything.

In the course of my discussion it will become clear that this mystery about the efficacy of moral facts undermines AMR in several ways. First, it is a reason to reject Brink's dialectical stance (2.3). One of the possible rationales offered by Brink for that stance cites facts about moral phenomenology, i.e., the fact that the world seems to contain real moral values. As such, it depends also on the principle that we are generally warranted in presuming that things are as they appear. But of course that principle could not reasonably be applied to the moral case if there were reason to believe that the appearance of real moral value in the world could not have been caused by real moral facts (2.3). Moreover, there is another epistemic principle that can and should be employed in the case of moral appearances, viz., when the experience in question is an experience of something whose existence is itself a matter of substantial controversy among reasonable persons, then one ought generally to suspend judgment about whether the item experienced really was present and caused one's experience. Because the truth of AMR is in dispute, a more

reasonable dialectical approach is thus to simply suspend judgment until one has examined the issue in greater detail (e.g., by assessing the arguments for and against the theory in the literature) (2.3; 2.4).

Second, the fact that we do not understand how moral facts could be causes also undermines Brink's account of moral belief justification. One way it does this is by undermining his case for the initial credibility of moral judgments. On Brink's view, some moral beliefs (considered moral judgments) are initially credible in virtue of their coherence with certain second-order beliefs about the type of beliefs they are and about why beliefs of that sort are credible. But presumably such a second-order account should include assumptions about the efficacy of the class of facts being affirmed by the beliefs in question. Lacking any understanding of how moral facts could cause us to make moral judgments, it appears we lack also any second-order rationale for regarding moral judgments as having initial credibility. But then the coherence of our other beliefs with considered moral judgments confers on them no warrant: we have no reason to regard the clusters of beliefs that are coherent with our considered moral judgments as anything more than internally coherent belief sets lacking any input from inquiry-independent moral facts.¹⁹ Now, it might be supposed that the worry about input should be ameliorated by certain of the arguments presented by AMR theorists. Sturgeon and Brink present arguments for the explanatory relevance and indispensability of moral fact-affirming assumptions. And Richard Boyd has defended an account of the meaning of moral property-ascribing terms that, if sound, makes plausible the idea that their use by inquirers is somehow "regulated" or "guided" by moral facts. But, as we shall see, those arguments are themselves threatened by the mystery to which Harman has pointed, and do not succeed in ameliorating it. Since we do not understand moral efficacy, statements

¹⁹ See Brandt 1979: 20; Little 1984: 376.

affirming moral facts cannot pass any (plausible) counterfactual tests of explanatory relevance (4.3); nor can their apparent "generality" make them indispensable to certain causal explanations (5.2). And Boyd's argument fails for similar reasons (5.5). All of this suggests that moral (fact-affirming) beliefs cannot possess the sort of coherence that on the AMR account is required for justification; i.e., that they cannot be coherent with plausible second-order beliefs about their etiology. And the failure of the AMR account of the epistemic justification of moral property ascriptions is arguably a sufficient reason to reject AMR itself.

Finally, Harman's observation undermines Brink's response to objections to ethical naturalism that are motivated by the semantic test of properties. On Brink's account, moral facts are facts about supervening, constituted (hence multiply realizable) upper level properties whose existence is discerned via the ordinary cognitive and perceptual processes employed in empirical inquiry. This construal of moral facts assumes a crucial role in the defence of naturalism against the semantical objection mentioned above. For it is plausible to suppose that moral facts, so construed, might very well obtain even if moral property ascriptions cannot be justified on the basis of an extension-fixing analysis of their constituent terms. However, appealing to these modal concepts would arguably not assist the case for AMR if it turned out that the modal features of moral property ascriptions could be explained without making any realist assumptions. In fact, I defend a modal analysis of moral property ascriptions that incorporates the concepts Brink mentions but that is not essentially realistic (Chapter Six). On my account, moral property ascriptions are a special sort of "conceptual" claim; i.e., they are determined by facts about what ascribers are and are not capable of conceiving (6.6). For instance, whether a person is disposed to say that Hitler is a bad person in light of his character and behaviour will depend largely on whether she is able to conceive of a world in which Hitler had that

character and behaved in that way without being bad. I contend that, in order for the AMR theorist to accommodate this account of moral ascriptions, it would have to be possible to explain certain nonmoral facts on moral assumptions (6.8). In particular, it would have to be possible to explain the particular cluster of conceptual abilities that motivate a person's particular moral judgment on the assumption that a moral fact, viz., the fact affirmed by the moral judgment, has somehow caused the person to have those particular abilities. Now, it is easy to see how Brink's account of moral belief justification might be adjusted so as to meet this requirement. On Brink's account, moral beliefs are justified if they can reasonably be regarded as having the appropriate sort of causal etiology; and it could plausibly be contended that they could and would have this etiology (even qua conceptual claims) if it were possible to explain a person's conceptual abilities on moral assumptions. For instance, one might explain a person's judgment about Hitler on the assumption she has a particular set of conceptual capacities (she finds it inconceivable that Hitler could have done the things that he did, and had the personality he had, but not have been a bad person), and explain the conceptual capacities in turn on the assumption that the moral fact affirmed by the judgment has somehow affected her so as to cause her to have those particular abilities or inabilities. Indeed, expressed in the idiom of Brink's account, one could say that moral beliefs are made coherent (with important second-order beliefs about reliability and so forth) precisely by the making of explanations of this sort. But this sort of response is untenable if moral facts are inefficacious and cannot figure in causal explanations. The causal impotence of moral facts is a reason to accept a very different construal of moral supervenience than that proposed by AMR theorists. It is a reason to construe moral supervenience not as an account of the nature of moral properties and how they relate to natural properties, but as an account of moral judgments and of the modal constraints to which they are subject. But of course that construal of moral supervenience does not serve

the realist's case against the semantic objection, for it does not explain how moral properties can "be" natural properties notwithstanding failures of semantic reduction.

CHAPTER TWO

THE BURDEN OF ARGUMENT

2.1: Introduction

Brink begins his inquiry by placing the burden of proof squarely on the moral antirealist's shoulders. He writes:

Moral realism should be our metaethical starting point, and we should give it up only if it does involve unacceptable metaphysical and epistemological commitments (W)e have reason to accept moral realism that can be overturned only if there are powerful objections to moral realism. (1989: 24)

Brink is not alone in making this dialectical move; it is made by advocates of diverse forms of moral realism.²⁰ In this Chapter, I examine whether the position has any merit. Since, as we shall see, Brink does not himself make his case for it entirely clear, I begin by giving a partial reconstruction of that argument, partly by noting Brink's own remarks, and partly by citing what seem to me to be some of the assumptions upon which any "burden of proof" argument depends. I argue that Brink is at least somewhat sympathetic to two rather different arguments for the onus shift. The first cites facts about our common moral phenomenology; the second, and evidently his more carefully considered view, cites those facts about moral inquiry that are commonly cited by ethical cognitivists in support of their theory. I argue that neither of these arguments warrants the stance in question.

2.2: What is Brink's Argument for the Onus Shift?

This section contains a brief account of how I think the argument can and should be interpreted, one that I hope distinguishes it from a number of claims with which it might be confused.

²⁰ See, e.g., McNaughton 1988: 40; Nagel 1986: 143; 1980: 100.

First, I take it that the dialectical stance in question must, to be epistemologically interesting and worth analyzing, affirm more than just a fact about our moral psychology or initially-held metaethical convictions. It must be a claim about what cognizers are entitled to presume about moral realism. So I take it that Brink is affirming not just that there is a pervasive presumption that moral realism is true, but that inquirers are warranted in making that presumption. This may seem obvious, but it is not always made clear in the literature. For instance, McNaughton, another realist, seems to confound the two claims when he writes:

The realist insists on an obvious, but crucial, methodological point: there is a presumption that things are the way we experience them as being - a presumption that can only be overthrown if weighty reasons can be brought to show that our experience is untrustworthy or misleading. Moral value is presented to us as something independent of our beliefs or feelings about it; something which may require careful thought or attention to be discovered. There is a presumption, therefore, that there is a moral reality to which we can be genuinely sensitive. (1988: 40)

Of course, it may turn out that the epistemic thesis follows from certain assumptions about our moral psychology, moral inquiry, and so forth, together with certain epistemic principles. Indeed, as we shall see, Brink (sometimes) seems to think that the pervasive acceptance of the presumption in question, together with a certain philosophical defence, adequately justifies the presumption. But it remains to be seen whether this is so and what sort of epistemic principle(s) could warrant such a position.

Second, Brink is presumably not just claiming that certain facts constitute prima facie or tentative evidence for moral realism, but that they constitute enough evidence to warrant the presumption in question. Indeed, his case must, I think, assume this form in order to do the work he wants, for establishing merely that certain facts are prima facie evidence for moral realism is not in itself sufficient to shift the burden of argument to the moral antirealist. Since a person can have evidence for p but also have equally weighty or stronger evidence against p, it is possible for a person to have evidence for p without it

being reasonable for that person to accept p until confronted with powerful reasons for thinking that p is false. What would be enough evidence to warrant a presumption of this sort? Presumably, evidence of p 's truth that is weighty enough to make it probable (indeed, even very probable) that p is true (notwithstanding other evidence a person may have to the contrary). This should, I think, be true of the presumption that moral realism is true. That is, one should be entitled to make this presumption and to hold it until one is given strong reasons for thinking that it is false, only if one has evidence that makes it very probable that the theory is true. But notice that accepting the distinction between just having evidence and having enough evidence makes it possible to agree with Brink and other moral realists that certain facts are evidence for moral realism, while also denying that the moral antirealist must shoulder the burden of argument. If, for instance, we had other evidence against moral realism, this might very well shift the burden in the other direction. I take this to provide the antirealist who rejects the realist's dialectical position much more room to manoeuvre than one might otherwise expect, for now he need not be committed to the (perhaps untenable) view that the phenomenological and other facts realists mention have no evidential standing whatsoever. It also makes the realist's task significantly more difficult than it otherwise might be: the realist must show that the facts he cites, whatever they are, constitute enough evidence to make his favoured theory more probable than any competing theory, but assessing the respective probabilities of AMR and its competitors would not seem to be an easy task.

Third, since to have evidence for a theory is to have evidence for that theory at a particular time, the thesis should be understood as a "temporally relative" one, i.e., as the claim that at a certain time t the overall evidence makes it probable that p is true. Put another way, since Brink's thesis is explicitly "dialectical," it should be construed as the claim that at a particular point in the ongoing process of moral inquiry (which would

presumably include the ongoing dialogue about moral realism), inquirers have enough evidence overall to make it probable at that point that moral realism is true. The "dialectical" nature of the thesis complicates matters still further, for it raises the question of precisely when, in their ongoing inquiry into the nature of ethics, are inquirers entitled to presume the truth of moral realism? There are presumably a plurality of possible positions on this. For instance, one might hold that a person is entitled to the presumption if it seems to that person that the world contains real moral values. On this view, inquirers seem entitled to presume the truth of moral realism at an early stage in their inquiry, e.g., before having taken stock of any of the arguments of metaethicists. On the other hand, one might claim that there is, at the present point in the ongoing dialogue about moral realism, enough evidence favouring that theory to warrant inquirers in presuming its truth as they pursue their inquiry further.

It is noteworthy that Brink is not entirely clear about exactly where along the dialectic he thinks the presumption becomes warranted. His remarks about the importance of the commonsense presuppositions of moral inquiry and moral phenomenology suggest that he is sympathetic to the first view. But he also appears to think that his rather sophisticated account of the action-guiding character of morality and of moral realism's practical implications lend support to the presumption. For instance, in a passage that occurs late in his book, he writes :

My defence of moral realism is now near completion. We have seen reasons to want to defend moral realism. Various features of moral inquiry presuppose or are naturally explained by a realist form of cognitivism (2.5-2.8); the practical or action-guiding character of morality supports moral realism (Chapter 3); and moral realism's practical implications seem superior to those of many antirealist theories (Chapter 4). These considerations create a presumption in favour of moral realism and place the burden of proof on the antirealist. (1989: 211)²¹

²¹ See also 1989: 132.

Of course, these dialectical positions may not be equally viable. In the following sections I present separate discussions of each of what I take to be the promising and exegetically sound versions of these positions.

Fourth, I take it that the presumption in favour of "moral realism" that Brink wishes to defend is a presumption in favour of Brink's version of moral realism, viz., AMR. Notice, in this connection, that although Brink initially construes his theory in a fairly broad way, i.e., as the thesis that

MR: (1) There are moral facts or truths, and (2) these facts or truths are independent of the evidence for them,

he makes it clear that he regards this as representing just a "core element" of his conception of moral realism, and that he ultimately intends to defend a very special variety of moral realism, one whose various features should become evident in the course of his defence (1989: 17-19). Indeed, anyone who accepts his ensuing arguments for parity between science and morality would, in virtue of accepting those arguments, be committed to a particular conception of what moral facts are like. For instance, one would accept that they are facts assumptions about which can service explanations of moral judgments, that moral facts are typically constituted out of natural property complexes, and so forth. All of this suggests to me that Brink's argument for the presumption favouring moral realism should be taken as an argument for the presumption that there exist certain moral facts that are like scientific facts in the respects he mentions. This further complicates the burden of proof issue. For it raises the question of whether the onus-shifting factors (whatever they are) warrant a presumption in favour of Brink's theory, or just, say, some form of metaethical objectivism.

Finally, while the structure of Brink's case for the onus shift is not readily apparent, I take it that any such dialectical case will depend (at least implicitly) on (a) certain factual assumptions, and (b) an epistemic principle in virtue of which those factual assumptions, if

true, warrant shifting the burden of argument in favour of a particular theory. McNaughton (1988), again, has given us an illustrative example of such a theory; one where the aforementioned structure is more apparent. McNaughton evidently thinks that the burden is shifted in light of certain facts about moral phenomenology (such as the fact that the world seems to a person to contain real moral values), and in virtue of Richard Swinburne's Principle of Credulity (PC).²² The PC affirms that:

It is a principle of rationality that (in the absence of special considerations) if it seems (epistemically) to a subject that x is present, then probably x is present; what one seems to perceive is probably so. How things seem to be is good grounds for a belief about how things are. (1979: 254)

According to Swinburne, if S has an experience that seems to him to be an experience of God's presence, then he is warranted in presuming that it is an experience of God's presence - unless the experience occurs under conditions of general epistemic unreliability, or there is some other positive reason to suppose that it was not an experience of God's presence. McNaughton's argument, then, is simply this:

- (1) We are generally entitled, by the PC, to presume that things are the way they seem (unless considerations such as those just described obtain);
- (2) The world seems to us to contain real moral value(s);
- ∴ (3) We are entitled to presume that the world contains real moral value(s) (unless special considerations obtain).²³

Now, I take it that Brink's argument is similar to McNaughton's in its general structure: Brink's argument, like McNaughton's, implicitly depends on certain factual assumptions and on a certain epistemic principle in virtue of which the facts affirmed by the assumptions warrant a shift in onus. But it is not clear whether Brink would want to accept Swinburne's principle. For he does not want to make his case for the onus shift

²² McNaughton 1988: 40-41, 65; Swinburne 1979: 254-260.

²³ McNaughton 1988: 40-41.

depend on how things seem to any particular cognizer (nor indeed on how things seem to cognizers in general), but on certain facts about moral inquiry and language.²⁴ He writes:

My appeal to commonsense moral thinking is not a prediction about the likely results of a Gallup poll on the issue of moral realism. Rather, my concern is with the philosophical implications or presuppositions of moral thought and practice. (Compare the way in which philosophers of science take the practices of working scientists as an important methodological constraint but largely discount scientists' philosophical views about the status of their research.) I claim that cognitivism seems to be presupposed by common normative practices of moral judgment, argument, and deliberation and that reflection on the nature of moral theorizing seems to support a realist view about these moral facts and truths. This claim may be false, but this is not shown by an appeal to common metaethical beliefs (or the lack thereof). (1989: 25)

We need not catalogue all of the features of moral inquiry Brink mentions in his book; they are just the familiar features cited by ethical cognitivists in support of their theory: e.g., the declarative form of moral sentences, their apparent ascriptive function, the presupposition germane to our moral thinking that moral mistakes are possible, and so forth.²⁵ Brink evidently thinks the shift in onus is warranted because he thinks that these features are best explained on the assumption that moral inquiry is conducted in order to discover inquiry-independent moral truths or facts. He writes:

If my argument has been correct, general considerations about the nature of inquiry and considerations of moral inquiry in particular are most easily explained on the assumption that moral inquiry is directed at discovering moral facts that obtain independently of our moral beliefs and at arriving at evidence-independent true moral beliefs. I take this conclusion to establish a presumptive case in favour of moral realism and to shift the burden of proof to the moral antirealist. (1989: 36)

All of this suggests that Brink's argument depends not on the PC, but on something approximating:

PC*: If certain general considerations about the nature of an area of inquiry are most easily explained on the assumption that the inquiry is undertaken in order to discover

²⁴ In fact, since Brink is "willing to admit that, about moral realism, common belief is silent, divided, or even antagonistic" (1989: 25), it appears that the actual prereflective views and experiences of cognizers are supposed to be irrelevant to the thesis.

²⁵ See Brink 1989: 25-35.

evidence-independent truths, then realism about that area is probably true and antirealists about that area of inquiry must shoulder the burden of argument.

So, Brink's argument for the onus shift in favour of AMR is this:

(1) we are generally entitled, by the PC*, to presume that realism about an area of inquiry is true if features germane to the area of inquiry are most easily explained on the assumption that the inquiry is undertaken in order to discover evidence-independent truths;

(2) features germane to moral inquiry are most easily explained on the assumption that the inquiry is undertaken in order to discover evidence-independent moral truths;

∴ (3) we are entitled to presume that moral realism (i.e., AMR) is true unless confronted with positive reasons for thinking otherwise.

This is, I think, a fair interpretation of Brink: one that at once elucidates the assumptions that I think anyone who makes this kind of argument would have to make (e.g., Brink himself does not explicitly say that the onus shift is warranted by considerations of probability, but it is difficult to see how else to understand the claim), and that is exegetically sound. Hence, a reasonable analysis of his position should, I think, proceed on the assumption that the epistemic principle on which he relies is the PC* and not the PC.

Nonetheless, there is some reason to examine a PC-based version of the argument. First, the question of whether the dialectical stance in favour of AMR could be warranted on phenomenological grounds is intrinsically interesting: since phenomenological considerations have seemed so important to other moral realists, one naturally wonders how much justificatory weight phenomenological considerations can carry in warranting the presumption in favour of Brink's theory. Second, Brink himself does not entirely dismiss "phenomenological" claims - he thinks they are "correct and important."²⁶ Indeed,

²⁶ For instance, he writes: "Moral judgments are typically expressed in language employing the declarative mood; we engage in moral argument and deliberation; we regard people as capable both of making moral mistakes and of correcting their moral views; we often feel constrained by what we take to be moral requirements that are in some sense imposed from without and independent of us. These phenomena are held to demonstrate

he seems to argue for a Swinburnian view of what we are entitled to presume in defending his view of the "initial credibility" of considered moral judgments (CMJs). For Brink, among the factors that accord initial credibility to CMJs are second-order beliefs about the reliability of beliefs formed in the way CMJs are formed; for Swinburne, the PC itself depends on the idea that we generally get things right, and is suspended or limited precisely when there is reason to believe that our cognitive or perceptual abilities have been compromised. Thus, judgments on both accounts seem to gain and lose initial credibility in the same way: viz., in virtue of the presence or absence of reasons to believe that they have been formed via generally reliable belief-forming processes.²⁷

For these reasons I have included a section that examines whether a presumption in favour of AMR might be warranted by the PC and by the phenomenological facts it recognizes, even though I think Brink's view relies on the PC*. I will argue that neither principle succeeds in justifying the shift in onus Brink defends.

2.3: The PC Argument

The PC, according to Swinburne, affirms a presumption that things are the way they appear unless we can cite some "special consideration" that shows otherwise (1979: 254). Expressed in terms of probability: if *e* means 'it seems to *S* that *x* is *w*', *h* means '*x*

the realist or cognitivist character of commonsense morality; morality seems to concern matters of fact that people can and do sometimes recognize and debate about. I think these phenomenological claims are correct and important, in part on their face and in larger part because they reflect, and are confirmed by, various philosophical presuppositions of inquiry in general, and moral inquiry in particular." (1989: 24)

²⁷ I have some reservations about the plausibility of the parallel proposed here. For instance, Brink never claims that CMJs can be presumed to be true; rather, he just claims that coherence with such beliefs is "evidential" (1989: 139). He also concedes that CMJs may not have the same degree of credibility as nonmoral observational judgments (1989: 138).

is w', k denotes irrelevant background knowledge (i.e., knowledge that has no bearing on the truth of h), and c denotes special considerations, then

$$PC: P(h/e \ \& \ k) > .5 \ \& \ P(h/e \ \& \ k \ \& \ c) \leq .5.^{28}$$

According to Swinburne, "in effect the PC is claiming . . . (that) men have a basic ability to detect how things are" (1979: 275). Accordingly, he justifies the PC by a reductio against general epistemological skepticism. In acquiring knowledge, we need to start somewhere; our initial convictions or perceptions will need to enjoy a presumption of truth. Rejecting the PC, he says, lands us in a "skeptical bog" (1979: 254, footnote). Apart from challenging the rationality of this principle, we can ask what special considerations might limit the PC. What are the considerations c that permit S to (at least) rationally suspend judgment about h, notwithstanding e?

Let us begin with Swinburne's own list:

1. One may show that the apparent perception was made under conditions or by a subject found in the past to be unreliable. Thus one may show that S's perceptual claims are generally false, or that perceptual claims are generally false when made under the influence of LSD, which is good inductive evidence for believing that a particular new perceptual claim made by S or made under the influence of LSD is false.
2. One may show that the perceptual claim was to have perceived an object of a certain kind in circumstances which similar perceptual claims have proved false. Thus if it seems to S that he has read ordinary-size print at a distance of a hundred yards, we can test him on a number of other occasions and see if he is able to read what is written at that distance; and if he is not we have good inductive evidence that the original claim was false.
3. Since to perceive x is to have one's experience of its seeming that x is present caused by x's being present, one can challenge a perceptual claim to have perceived x either by showing that probably x was not present or
4. . . . by showing that even if x was present, it probably did not cause the experience of it seeming that x was present . (1979: 260-261)

²⁸ Swinburne 1979: 260.

These are "limits" to the principle of credulity. We need not award any presumption of truth to perceptual claims when at least one of these conditions obtain. Generally, they concern the availability of inductive evidence for the falsity of the perceptual claim (e.g., we need not award that presumption to claims made by persons who are compulsive liars since there is good inductive evidence for the claim that any testimony they give is false). Swinburne claims that none of these considerations typically obtain in the case of religious experiences.

Might any obtain, however, in the case of moral experiences? Let us assume that the experience at issue is the experience of 'real moral properties', and assume that Brink effectively describes an important element of this experience when he says that moral value presents itself as something "imposed from without and independent of us" (1989: 24). Accordingly, when Spike (who is sober, sane, and not a pathological liar) rounds a corner and encounters hoodlums setting fire to a cat, he perceives the badness in their action as if it were a real property; further, he feels compelled, if not to intervene directly, to summon the police (or the SPCA), and he perceives his obligation as a moral constraint imposed on him from without by real moral features of the situation. It seems to Spike that the hoodlums are really, in Brink's sense of that term, bad.

Does Spike's judgment pass Swinburne's test? His judgment passes, I believe, 1, and 2. Spike is not a pathological liar or a drunk or mentally ill, and he possesses ordinary perceptual and cognitive abilities. Moral judgment obviously does not require, or typically involve, being intoxicated or mentally ill or otherwise perceptually or cognitively impaired. However, it might be argued that 3 and 4 limit the PC in the case of moral judgments because it is probable that there are no such properties, and moreover that even if there are such properties, it is improbable that they are responsible for moral perceptions. In fact, antirealists such as Harman have argued for precisely the second of these claims. This

suggests that Harman's argument and recent defences of it can be brought to bear on the dialectical question in a way that defeats the realist's dialectical stance (at least to the extent that that stance is based on Swinburne's principle). Assuming those antirealist arguments are sound, and that we should accept the limits that 3 and 4 impose on the PC, it appears we should not apply the PC to moral judgments.

Moreover, there is another plausible limit to the PC that Swinburne does not mention, and whose acceptance appears to require limiting the PC in the case of moral judgments. That is, a cognizer could be not be warranted in presuming that things are the way they seem if the appearances concern a matter of general controversy. On the contrary, it would seem that in such cases the most reasonable response is to simply suspend judgment about whether things really are the way they initially appear. This point appears in the literature on Swinburne. For instance, R. W. Clarke (1984) argues that the PC should not be applied if there is "significant controversy among rational human beings" on an issue.²⁹ So, if I see what seems to be a flying saucer, or an apparent exhibition of ESP, or anything else whose existence is a matter of substantial controversy among rational persons, I could not be warranted in presuming that the object of the experience is as it appears. It is arguably likewise in the moral case: since whether moral realism is true is a matter of substantial controversy, the person to whom the world seems to contain real moral values could not reasonably presume on this basis alone that this is so, but should

²⁹ He writes: "If, as seems clearly to be the case, a person can in all ways be a normal observer (relative to our culture, of course) and still not have a glimpse of God, or other religious experience under any known objectively specifiable set of circumstances, then there is no reason to suppose that a general defence of the empirical foundations of knowledge requires that we assign any evidential value to religious experiences. In place of saying that a person's alleged glimpse of heaven provides some reason to believe that heaven exists, we can say that the person's experience is something to be interpreted and explained in light of all relevant knowledge about the world (based upon experiences which people have under specifiable normal conditions of observation), and that until a particular interpretation is settled upon the person's experience remains evidentially neutral." (Clarke 1984: 192)

just suspend judgment about the matter until having investigated the issue further (at least until they are in a position to form a somewhat educated opinion on the topic in question).

The claim that inquirers should suspend judgment about controversial experiences rather than applying the PC is, moreover, not just plausible, but defensible. Presumably, a person's claim about how certain things seem will be controversial just because that person's phenomenological response is not shared by all rational inquirers, i.e., because even though it seems to a rational inquirer (or sub-group of such inquirers) that x is present, there are other rational inquirers to whom it seems (under relevantly similar circumstances) that x is not present (or who can marshal other reasons to deny the presence of x). But the PC implies only that, in such situations, each inquirer is warranted in presuming that things are how they seem to them; hence where there are phenomenological differences across individuals, it would simply warrant multifarious incompatible presumptions. As William Forgie has argued:

(The PC is impotent) as a polemical weapon for settling any controversial question of existence. Suppose x is something whose existence is a matter of controversy. And now suppose N claims to have perceived x. Those who are skeptical of the existence of x will announce their belief that although N's claim is no doubt sincere and may well be prompted by an accurate perception of something or other, his identification of that something as x is unjustified. As neutral observers we will find ourselves in a position of stalemate. Accepting the PC, we will have prima facie evidence (from N's perceptual belief) that any auxiliary beliefs N used in identifying the object of his experience justify that identification. But we will also have prima facie evidence (from the skeptics' belief) that they do not. Again the stalemate can be overcome only by a careful look at the quality of N's grounds for his identification. The PC, having helped to produce this stalemate, cannot resolve it. (1986: 156)

Now, this difficulty would never arise in the moral case if Brink's assumptions about our moral phenomenologies are correct. For Brink evidently thinks that moral realism is the natural and universal stance. Says Brink (apparently speaking for us all): "we are led to moral antirealism," and away from that natural stance, because we "come to regard the moral realist's commitments as untenable, say, because of the apparently occult nature of

moral facts or because of the apparent lack of a well developed and respectable methodology in ethics."³⁰ Accounts of the prereflective stance that are at odds with this "dialectical picture" are "revisionist."³¹ But these phenomenological assumptions seem plainly false: i.e., the world does not appear to everyone the way it may appear to Brink. There is, we might say, considerable phenomenological pluralism or heterogeneity among ordinary moral inquirers.³² I believe this in part because the world does not appear to me the way it apparently does to Brink; but there is evidence for pluralism in the literature as well. Consider Harman's confession:

I have always been a moral relativist. As far back as I remember thinking about it, it has seemed to me obvious that the dictates of morality arise from some sort of convention or understanding among people, that different people arrive at different understandings, and that there are no basic moral demands that apply to everyone. For many years, this seemed so obvious to me that I assumed it was everyone's instinctive view, at least everyone who gave the matter any thought 'in this day and age'. (1984: 27)

Harman, in fact, goes on to suggest a dialectical picture that seems precisely the reverse of Brink's:

In the 1960s I was distressed to hear from various people teaching ethics that students in their classes tended to proclaim themselves moral relativists until they had been shown how confused they were about ethics. (1984: 28)

He concludes that some people just seem pre-reflectively inclined toward moral relativism, others toward moral "absolutism." Similarly, Panayot Butchvarov has claimed that:

The truth is that thoroughgoing realists in ethics, such as Plato and Moore, have rested their realist views chiefly on their belief that they are aware of such a property as

³⁰ Brink 1989: 23.

³¹ Brink 1989: 23, footnote.

³² This is not to say that there will not be widespread agreement among cognizers about certain moral judgments. Presumably, nearly everyone will agree that burning a cat for no reason is wrong; the point is just that people will disagree about whether this fact appears to obtain in the way the realist thinks it does.

goodness. And the truth is also that their opponents have rested their views chiefly on their belief that they are not aware of such a property. (1989: 58)

If this is correct, then pace Brink and McNaughton, there is no single moral phenomenological response, and hence no single natural moral stance. It suggests, instead, that moral phenomenologies vary across individuals (and perhaps also that there may be some verbal disagreement involved).³³

Moreover, not only do moral phenomenologies vary across individuals, individual phenomenologies may themselves contain conflicting elements where realism is concerned. For I take it to be a simple and obvious fact about moral psychology that morality may seem to the same person to be both objective (in some respects) and subjective (in other respects). Indeed, Brink himself seems to possess this sort of moral phenomenology. For he refers to one of the factors that commonly motivates antirealism, i.e., the putative "occult" nature of moral facts, as "apparent." But that is exactly the status that he ascribes to the putative inquiry-independent reality of those facts. And if both features are apparent to a person, then there would seem to be little basis for regarding one element as germane, and the other foreign, to that person's moral phenomenology. Of course, the apparent reality of such facts may strike a person more forcefully (i.e., as 'more apparent') than their apparent occult nature. But perhaps also they may each seem equally evident to a person, while other factors, having little to do with phenomenal appearances, may drive a person to realism (perhaps a person just feels better believing values are real).

³³ It is noteworthy that some theorists seem to adopt a purely cognitive construal of moral commitments and their surrounding psychology, to explicitly exclude moral 'experiences'. For example, Mark Platts (1980: 79) has argued that it is precisely the lack of a phenomenological feel that distinguishes the non-appetitive desires characteristic of moral sensitivity from appetitive ones, and that lacking phenomenological fortification our moral commitments require a realist construal to be sustained.

In the latter case, one could plausibly regard Brink's project as subversive of the person's 'natural' inclination to see in moral facts an occult mystery.³⁴

Phenomenological heterogeneity is further evidenced by the diversity of objectivist metaethical theories that have been proposed, many of which are regarded by their proponents as a vindication of commonsense morality or moral experience, etc. But what is the nature of morality according to these writers? For Brink, moral values are properties of objects, not unlike the nonmoral macrophysical properties that are realized by their underlying microphysical composition. For David McNaughton and John McDowell, values are akin in objectivity to secondary properties, such as colours. For Thomas Nagel (1986: 140-143), they are "normative reasons" for action. For Alan Gewirth (1979), morality is objective because there are certain impartial demands placed on us by our status as agents of a certain sort. While these writers will probably agree that morality is, in some sense, 'objective', we have no reason to suppose that they are giving multifarious interpretations of a common phenomenology. In general, regarding a reason for action as valid is phenomenologically different from perceiving a property as real. One should therefore expect the feeling of being compelled to do something by 'normative reasons' to be different from that of perceiving supervening moral properties, and thus that these writers arrive at their considered views from different phenomenological starting points. For Nagel, morality really does 'feel like' the constraint imposed by rationality, while for Brink, the demand is more like that which is imposed on us in confronting the world from a roughly empirical, scientific viewpoint. The phenomenological evidence, though broadly

³⁴ Part of the difficulty one encounters in trying to settle this issue involves the question of how 'sophisticated' theoretical insights have to be in order to be regarded as foreign to the natural stance. How deeply and carefully must one reflect on one's 'initial' ethical position in order for the result of the reflective process to be regarded as foreign to that position? There is probably no clear answer to this question, but it does not seem that a correct phenomenological account of that position must always, and for everyone, regard the perception of an occult mystery in moral facts as foreign.

'objectivist', nonetheless conflicts; even supposing that objectivism in general is a pervasive natural inclination, it is not open to any of these writers to suppose that his prereflective inclinations are 'natural', nor that all of us ought rationally to presume the truth of any particular objectivist view in light of that broadly objectivist natural stance.

Some philosophers have argued that sustaining a moral commitment requires that one understand oneself to be making a claim about an inquiry-independent moral reality.³⁵ This is a psychological claim, and not a claim about the rationality of moral commitments. Its relevance as a psychological claim to the discussion here is that, if it is true, moral commitment becomes evidence of a pervasive commitment to moral realism. But there is evidence that it is false. First, many philosophers seem satisfied to show that moral convictions are objective via means other than realism (e.g., that they are reasonable to hold). Thus, even if some type of objectivism is required to sustain moral commitments, it need not be of the realist type.³⁶ Second, many persons manage to sustain their moral commitments even after protracted reflection on the nature of those commitments has led us to, or perhaps confirmed a nascent inclination toward, moral antirealism (as does, for example, Mackie 1977: 105ff.). And if we are able to do that, it seems we should with even greater ease sustain our initial moral convictions in the face of worries about apparent occultism. If we assume that this ability is pervasive, then the fact that a person is able to sustain moral commitments cannot itself be regarded as evidence that she possesses a realist moral phenomenology.

³⁵ See Wiggins 1976: 341; Platts 1980: 79; Lovibond 1983: 9.

³⁶ In fact, the requirement, even as Wiggins presents it, does not seem to be a belief in moral realism. For Wiggins, it is "objective" reasons that the will "craves" (Wiggins, 1976: 341). But objective reasons need not cite real moral properties. They may be Kantian reasons, or Gewirthian reasons, etc.

All of this suggests that there can be, and is in fact, considerable moral phenomenological heterogeneity. Given this, one can expect an application of the PC to the moral case to lead to the kind of dialectical stalemate that Forgie describes. Suppose, for example, that Harman and Brink round a corner and encounter hoodlums about to set fire to a cat. It seems to Brink that the hoodlums, and their actions, possess the real property of badness, whereas it seems to Harman that they are bad, but not in virtue of their possessing any real moral property. By the PC, both Brink and Harman are entitled, given their respective phenomenological responses, to presume that their experiences are veracious: Brink can presume that the children really are bad (in the realist's sense), and Harman can presume that they are not bad in that sense. But the PC provides no basis for either of these observers to insist that observers whose experiences are at odds with their own must adopt their dialectical stance. And presumably realists such as Brink and McNaughton are making a more ambitious claim than this; presumably, they seek to establish a thesis about what any reasonable person interested in metaethics is warranted in presuming.

Finally, it is noteworthy that Swinburne regarded how things seem not to be as evidence of how things are only given the satisfaction of a certain special requirement. He writes:

Note that the principle is so phrased that how things seem positively to be is evidence of how they are, but how things seem not to be is not such evidence. If it seems to me that there is present a table in the room, or statue in the garden, then probably there is. But if it seems to me that there is no table in the room, then that is only reason for supposing that there is not, if there are good grounds, for supposing that I have looked everywhere in the room and (having eyes in working order, being able to recognize a table when I see one, etc.) would have seen one if there was one there. An atheist's claim to have had an experience of its seeming to him that there was no God could only be evidence that there was no God if similar restrictions were satisfied. They could not be - for there are no good grounds for supposing that if there is a God, the atheist would have experienced him. (Swinburne 1979: 254-255)

This is an important move in Swinburne's argument because it allows him to avoid the sorts of dialectical conflicts described above. Given this, an AMR theorist might contend

that the requirement described here is not satisfied in the case of the moral antirealist either, and might on this basis claim that the aforementioned dialectical conflicts never arise in the debate about moral value. But the requirement mentioned here is satisfied in the case of the moral antirealist. For there is every reason to think that moral antirealists are as capable of detecting moral facts as moral realists, and every reason to expect moral antirealists to have detected moral facts, if they were there and if they were the sorts of things that could be detected by persons with ordinary perceptual and cognitive abilities.

2.4: The PC* Argument

If the arguments of 2.3 are sound, the PC-based argument for the realist's dialectical position is unsound. However, as I mentioned in 2.2, there are exegetical grounds to suppose that Brink never intends to base his position on the PC. Since Brink is "willing to admit that, about moral realism, common belief is silent, divided, or even antagonistic" (1989: 25), it appears that how things appear to moral cognizers where moral objectivity is concerned is irrelevant to his considered view. Rather, his considered view is best represented by the PC* argument, not the PC argument.

Does the PC* argument for the onus shift fare any better than the PC one? There appear to be two difficulties with this argument. First, it is unclear how moral inquiry could be conducted "in order to discover inquiry-independent moral facts," unless inquirers themselves had some sort of commitment to moral realism. Brink himself seems aware of this difficulty, for he often does ascribe a pervasive commitment to moral realism; for instance, he claims that features of moral practice and inquiry "reflect our belief in, or commitment to, the objectivity of ethics."³⁷, and that this belief is germane to

³⁷ Brink 1989: 24; emphasis added. Of course, in light of the previous comments, Brink's claim is more plausible if "objectivity of ethics" is replaced with "the existence of real moral properties."

"commonsense moral thinking." He also says that just as we "begin as realists about the external world (W)e begin as (tacit) cognitivists and realists about ethics" (1989: 23). But he also suggests that the commitment is not a "common belief" about ethics, and that it is not an "unreflective and untutored metaphysical or metaethical view" (1989: 25). Such remarks leave the exact nature of the starting-point commitment to realism unexplained. Given this, it is difficult to assess whether the commitment Brink has in mind is something whose existence is incompatible with the phenomenological facts Harman and Butchvarov describe, and whether moral inquiry can properly be said to be conducted for the purpose of discovering inquiry-independent moral facts in light of that commitment.

Moreover, there is, I think, a more serious difficulty afflicting the PC* argument: unlike the PC, the PC* is not even a plausible epistemic principle. For suppose that certain general considerations about the nature of an area of inquiry are most easily explained on the assumption that that inquiry is undertaken in order to discover evidence-independent truths; this fact does not in the slightest raise the probability that realism about that area is true. The probability of realism about an area of inquiry would presumably be raised in light of assumptions about the area of inquiry only if those facts could not be easily explained except on realist assumptions. But there is no need whatsoever to assume that the inquiry-independent facts or properties that are putatively sought in the conduct of an inquiry really exist in order to explain the ostensibly "cognitivist" features of that area of inquiry. The ostensibly realist and cognitivist features of moral inquiry do seem adequately explained on the assumption that the world has seemed to contain real moral values - at least to those who have had the most influence on the shape of that inquiry over time.

Finally, as we have already noted, although moral realism is a form of moral objectivism, there are nonrealist forms of moral objectivism.³⁸ So even if inquirers are in

³⁸ See, for example, Hare 1981: 206-228; Lear 1984.

general warranted in presuming that some form of moral objectivism is true in light of certain features of moral inquiry, it is not obvious that they are in light of those facts warranted in presuming the truth of the theory Brink favours (i.e., moral realism). Against this, it may be claimed that Brink's construal of moral realism, viz., as just the theory that there are (knowable) "inquiry-independent facts or truths," is itself broad enough to describe almost any form of moral objectivism, and hence that the presumption in question is a presumption in favour of Brink's theory. However, this neglects the fact that during the course of defending his theory Brink succeeds in articulating a very precise account of the objectivity of moral values. On his view, it is clear that morality is objective because there are a posteriori, empirically-discernible, supervening and naturalistically-constituted moral facts. And it would not seem that inquirers could be warranted in presuming the truth of this particular form of moral realism in light of the facts about inquiry in question even if they are entitled to assume the truth of some form of moral objectivism.

2.5: Conclusion

If the arguments of this chapter are sound, Brink is at least tentatively sympathetic to two possible arguments for imposing the burden of argument on the moral antirealist, viz., the PC Argument, and the PC* Argument. Neither of these arguments justify Brink's dialectical stance. The PC Argument fails largely because it is reasonable to limit the PC, and rationally reserve judgment, when dealing with perceptions germane to areas of controversy, and because it is indisputable that controversy exists about the existence of real moral properties. The PC* Argument fails because the PC* is not a viable epistemic principle (the fact that an area of inquiry is generally conducted in order to discover evidence-independent facts does not itself make realism about that area very probably true), and because a key assumption on which it depends, viz., the assumption that moral inquiry

is generally conducted with the goal of discovering real moral facts (of the sort affirmed by AMR), is arguably false.

CHAPTER THREE

ARE MORAL THEORIES EMPIRICALLY TESTABLE?

3.1: Introduction

According to the moral input objection, a cognizer's moral beliefs do not require any input from an inquiry-independent moral reality in order to be coherent. They may be internally coherent, but false. And since different cognizers can and do hold moral beliefs that are equally coherent, but incompatible (a fact that seems best explained without mentioning moral facts) their moral beliefs are probably false. On the AMR view, the moral input objection fails because moral beliefs can be coherent with other (moral and nonmoral) beliefs in ways that tell in favour of their truth (realistically construed). Particular moral judgments can count as evidence for our moral theories because they are best explained on the assumption that the facts they affirm are present and cause us to judge as we do; and moral theoretical convictions are empirically testable in the way scientific theories are. This chapter and subsequent chapters examine these epistemological claims. Sections 3.2 and 3.3 examine the testability thesis as it has been presented by Sturgeon and Brink, focusing on some persuasive examples they give of empirical tests of moral principles. Chapter Four examines whether assumptions about moral facts are ever relevant to reasonable explanations of nonmoral facts. Chapter Five examines the claim that moral assumptions are indispensable to such explanations.

3.2: The Case for Testability

Sturgeon and Brink concede that moral theories are not testable in isolation. However, they argue that since the same is true of scientific theories, this fact about moral

theories constitutes no special flaw. In science, to "test" a theory is to subject it to a procedure that either falsifies or confirms it through observation. Yet it is widely recognized that many important theoretical claims in science have no observational consequences, and are hence untestable, on their own.³⁹ Testing them involves conjoining them with "auxiliary hypotheses," and deducing certain predictive or retrodictive statements. Consider as an example the universal theory of gravitation:

(ST) a law that every body *a* exerts on every other body *b* a force F_{ab} whose direction is toward *a* and whose magnitude is a universal constant g times $M_a M_b / d^2$, and Newton's three laws.

ST is supposed to govern the motion of every physical body (or, more exactly, every "point mass"). Nonetheless, on its own it is compatible with virtually any observed motion; it is not falsified by bouncing balls, rocket flight, perturbations in the orbit of Uranus, and so on. We can test ST only by conjoining it with auxiliary claims that are independently supported. For example, conjoining ST with

(SAS₁) Venus and the sun are subject to no forces except mutually induced gravitational forces, and

(SAS₂) both bodies exist in a hard vacuum,

we can deduce:

(O₁) The motion of Venus will conform to Kepler's three laws of celestial mechanics. The argument functions as an empirical test of ST because the truth value of its consequence, O₁, can be determined empirically. If O₁ fails to obtain (because there turn out to be perturbations in the orbit of Venus not calculable on the basis of ST, SAS₁, and SAS₂) we have reason to reject at least one of the conjoined claims. Which we reject will depend on the support conferred on ST relative to SAS₁, and SAS₂. If many confirmed claims have been deduced from arguments that assume the theory, as is the case with ST,

³⁹ The point is due to Duhem 1962: 186-190.

but one or more of the auxiliary statements is ad hoc, it would be more reasonable to reject the auxiliary claim than ST. Any evidence against one of the auxiliary statements is a reason to discard that statement in favour of one that is better supported, one that conjoined with ST entails an accurate claim about the motion of Venus. If, on the other hand, all of the auxiliary claims have been independently confirmed, O₁'s disconfirmation may pose a genuine problem for ST.

According to Sturgeon and Brink, moral theories are tested in the same indirect way: not in isolation but in conjunction with auxiliary claims that are, as Sturgeon puts it, "of comparable status" (Sturgeon 1984: 51). Sturgeon offers the following example:

(MT) An action is wrong just in case there is something else the agent could have done that would have produced a greater net balance of pleasure over pain;

(MAS) It is always wrong to deliberately kill a human being.

∴ (O₂) Deliberately killing a human being always produces a lesser balance of pleasure over pain than some available alternative act.⁴⁰

Does the procedure Sturgeon describes do what he says it does? A good starting point for a discussion of Sturgeon's argument is a criticism of the testability claim that has recently been given by Jobe (1990: 268-269). According to Jobe, Sturgeon's procedure

is not at all sufficient to show that a principle can be tested and confirmed, for (it) is applicable to vast amounts of nonsense. To take an example, consider the principle (a) "If a person is happy on a Sunday then the Absolute is smiling on him at that time." When combined with a "comparable principle" such as (b) "If the Absolute smiles on a person on a given day then that person will be happy on the next day," we can deduce the consequence that if a person is happy on a Sunday then he will be happy on Monday. So, if Jones, say, is happy on Sunday but not on Monday we can conclude that not both (a) and (b) can be true. But the test outcome does not in the slightest raise the probability of either component claim, and so we have no more reason than before for accepting either one.

That the test outcome Jobe mentions need not result in confirmation of the conjoined sentences may be conceded without also conceding that the sentences in question

⁴⁰ Sturgeon 1984: 51.

are unconfirmable; for the outcome amounts to predictive failure, and normally it is predictive success that is held to result in confirmation. Nonetheless, his objection to the claim that the test result succeeds in falsifying the sentences appears to count against their confirmability too. According to Jobe, the result does not falsify (a) and (b) because it cannot ensure "that the sentences in question even possess a truth value." On this basis, Jobe concludes that Sturgeon's procedure cannot contribute to a rebuttal of noncognitivist moral skepticism. If his reasoning is correct, it seems that no test result could succeed in falsifying or confirming either sentence.

There are problems with Jobe's position, however. One is strategic: even if it is accepted that the counterexample reveals a difficulty with the testing procedure, it is not at all clear that the argument is useful as a critique of AMR. For the procedure described seems common to both moral and scientific testing: both involve deducing an observable consequence from the conjunction of two or more "comparable" principles (i.e., principles germane to the same discipline, employing the same terminology, and so on). Given this, it seems that Jobe's counterexample to it must count equally against scientific testing, and that it does not reveal a problem peculiar to morality. It is true that Jobe claims that his counterexample shows only that Sturgeon's argument is "not sufficient to show that a principle can be tested and confirmed"; but without an account of what, in addition, is necessary for genuine testability (something not provided by Jobe), the argument gives us no reason to regard moral testing as particularly defective.

Perhaps more serious is the fact that Jobe has not cited a genuine counterexample to Sturgeon's procedure; that is, a possible employment of that procedure that nonetheless fails to ensure that the constituent premises bear truth value. Sturgeon's procedure essentially involves making a deductively valid inference from certain conjoined principles. But the notion of deductive validity applies only to meaningful, truth-value bearing

sentences. If, as Jobe appears to believe, references to "smiling Absolutes" are nonsense, then the sentences themselves must lack truth value. But if so, they could not figure in a deductively valid argument. It follows that the piece of reasoning Jobe cites is not a genuine counterexample to Sturgeon's procedure.

It is, moreover, simply a mistake to hold that empirical testing procedures must establish that the tested claims possess truth value. For it is a requirement that no empirical testing procedure could satisfy. Consider the directly testable claim (A) "There is an apple on the table." Suppose the "procedure" in question for testing (A) asks us simply to determine empirically whether there is an apple on the table. Finding no apple, we can conclude that (A) is false. But while this outcome (and the alternative) recommends ascribing some particular truth value to the sentence, neither reveals whether (A) possesses a truth value. Whether the claim bears truth value depends on its meaning, and this is revealed by conceptual analysis, not by empirical tests. So it seems that this procedure too fails to satisfy Jobe's requirement. But this is hardly a reason to reject the procedure or to regard (A) as untestable.

Given that in general truth-value possession depends on meaning, and that meaning is not subject to empirical test, it is unfair to condemn Sturgeon's argument on the ground that it does not support a case against noncognitivist moral skepticism. Moral noncognitivism is a theory of the meaning of moral sentences, and is refuted by factors that bear on issues of meaning and not by the results of any empirical tests. The situation is similar in science, where analogous to moral noncognitivism is the skeptical position known as "semantic instrumentalism."⁴¹ Semantic instrumentalism is an account of the meaning of scientific claims that denies that scientific theories bear truth values; it says that

⁴¹ Newton-Smith uses this term to describe a position he attributes to Ernst Mach. See Newton-Smith 1983: 30; Mach 1960: 577.

they function as mere predictive devices. As in the moral case, its truth is tested conceptually, not empirically. Thus it can be rejected if it fails "as a claim about how we do in fact regard (scientific) theories."⁴² In neither area does the failure of an empirical testing procedure to contribute to the case against this form of skepticism reveal a flaw in the procedure.

3.3: MT and the Independence Requirement

Notwithstanding the failure of these criticisms, one may still question whether the utilitarian principle Sturgeon cites can be tested in the way that he supposes. One potential problem is posed by the fact that testing a theory properly using the proposed method requires that the auxiliary statements be independently supported. (This is because deducing a false prediction shows only that at least one of the conjoined statements is false; having independent support for auxiliary statements shifts the blame for predictive failure to the theory in question.) But while scientific principles such as the ones Sturgeon mentions can meet this condition, it is not obvious that his utilitarian principle can.

To appreciate why this is so, consider the nature and rationale of the independence requirement. What is normally thought to be required when testing scientific theories is support for auxiliary hypotheses that one can reasonably accept while suspending judgment on the theory in question. Of course, in science, the support accorded to auxiliary claims will normally be observational; it will be reasonable to accept an auxiliary claim if it passes certain observational tests, and since hypotheses are typically tested in bundles, those tests will themselves require making still other auxiliary assumptions. Satisfying the independence requirement is thus normally a matter of conjoining the theory to be tested with auxiliary claims that succeed in passing further observational tests in which they are

⁴² Newton-Smith 1983: 30.

themselves conjoined with yet further assumptions - assumptions that are logically independent of the original theory to be tested. As Kitcher has remarked:

The crucial point is that, while hypotheses are always tested in bundles, they can be tested in different bundles. An auxiliary hypothesis ought to be testable independently of the particular problem it is introduced to solve, independently of the theory it is designed to save.⁴³

To see why such tests must employ auxiliary claims with support of this sort, consider what must happen if one were to attempt to apply the test by employing auxiliary claims that lacked such support. Suppose we wish to test a theory T by conjoining it with an auxiliary claim AS and deducing an observational claim O. In order to blame predictive failure on T and not AS, we must have reason to suppose that AS is true. But suppose that we could not assess whether AS is true without assuming a position on the truth of T; i.e., that our only ground for accepting AS is that it passes a test (or a number of tests) in which T (or its denial) figures as an auxiliary assumption. Suppose further, to take one of the two alternative interpretations of T that might figure in a test of AS, that the only rationale for AS is that it has passed a test in which it is assumed that T is true. Deducing a false observational claim from the conjunction of T and AS would not falsify T; for to establish T's falsity would be, ex hypothesi, to remove our only basis for keeping AS in case of predictive failure. If, on the other hand, the only rationale for AS is that it has passed a test in which it is assumed that T is false, then the test of T would be otiose: one would be justified in rejecting T on the basis of O's failure to obtain only if one is already justified in supposing that T is false.

Clearly, scientific principles such as the ones Sturgeon mentions can meet this condition. In the test of ST, we can assess whether SAS₁ and SAS₂ are true without

⁴³ Kitcher 1982: 48. The availability of independently testable auxiliary statements distinguishes what he calls "successful" or genuine science from "pseudo-sciences" such as creationism.

assuming any particular position on ST: for instance, spectrographic investigation may indicate the presence or absence of any particles of matter in interplanetary space, and telescopic exploration should be capable of revealing the presence or absence of any bodies that might affect the motion of Venus with respect to the Sun. Neither of these investigations require assuming a position on ST .

Or, to take another example, suppose that a theory about the behaviour of bodies within the earth's gravitational field contains the following law giving the distances which a body falls from rest in that field.⁴⁴

$$L: s = 1/2gt^2$$

While L is untestable on its own (since, as it stands, it has no capacity to predict how any particular body will behave), we can test it using the proposed procedure by conjoining it with certain independently-supported assumptions and deducing a predictive claim about some particular body B. For example, we can assume that B is and will remain during the course of the experiment within the earth's gravitational field, that certain initial spatial and temporal conditions define B at rest on a particular occasion, that the gravitational constant g has a specific value, and that our measuring devices are accurate.⁴⁵ These auxiliary claims can be assessed without making any assumptions about the truth of L. We can empirically check the initial conditions that define B at rest by using a ruler and a clock; we can verify that our ruler is rigid and that our timepiece is accurate; there is also an experimental procedure (the Cavendish experiment) we can use to determine the value of g , one whose results we can assess without knowing whether L is true.

⁴⁴ The example is from Newton-Smith 1983: 71.

⁴⁵ The latter, of course, will be determined by some other area of gravitational theory.

However, it is not obvious that Sturgeon's proposed test of MT meets this condition. One might reject his test of this principle on the following grounds.⁴⁶ Having an independent case for the claim that killing is always wrong must, it seems, involve citing wrongmaking aspects of deliberate killing not including the property mentioned by MT, and citing a principle or set of principles in virtue of which these features are wrongmaking (principles that are logically independent of MT). For instance, one might claim that deliberately killing humans is always wrong because it interferes with a person's ability to pursue their rational creative projects, and because such interference is, in principle, wrong. However, since MT says that an action is wrong if and only if it possesses that property, it denies that an action might possess any wrongmaking properties apart from the one it mentions.⁴⁷ Hence, to cite wrongmaking features of deliberate killing other than the mitigation of pleasure would be to cite considerations that show that MT is false. But this botches the test, for accepting such evidence for MAS requires rejecting MT regardless of whether O₂ obtains.

This objection, as it stands, is untenable, for it overlooks the fact that MT does not logically rule out the possibility of there being wrongmaking properties that are coextensive with the property it mentions. Thus, suppose G is the property mentioned by MT. MT does not rule out the possibility of there being another wrongmaking property H that always occurs with G (in the way that the property "is a creature with a kidney" always occurs with "is a creature with a heart"). Given this, one might very well support MAS via an appeal to the fact that deliberate killing bears H, and to some (logically independent) moral principle in virtue of which H-bearing actions are wrong - while suspending

⁴⁶ See Yasenchuk 1994.

⁴⁷ I take it that the phrase "just in case" expresses a biconditional relation (equivalent to "if and only if"). This, at any rate, is how it is normally taken. See, e.g., Bergmann, Moore, and Nelson 1980: 35; Hodges 1977: 99.

judgment about MT. For instance, one might cite the "Kantian" principle MT* (which is logically independent of MT) in virtue of which deliberate killing is wrong, i.e.,

(MT*) An action is wrong if it treats humanity as a means rather than as an end in itself.

Then, provided there are grounds to suppose that

(A) There is something their agent could have done that would have produced a greater net balance of pleasure over pain just in case the act in question treats humanity as a means rather than as an end in itself,

MAS would have support of the required kind. Now, the association relied on here (i.e., [A]) is clearly improbable; too often the most effective way of promoting utility is to treat others as a means, and so on. Indeed, it is not implausible to suppose that there are simply no such associations of moral properties, for this supposition seems to explain why cognizers who must decide what to do in particular situations so frequently must decide among conflicting moral considerations. Nonetheless, some moral philosophers seem to recognize the existence of contingently-associated "clusters" of goods, and they would presumably accept such associations of other properties (including wrongness) as well.⁴⁸ On such a view, certain basic properties are unified by certain social and psychological causal mechanisms, and hence always (as a matter of fact) occur in tandem. So it seems that, given such a view, some MT-type principles might be tested in a way that satisfies the independence requirement. Moreover, such an account appears to place MT on par with its scientific cousins in terms of testability. For similar considerations explain why we can often test scientific principles in a way that satisfies the independence requirement even if

⁴⁸ Something like this idea appears in Boyd's account of his "homeostatic consequentialism" (1988: 203). According to Boyd, goods in homeostasis with one another are united by "psychological and social mechanisms." I doubt that Boyd would agree that all of the properties that make a homeostatic cluster are inextricably linked in virtue of their membership in that cluster, so that they all always occur together. Still, it is not implausible to suppose that at least some homeostatic properties are connected like this to one another.

they are formally equivalent to MT. Natural kinds are often, even typically, understood as possessing a plurality of coextensive observable properties. For example, it is generally accepted that water, and only water, has the capacity to sustain aquatic life, and also that water, and only water, expands when frozen. So, consider the following configuration of sentences:

(WT) A solution is water just in case it expands when frozen.

(WAS) The liquid in the flask is water.

∴ (O₃) The liquid in the flask will expand when frozen.

A scientist might support WAS by citing certain of the other observable properties that water, and only water, is known to possess, while suspending judgment about WT. For although WT says that water and only water possesses a particular observable property, it does not rule out the possibility that water might possess yet other properties that are coextensive with the property it mentions. So, for instance, one could support SAS via the claim that water, and only water, supports aquatic life, together with the claim that the liquid in the flask supports aquatic life, without having to suppose that WT is false.⁴⁹

However, the MT/WT analogy breaks down for the following reason. Notice that certain moral theorists, notably, monistic hedonists, could not test MT in this way. For presumably such persons not only think that actions that produce the least overall pleasure in a given circumstance (and only those actions) are wrong; they also think that that

⁴⁹ Some water properties, of course, are not coextensive with the property of expanding when frozen (e.g., water and ethyl alcohol have the same appearance under many conditions; so the property of being "clear at room temperature" is not coextensive with the property of expanding when frozen), yet the fact that the liquid in the flask bears such a property can be prima facie evidence that it is water. Nonetheless, to test WT, a scientist would need relatively strong evidence for WAS (so that whatever observational result occurs could not plausibly be explained on the assumption that WAS is false). And she would have such evidence if she knew that the liquid in the flask bears a property (or a combination of properties) that water, and only water, possesses.

particular property is the only wrongmaking property an action might possess. But even though MT does not completely represent the theory of hedonistic monism, it does assign wrongness to exactly those actions that a hedonist would want to condemn; hedonistic monism arguably implies MT. So determining whether hedonistic monism is true would appear to require assessing MT. Yet because the hedonistic monist recognizes no wrongmaking properties other than that mentioned by MT, he could not test MT in a way that satisfies the independence requirement. By contrast, no scientist is likely to deny that the presence of water may be marked by a property (or specific complex of properties) other than the property mentioned by WT; indeed I believe that there is no candidate theory of any natural kind that holds that the presence of that kind can only be marked by a single observable property. Now, one might think that this difference should hardly trouble an AMR theorist. For it might be urged that moral monism of any sort is simply untenable (i.e., that monists simply have not reflected carefully enough on the matter and that if they did they would discover that monism of any kind just fails to cohere with their considered moral judgments), and that, this being so, the objection seems to establish at most the untestability of a class of intuitively unappealing moral theories. And it may be thought that since we ordinarily pursue our moral inquiries by assessing interesting theories (i.e., those theories that appear to be capable of accommodating our considered moral judgments), not uninteresting ones, the purported untestability of monism (or, more exactly, the inability of monists to test their own theories) should not prevent more reflective inquirers from conducting moral inquiry as a partly empirical one. However, one should recall that while hedonistic monism and other monisms may currently be in disfavour, this has not always been the case; monisms were once taken very seriously. And it would seem to be a deficiency of the proposed method of theory assessment that it could not have been relied on in order to assess a class of theories that once appeared

plausible. For I can think of no example of a scientific theory of a natural kind (e.g., water, gold, etc.) that at some point in time appeared plausible, but that has faced, or faces, this sort of difficulty. And if, as Brink and Sturgeon appear to think, scientific theories are reasonably refuted and confirmed via this method (or at least on the basis of scientists' educated expectations of what they would observe under certain conditions), the problem would appear to lie with the nature of moral inquiry and not with the Duhemian testing method.⁵⁰

Moreover, the difficulty arguably afflicts another moral theory type that is not clearly counterintuitive and in current disfavour, i.e., the unifying principles of "unified pluralistic" moral theories, including the theory favoured by Brink, Objective Utilitarianism (OU). According to Brink, unified moral theories "explain why things to which we attach moral value are all of moral value and attempt to organize different moral considerations in some fairly systematic order" (Brink 1989: 250).⁵¹ OU is unified but pluralistic because it recognizes a variety of distinct goods (e.g., the pursuit and success of rational creative projects, friendship, etc.) whose value is explained by the fact that they are components of

⁵⁰ Of course, even if a moral theory cannot be tested using the Duhemian method, it might be empirically supported in another (albeit even more indirect) way: moral theories are clearly assessed against considered moral judgments (if a theory cannot accommodate an appropriately broad range of CMJs, one will hardly bother testing it against observations!). And it may be contended that CMJs are accepted and rejected on the basis of input from an inquiry-independent moral reality. This would be plausible if, for instance, the best explanation of why a moral property-ascribing CMJ is accepted is that the property it ascribes is present and has caused the cognizer to make the judgment. Given this account, a theory such as MT could be said to have empirical support in light of its coherence with those CMJs for whom such an account is plausible.

⁵¹ Fragmented theories, by contrast, recognize a plurality of "distinct moral considerations that are connected by no common moral principles." Fragmented theories are always pluralistic. But unified theories can be pluralistic or monistic. Hedonism, for instance, is unified (by the single value it recognizes), and monistic (because it recognizes only this value), but OU is unified but pluralistic because although it recognizes a plurality of distinctly good things, their goodness is explained on the assumption that they are components of a more general good. See Brink 1989: 250-251.

human welfare (Brink 1989: 250-251). Presumably, the unifying principle of Brink's moral theory is just the claim that

(OU) An action is right just in case it realizes more human welfare than any available alternative action.

But it is not obvious that this principle could be conjoined with a substantive judgment that enjoys "independent support." Suppose we wished to test OU using Sturgeon's Duhemian method. OU, conjoined with the (plausible) moral judgment that

(MAS₂) It is always right to help little old ladies across busy intersections if they ask for such assistance,

entails

(O₄) Helping little old ladies across busy intersections when they ask for such assistance always realizes more human welfare than any available alternative act.

Having an independent case for MAS₂ would seem to require citing rightmaking aspects of helping little old ladies quite apart from the capacity such actions might have in advancing human welfare, and appealing to (logically independent) principles in virtue of which these features are rightmaking. But now a parallel objection to the one raised above may be raised here: although OU, like MT, does not rule out the possibility of there being such rightmaking properties (since it does not rule out the possibility of coextensive rightmaking properties), no proponent of OU will likely recognize such properties. So it seems that such persons could not test their favoured theory. Moreover, many philosophers write as though morality essentially involves a concern for human welfare (that this is part of the concept of morality); so there may very well be a 'second-order' conceptual constraint that precludes the satisfaction of the independence requirement in these cases. Finally, even if this is not so, and a cognizer (perhaps inspired by the work of so-called "deep-ecologists") may recognize the value in actions that promote goods other than that of human welfare, satisfaction of the independence requirement would appear unlikely. For its satisfaction requires not only the recognition of rightmaking properties other than that mentioned by

OU, but also the supposition that any extraneous rightmaking properties are coextensive with that of advancing human welfare. And this last supposition seems implausible: often the pursuit of human welfare requires the sacrifice of ecological goods, and vice versa.

My objection to Sturgeon's test might still be challenged in another way. Notice that in the water case, all the scientist needs to support WAS and thereby test WT is evidence that the sample of liquid is water, i.e., something that indicates the presence of water not including the property mentioned by WT. And this need not be (and would not even normally be) some fact about the constitution of the sample in virtue of which it is water, such as its molecular structure, and so forth (i.e., some water-making property). It could just be an observable property (or set of properties) that (he has reason to think) is possessed by water and only by water (or by water and only a few other rarely encountered substances), and whose possession thereby makes it quite likely that the sample is water. Accordingly, it might be suggested that the AMR theorist need not support MAS by citing any wrongmaking property, and that she only needs to cite a fact that can serve as evidence of the truth of MAS. Now, an exact analogy with the WT case would appear to demand that she cite some observable feature of the action-type mentioned by MT that is not wrongmaking but that is nonetheless known to be peculiar to wrong actions (such a demand would be satisfied if, for instance, wrong actions were known to have, say, a distinctive odour, or a certain social effect, etc.). Yet it is not obvious that there is any such property; indeed, it seems that any observable feature that an act might possess that might be regarded as germane to all and only wrong actions would also appear to be something that helps to make the action wrong. Nonetheless, an AMR theorist could still point to evidence for claims such as MAS, albeit of a slightly different sort. Recall that, on the Brink-BonJour account, CMJs - like observational judgments - "have initial credibility and are not just initially believed" (1989: 136). On such an account, it seems that MAS might

be (at least partly) justified on the basis of second-order beliefs about its status as a considered moral judgment and about the reliability of such judgments. The fact that MAS survived a process of sober and impartial contemplation and seemed compatible with one's other moral beliefs would itself be evidence of the truth of MAS, and would be support that it is independent of MT in the required respect (since it includes no claim about the property mentioned by MT).

This seems to cohere nicely with the general coherentist spirit of the Brink-BonJour account. Moral principles like MT are testable, on this view, not directly but only when conjoined with auxiliary moral assumptions, and only given certain second-order assumptions about those auxiliary claims together with (plausible) epistemological principles that accord credibility to such claims. Two problems remain, however. First, as I suggested in Chapter One, since we have no clear understanding of how CMJs could be accepted as a result of input from the moral properties they affirm, there is no viable second-order account of CMJs that establishes their initial credibility. Second, the proposed view of moral theory testing requires that one have relatively strong grounds for accepting the auxiliary claim - at least, those grounds must be stronger than any one might have for accepting the theory one wants to test. This is, again, because one needs to lay blame for any predictive failure on the theory in question and not on the auxiliary assumption. But suppose one's basis for accepting the auxiliary claim MAS is, as the present account claims, just that it is a CMJ and that as such it is reliable - reliable, as Brink explains, because as a considered moral belief it has been formed under conditions of general cognitive reliability, and so forth. Having relatively less support for MT would seem to require that MT be accepted only under comparatively less favourable conditions of belief formation, and that it not be accepted when those conditions are met.⁵² But the

⁵² It might be claimed that we could test spontaneous moral beliefs using the proposed method because these are not as well grounded as CMJs. But suppose that a

fact that a belief has this etiology is arguably sufficient grounds for discarding it. And this seems to point to a serious constraint on the test, for now it seems that we can apply the test only to those MT-type principles that are untenable anyway, but not to plausible ones (as we do in science).

3.4: Can Any Moral Principles be Tested Using Duhem's Method?

Of course, the considerations raised here do not show that no moral principle can be tested against nonmoral observations. If sound, they establish at most that (a) whether a moral theory can be tested depends on whether it can be conjoined with independently supported moral judgments, (b) whether this can be done depends partly on the sort of principle or theory it is (i.e., on its formal features), and (c) certain possibly attractive moral theories (including Brink's) are not testable along the lines proposed. This alone seems somewhat damaging to AMR, for (as I have already mentioned) it implies that many moral theories that are of current interest to serious moral inquirers, as well as many theories that have inspired great interest in the past, could not be tested in the way (formally equivalent) scientific principles can be. And this seems to represent a significant constraint on the conduct of moral inquiry; one that is not present in the scientific case.

Nonetheless, clearly many (perhaps most) moral principles are unscathed by the objection of 3.3. Consider as an example Brink's proposed test of the principle "Good people keep their promises when doing so involves great personal sacrifice":

belief p is a spontaneous moral belief at time t_1 , that we subject it to a Duhemian test at that time by conjoining it with a CMJ, and that it passes the test. Our next epistemic task would be to consider p (by assessing it under conditions of general cognitive reliability). So, suppose that at t_2 p has been considered, that it passes this assessment, and that it is therefore a CMJ at t_2 . We would, it seems, have to assess the observational test conducted at t_1 in terms of p 's status at t_2 (to ignore the additional information that has been adduced about p after t_1 would seem to be epistemically irresponsible). But at t_2 p is, ex hypothesi, as well supported as the auxiliary CMJ. So the observational evidence adduced at t_1 could not be regarded as germane to a test of p at t_2 .

My belief that good people keep their promises when doing so involves great personal sacrifice has no observational consequences when taken in isolation. But when I conjoin it with my independently supported moral belief that Zenobia is a good person (i.e., my evidence not including Zenobia's promise-keeping behaviour), I can obtain the observational consequence that Zenobia will keep her promise to Zelda, even though doing so will involve great personal sacrifice on Zenobia's part. (1989: 137)

The conjoined belief in this example can apparently meet the demand for independent support. One could, it seems, cite some aspect of Zenobia's character or behaviour (apart from her inclination to keep her promises) as evidence of her goodness while suspending judgment about whether promise-keeping makes a person good. Now, if this and other moral principles can be tested empirically, then the desired parallel with science would appear sustainable, notwithstanding the criticisms expressed in 3.3. For it can plausibly be maintained that the desired analogy requires only that some moral principles can be tested empirically, and that other moral claims just be coherent with the principles that pass such tests; it does not imply that every moral principle be testable using the Duhemian method. For although some scientific claims have empirically determinable consequences (at least when conjoined with certain auxiliary statements) and are thus empirically testable, many scientific claims do not (e.g., certain claims about the origin of the universe or about subatomic particles). So it might be contended that in both science and in ethics there are constraints, of a very general sort, on which claims are testable - constraints that are quite compatible with a realist coherentist account of those disciplines, since on such an account, as long as some of the constituent claims in a discipline are empirically testable, other beliefs in that discipline may be justifiably held (and taken to "correspond" to features of reality) in so far as they cohere with those beliefs.

However, there do seem to be grounds for accepting the much stronger claim that no moral claim can be tested using the method in question. For there is more to testing a scientific theory than simply conjoining it with independently supported auxiliary

statements and deriving observational claims. A theory and a conjoined set of independently supported auxiliary statements may entail true predictive claims without being confirmed. The reason is that there are an infinite number of generalizations compatible with any given set of observational data; even rival theories may (in conjunction with certain auxiliary statements) entail the same observational prediction. Confirming a theory is thus normally a matter of selecting it from among a plurality of contenders on the basis of its capacity to explain the observations that it (and certain independently-supported auxiliary statements) entails.⁵³ In short, in science, a theory is confirmed by an observation via the Duhemian method only when the best explanation of the observation includes the assumption that the theory is true. A similar point can be made about falsification. Scientific theories may be conjoined with independently-supported auxiliary claims without being falsified. Since, as a matter of fact, virtually every important scientific hypothesis has had to contend with some observational deviations, predictive failure is normally taken to falsify a theory only when the best explanation of the observation is the falsehood of the theory.⁵⁴ Now, one would expect the observational claims that figure in any putative moral test to be similarly overdetermined: for any given observational claim there should be a multitude of (often incompatible) theories and auxiliary claims (both moral and nonmoral) that would yield the same observational consequence. So the analogy would seem to require that the primary basis for theory selection be, as it is in the scientific case, the relative explanatory power of the contending theories; i.e., that in morals, as in science, falsification occurs only when the best explanation of an observation is that a theory is false. Thus, in inquiring whether Brink's promise-keeping principle is testable in the proposed manner we must determine not only whether Zenobia will indeed keep her

⁵³ See Goodman 1954: 87-120.

⁵⁴ Miller 1983: 758-759.

promise to *Zelda* at great personal cost, but also whether the best explanation of *Zenobia's* behaviour includes the assumptions that good people keep their promises and that *Zenobia* is a good person. Now, it remains to be seen whether nonmoral facts such as the fact that *Zenobia* keeps her promise(s) to *Zelda* can be best explained on such moral assumptions; if *Zenobia* does keep her promise(s) to *Zelda*, perhaps the best explanation of this is just that *Zenobia* has been raised in a certain way, i.e., to have certain moral or religious beliefs and that people given this upbringing will behave the way *Zenobia* does. If this nonmoral account turned out to be the best explanation of *Zenobia's* behaviour, one could not regard the deduction Brink mentions as an empirical test of the promise-keeping principle. Of course, if Sturgeon and Brink are correct in supposing that nonmoral facts can be best explained on moral theoretical assumptions, then moral theories can meet this further requirement for testability. But if the arguments of Chapter Four are sound, moral theories are irrelevant to explanations of the nonmoral facts, and could not, therefore, be regarded as testable against nonmoral observations.

3.5: Conclusion

If the arguments of this chapter are sound, many of the moral beliefs an inquirer might hold will not be coherent in the ways that are thought, on the AMR account, to be evidence of their truth (realistically construed). I have identified a number of technical reasons to reject the view that we can generally employ the Duhemian testing method to assess moral principles. Whether a moral inquirer can test her favoured theory depends on the logical features that theory possesses, and on the specific moral commitments she is inclined to make or to avoid. As a matter of fact, many moral theories, including the example Sturgeon mentions, and Brink's favoured moral theory could not be tested in the proposed manner. And if the arguments of Chapter Four are accepted, no moral theory can be tested empirically.

CHAPTER FOUR

EXPLANATORY RELEVANCE

4.1: Introduction

The arguments of Chapter Three, if sound, provide good reasons to suppose that many of the moral principles an inquirer may want to test, including the one contained in Sturgeon's proposed example, cannot be tested in a way that satisfies the independence requirement. More importantly, they suggest that whether a moral theory is empirically testable depends on whether it is possible to explain certain nonmoral observations (such as the observation that Zenobia has kept her promise to Zelda), by citing moral facts (such as the fact that Zenobia is a good person and that good people keep their promises). Thus, assessing whether moral theories are empirically testable must involve assessing their capacity to explain these and other nonmoral observations. Now, as I indicated in Chapter One, the question of whether moral theories have certain explanatory powers bears on the integrity of AMR's broader justificatory strategy. For on the AMR account, moral claims are justified by their coherence with other moral and nonmoral claims, and such coherence is largely constituted by the ability of moral claims to explain various moral and nonmoral facts. And since we should expect the belief sets we would adopt in pursuing a coherentist method of justification to include plausible second-order beliefs about what kinds of beliefs moral beliefs are (especially about the conditions under which they are adopted) and about the general reliability of beliefs of this kind, the membership of moral beliefs in such coherent belief sets is itself a reason to suppose that those beliefs are true (i.e., that the facts they affirm obtain). So, exploring the question of whether moral claims are empirically

testable leads us naturally to assess the general account of justification proposed by AMR theorists.

Now, the heart of the moral explanations issue is the question of whether moral claims can be justified via a particular sort of abductive inference, viz., via an inference from the claim that certain nonmoral facts or observations are best explained on the assumption that a moral fact-affirming statement is true, to the truth of the moral assumption.⁵⁵ In the rest of this thesis, I shall grant the AMR theorist the assumption that the abductive justificatory strategy is generally sound, i.e., that if certain observations could be best explained on a moral assumption, this would be a good, normally sufficient, reason to think that the assumption is true, and concentrate on the question of whether any observations are best explained on moral assumptions.⁵⁶ AMR theorists deal with this question in two ways. First, since explaining nonmoral facts on moral assumptions clearly requires that moral assumptions be relevant to reasonable explanations of nonmoral facts, they attempt to rebut recent arguments for the claim that moral assumptions must be irrelevant to such explanations. Second, they argue that moral assumptions are indispensable to reasonable explanations of nonmoral facts. In this chapter, I examine recent arguments for and against the claim that moral facts are or can be relevant to the explanation of certain nonmoral explananda (including moral judgments). Chapter Five assesses the main arguments given by Sturgeon and Brink for the indispensability claim.

⁵⁵ For an analysis of the inference to the best explanation, see Harman 1965; Lipton 1993.

⁵⁶ It is noteworthy that some (non-AMR) moral realists reject the "best explanations" approach entirely. See e.g., Butchvarov 1989: 6-7.

4.2: Harman's Case for Irrelevance

The claim that moral facts are relevant to the explanation of nonmoral facts is precisely what Harman (1977; 1986) has denied. Harman claims that the central "problem with ethics" is that it contrasts with the scientific case in the following way:

The observation of an event can provide observational evidence for or against a scientific theory in the sense that the truth of that observation can be relevant to a reasonable explanation of why that observation was made. A moral observation does not seem, in the same sense, to be observational evidence for or against any moral theory, since the truth or falsity of the moral observation seems to be completely irrelevant to any reasonable explanation of why that observation was made. The fact that an observation of an event was made at the time it was made is evidence not only about the observer but also about the physical facts. The fact that you made a particular moral observation when you did does not seem to be evidence about moral facts, only evidence about you and your moral sensibility. (Harman 1977: 6)

To make his point, Harman contrasts the case of a physicist (call him Kildare) who sees a vapour trail in a cloud chamber and thinks "there goes a proton," with the case of a moral appraiser (call him Spike) who encounters children on a corner about to set fire to a cat, and who thinks "those children are doing something morally wrong." According to Harman, the factors that are relevant to the explanation of Kildare's judgment include both his "psychological set" (esp. his commitment to certain physical theories), and, crucially, the fact that a proton passes through a cloud chamber. By contrast, Spike's judgment is sufficiently explained on the assumption that his moral sensibility is such that he is inclined to condemn cat burning (i.e., his acceptance of certain more or less well-articulated moral principles), and that he sees an instance of cat-burning. In fact, it seems to Harman that any moral facts, such as the fact that cat burning is "really wrong," are completely irrelevant to the explanation of the judgment. It is, he claims, "completely irrelevant to our explanation whether your intuitive immediate judgment is true or false" (Harman 1977: 7).⁵⁷

⁵⁷ This construal of the judgments in question may suggest that Harman and Brink are discussing different classes of candidate explananda items. For it might be urged that

Why are moral properties irrelevant to the explanation of moral observations? The source of the difficulty seems to be that moral properties, in contrast to observable macrophysical "properties" such as vapour trails, are epiphenomenal; they can be effects but not causes. Harman assumes that in order for a candidate explanatory factor to be relevant to the explanation of a person's judgment, it must be possible for it to be part of the causal history of that judgment.⁵⁸ This requirement is met in the proton case: the proton causes the vapour trail, which in turn (in conjunction with the fact that Kildare accepts certain scientific theories regarding protons) causes Kildare to judge as he does.⁵⁹ In fact,

"immediate intuitive judgments" are more like initial moral judgments (IMJs) than considered moral judgments (CMJs). The familiar idea is that we collect an individual's initial prereflective moral responses and filter them to include only those judgments that one would accept under conditions generally conducive to avoiding errors (e.g., when one is calm, sober, and adequately informed) (see, e.g., Daniels 1980). So Harman's thesis is perhaps just that assumptions about moral facts or truths are irrelevant to a reasonable explanation of IMJs. But Brink makes no claims about the reliability of IMJs. His response to the moral input objection is that CMJs have initial credibility (in virtue of their own peculiar etiology), not that IMJs do. Now, this may very well commit Brink to some thesis about the capacity of assumptions about moral facts to explain CMJs, but it commits him to no such thesis about IMJs. So if the appropriate analogue to Harman's scientific observations are just IMJs, then perhaps Brink can concede Harman's thesis, but just insist that moral assumptions are relevant to, and required for, reasonable explanations of CMJs. But while I think that these concerns are worth mentioning, it would seem a mistake to suppose that the appropriate analogue to Harman's scientific observations are IMJs and not CMJs, for Harman's proposed explanation of Kildare's observational judgment is only plausible on the assumption that Kildare makes his observation under conditions that are in general conducive to the avoidance of error (e.g., that Kildare is sober, and so forth), and those are largely the conditions germane to the formations of CMJs. So although Kildare's judgment is an "immediate" one, it is arguably not like just any "initial" moral judgment, which might be formed under less reliable circumstances.

⁵⁸ Harman writes: "Facts about protons can affect what you observe, since a proton passing through the cloud chamber can cause a vapour trail that reflects light to your eye in a way that, given your scientific training and psychological set, leads you to judge that what you see is a proton. But there does not seem to be any way in which the actual rightness or wrongness of a given situation can have any effect on your perceptual apparatus" (1977: 7-8). See also Harman 1986: 62-64.

⁵⁹ It is true that Harman refers to this as an "explanatory chain," as opposed to a causal one, but the relevant explanation is clearly presumed to be causal. Wrongness, he

with the exception of the psychological forces at work in Kildare's brain (which are also causal factors), we have a fairly precise and detailed understanding of those causal connections.⁶⁰ We think we know why and how a free proton would cause a vapour trail under the conditions present in the cloud chamber; and since the vapour trail is itself observable, we understand how Kildare can interact causally with the vapour trail so that he judges as he does given the theories that he holds. It is true that one of the elements of the explanatory chain (the proton) is unobservable, but its existence is confirmed by its explanatory role, and it does not need to be observable in order to assume that role. By contrast, we do not understand how the wrongness of the children's behaviour could be a cause of Spike's judgment. Presumably, on the view in question, wrongness is unobservable (otherwise its presence would not need to be confirmed via an abductive inference),⁶¹ so it could not be a cause of Spike's judgment in the way observable objects (such as vapour trails) are the causes of our observational judgments.⁶² But nor would it seem to be like the proton in being something that can be postulated as a cause of something Spike can observe: Spike sees the children burning a cat, but the wrongness of the action is not, like the proton was for Kildare, plausibly postulated as the cause of this or

claims, breaks the explanatory chain because it can have no effect on one's perceptual apparatus (1977: 8).

⁶⁰ Jobe notes that we are least well equipped to give a detailed account of the psychological and neurophysiological processes that explain why a person judges as he does under certain conditions, and that this deficiency afflicts both moral and nonmoral explanations (Jobe 1990: 269).

⁶¹ The view of Sturgeon and Brink contrasts, in this respect, with that of realists such as McNaughton 1989, and McDowell 1981. The latter writers appear to argue for the direct perception of moral properties while recognizing the normative dimension of perception in general (which makes the recognition of a property depend in part on the background and training of the observer, the conditions in which the observation is made, and so on).

⁶² See Harman 1977: 8; 1986: 63; Jobe 1990: 269-270.

anything else Spike sees. So it seems that there is no account of how the wrongness of the action can be a cause of the judgment, and that there is a good reason to suppose that it cannot be a cause. And if it cannot be a cause of Spike's judgment, it cannot be relevant to a causal explanation of it, as the proton is relevant to the causal explanation of Kildare's judgment. We are left to explain Spike's judgment by appealing to factors that clearly can causally affect him: i.e., his beliefs and the fact that there is a cat being burned in front of him. And in fact such "nonmoral" explanations seem advantaged by the fact that they cite factors that are capable of relating to each other in such a straightforwardly causal way.

4.3: The Counterfactual Argument

Sturgeon notes that it is natural to suppose that whether a particular fact-affirming assumption is relevant to the explanation of a fact depends on whether "the fact would have obtained, and we could have explained it just as well, even if the assumption had been false."⁶³ But he insists that Harman's irrelevance thesis appears plausible only if this test is applied by considering whether the item to be explained would have occurred in a scenario that is similar to the actual one in every natural respect, but in which the moral property affirmed by the assumption is absent. For example, on this account, one would assess the relevance of the fact that Hitler was morally depraved to the explanation of our thinking he was depraved by considering whether the following counterfactual statement is true: "if Hitler were just as he was in all natural respects, but not morally depraved, we would still have thought he was depraved."⁶⁴ Since, as Harman claims, we do not see how the property affirmed by such moral assumptions could have affected anyone's perceptual

⁶³ Sturgeon 1984: 65; see also Brink 1989: 190.

⁶⁴ Sturgeon 1984: 78, n. 30.

apparatus, we should expect assessors to go on judging Hitler as they do regardless of any adjustment in the moral standing of his behaviour.

Sturgeon evidently objects to this version of the test because, as he sees it, how one would judge if Hitler were not morally depraved is itself irrelevant to the question of whether the fact that he was depraved is relevant to the explanation of anyone's judgment that he was morally depraved. Oddly, most of Sturgeon's argument for this crucial claim is hidden in a footnote. Since it is somewhat convoluted, I will paraphrase it instead of quoting it directly:

(1) The principle of moral/natural supervenience, conjoined with the substantive claim that Hitler really was morally depraved, together imply that "there is no possible world in which Hitler has just the personality he in fact did, in just the situation he was in, but is not morally depraved" (69); hence they imply that the antecedent of the aforementioned conditional is "necessarily false" (78, n. 30).

(2) Since the antecedent is necessarily false, the scenario it affirms can be, at best, only "epistemically" possible. That is, it assumes the same modal status as the scenario in which water is not H₂O, given the Kripkean view that water is essentially composed of whatever happens to compose it, and the assumption (which "we have every reason to accept") that water is in fact composed of H₂O (78, n. 30).⁶⁵

(3) Counterfactuals about epistemic possibilities are irrelevant to "questions about what causes what (and hence, about what explains what)"; only counterfactuals about "real" possibilities are relevant to these questions (78, n. 30).⁶⁶

⁶⁵ See Kripke 1980: 125.

⁶⁶ Sturgeon attributes this view to Richard Boyd. But Boyd himself seems to have published no argument for this claim, and in fact sometimes expresses the contrary view. For example, he has suggested that subjecting a theory T to theoretical criticism should involve asking "in the light of the best available theories, what alternatives are there to the mechanisms/processes posited or required by T? What mechanisms, known on the basis of other theories, might interfere with the operation of the mechanisms which T posits or requires? Does the plausibility of T depend upon theoretical considerations which themselves rest on a theory now in dispute? What weakness might that indicate in T? etc." (Boyd 1981: 620). Boyd's view here is apparently that a properly conducted scientific inquiry (into such questions as microphysical causation) requires considering the theoretical alternatives - alternatives that might, in fact, turn out to be false, and (given essentialism) even necessarily false.

∴ (4) "(A)nyone who accepts a causal theory of knowledge (or any theory according to which the justification of our beliefs depends on what explains our holding them)" must find the truth of counterfactual conditional statements whose antecedents affirm such possibilities irrelevant to "the question of how much we know, in science or in morals" (78, n. 30).

Sturgeon apparently intends "questions about 'what causes what'" to include questions about what is relevant to reasonable causal explanations; his suggestion is apparently that epistemic scenarios are in general irrelevant to such questions, or, put more succinctly, that (counterfactuals about) epistemic scenarios are irrelevant to assessments of explanatory relevance. Given this, according to Sturgeon, we must eschew Harman's version of the counterfactual test and assess the relevance of our moral assumptions by considering only what is "really" possible. Instead of testing the relevance of our assumptions by considering scenarios that are ruled out by our substantive judgments, and by our commitment to moral supervenience, we must test them by considering whether the item to be explained would have occurred if, counterfactually, the natural facts upon which the moral fact supervenes were themselves different enough to change the moral fact. For instance, on this account, the test of relevance in the cat-burning example would be:

Even if the children had been doing something else, something just different enough not to be wrong, would you have taken them even so to be doing something wrong?
(1984: 66)

According to Sturgeon, assuming a moral observer holds no bias against children (which would presumably distort his receptivity to the facts), we could not expect him to think the children are wrong under these circumstances. Hence, he concludes, the wrongness of the children's behaviour is relevant to the explanation of the person's judgment (1984: 67).

One possible objection to Sturgeon's case can be dismissed at the outset. It will be noticed that, in contrast to the Kripkean view (on which the basic properties postulated by true scientific theories are real), moral supervenience might be explained antirealistically (as merely a set of constraints imposed on our judgments by certain nonmoral features of the

world and of our psychology).⁶⁷ Given this, it could hardly follow from a claim of moral supervenience simpliciter that Hitler really was, as a matter of necessity, morally depraved, or that scenarios in which Hitler was not morally depraved are just "epistemic" possibilities in the Kripkean sense. However, since Sturgeon is arguably entitled (and even required) to suppose that certain moral facts exist in responding to Harman's irrelevance thesis, it appears that he is also entitled to employ a realistic conception of supervenience in responding to Harman.

The argument nonetheless faces other difficulties. One problem is that it is just not clear what "epistemic" possibilities are: if they are scenarios that are ruled out by whatever (scientific and moral) theories happen to be true, then in what sense are they possible? Sturgeon's answer would apparently be: "in Kripke's sense." But when we look to Kripke we discover at least two mutually incompatible conceptions, and on neither does Harman's scenario qualify as epistemic. Sometimes, for example, Kripke appears to identify the epistemic possibility that "gold might turn out not to be an element" with the possibility of cognizers having inadequate evidence for, or being simply ignorant of, the truth of a microphysical claim:

Gold apparently has the atomic number 79. Is it a necessary or contingent property of gold that it has the atomic number 79? Certainly we could find out that we were mistaken. The whole theory of protons, of atomic numbers, the whole theory of molecular structure and of atomic structure, on which such views are based, could all turn out to be false. Certainly we didn't know it from time immemorial. So in that sense, gold could turn out not to have atomic number 79.⁶⁸

But Harman's scenario is that in which the theory itself is false, not that in which cognizers are ignorant of, or lack evidence for, its truth. Elsewhere, Kripke construes the claim that

⁶⁷ Zimmerman 1984: 89 makes this objection to Sturgeon. I defend an antirealistic construal of moral supervenience in Chapter Six; see also Blackburn 1971; 1984; 1985.

⁶⁸ Kripke 1980: 123; see also 141, 143 n. 72; Putnam 1975: 233.

gold might have turned out to be a compound as the claim that "it is logically possible that there should have been a compound with all the properties originally known to hold of gold" (1980: 142-143). The epistemic possibility that gold might not be Au (which is the base gold is known to have) is here identified as the possibility of there being some other substance (e.g., a compound) that happens to bear all the macroproperties that gold in fact bears. But on this account, too, Harman's scenario apparently fails to qualify as epistemic. For the moral analogue to the gold scenario just mentioned would seem to be the rather strange possibility of there being, for instance, an action other than cat burning that happens to bear all the observable and moral properties of cat burning; i.e., a wrong action that looks (and so forth) like an instance of cat burning, but is not such an instance.⁶⁹ But the possibility Harman would presumably want us to consider is that in which our moral view of cat burning is incorrect, not that in which there is an action other than cat burning that resembles cat burning in its moral and phenomenal qualities.

Kripke exegesis aside, what is the correct view of the matter? There is no question that the truth of certain scientific claims imposes a significant constraint on the conduct of scientific inquiry. In particular, the fact that a scientific theory is true makes it (physically) impossible to conduct any experiment involving a scenario that falsifies the theory. Thus, if water is H₂O, then one could not conduct an actual experiment comparing a scenario in which water is not H₂O to one in which it is, and assess whether inquirers make the same judgment in both.

But the assumption on which Sturgeon's argument depends is not that the truth of a scientific theory makes certain experiments impossible to conduct, but the quite different

⁶⁹ Evan Jobe's scenario in which children burn a cushion in the shape of a cat may seem to represent such a scenario (see Jobe 1990; my discussion in 4.4). But that could not be the parallel scenario since burning a cushion in the shape of a cat is not wrong. A more plausible parallel scenario would, I think, be one in which children burn an animal other than a cat (e.g., a dog that looks like a cat). This action is not an instance of cat burning, but it arguably might share the phenomenal and moral properties of cat burning.

assumption that the truth of a scientific theory makes what we would expect to observe if the theory were false irrelevant - irrelevant, that is, to whether the truth of the theory is relevant to the explanation of our actual observations. And the second assumption is plainly false: we normally assess the relevance of physical laws (which, of course, hold with physical necessity) using the same procedure we employ for claims that are not physically necessary, viz., by considering whether the fact to be explained would have obtained if the claim in question were false. Suppose someone tried to explain the behaviour of a falling body B under certain initial conditions $\{C_1 \dots C_n\}$ by citing a physical law L, and that the question arose whether L is really relevant to the explanation of how B behaves under those conditions. Applying Sturgeon's reasoning to this example, if L is true, we could presumably not assess its relevance by considering how B would have behaved if, counterfactually, L were false. For if, as we are assuming, L is a physical law, then it will presumably be like the identification "water = H₂O" in holding as a matter of scientific necessity. But this hardly seems to be the correct approach: were someone to deny that L is relevant to an explanation of B's behaviour, a more reasonable reply would seem to be that L is relevant because, did L not obtain, B would have behaved differently. Indeed, any counterfactual assumption that is permitted by our scientific theories would appear to be entirely irrelevant to the assessment of the relevance of L. For instance, it may be scientifically possible to conduct a variant of the experiment by changing some or all of the initial conditions. But how B would behave under different initial conditions seems incapable of indicating anything about the relevance of L to the explanation of B's behaviour under $\{C_1 \dots C_n\}$; at most, it seems to suggest something about the relevance of the original conditions to the explanation of B's behaviour. The same point can be made in the cat-burning case: our moral theories might allow there to be an action that is different enough from cat burning not to be wrong. But varying the natural and social properties of

the action in this way will tell us whether the natural and social properties can be relevant to a reasonable explanation of Spike's moral judgment and not whether any moral property is relevant.

This is, of course, just an account of how things seem. But I believe these appearances can be explained in a way that vindicates the account. As I see it, the reason that the theoretically "safe" counterfactual scenarios mentioned above seem irrelevant to the assessments in question is that they make for skewed tests of relevance. In science, the notion of test skewing is ordinarily applied to observational tests of causal hypotheses: determining whether something A causes (and explains) something B is partly a matter of constructing an experiment in which A is the only possible cause of B; i.e., one in which other candidate explanatory hypotheses can reasonably be ruled out. If they cannot be ruled out, an experiment is skewed. For instance, a physician might attempt to diagnose whether her patient has a strep infection on the basis of an in vitro culture test. Normally, a positive in vitro test is evidence of (because it can be explained by) a strep infection. But this will not be so if the test result could plausibly be explained in another way - such as on the assumption that the petrie dish was infected by a technician's dirty hands. In that case the test is botched. Now, assessments of explanatory relevance, it seems, can be skewed in much the same way. Determining whether A is relevant to the explanation of B is partly a matter of considering whether B would have obtained if A did not. If B would not obtain if A did not, then we may regard A as relevant to the explanation of B. But this conclusion does not follow if there is an independently plausible explanation of why the scenarios differ with respect to B. The availability of another explanation would skew the test.⁷⁰

⁷⁰ Suppose for instance that epidemiologists discovered that vegetarians get fewer cancers than omnivores and carnivores. The discovery would no doubt suggest that the fact that the vegetarians eschewed meat is relevant to the explanation of the difference in cancer rates; for it would suggest that if the vegetarians had eaten meat, their rate would have been much higher. But now suppose that it is also discovered that vegetarians are also much less inclined to engage in other risky behaviours such as smoking and drinking

Now, Harman's way of testing the relevance of moral properties, whatever else might be wrong with it, is clearly not skewed in this way. Since, ex hypothesi, the only difference in the two contrasted scenarios, as Harman views them, is that wrongness is present in one but not the other, any difference in the way Spike judges could be explained on the assumption that the moral property is affecting Spike (even if it is not understood how this occurs⁷¹); other explanatory hypotheses are ruled out. By contrast, conducting the test Sturgeon's way, we have no reason to exclude the explanatory hypothesis that attributes the difference in judgment to differences in natural properties.

Finally, even if it is true that the scenarios that are ruled out by whatever moral and scientific theories happen to be true are irrelevant to questions about explanatory relevance, Sturgeon's modal and substantive assumptions hardly imply that he, or anyone else, is in a position to discern which scenarios these are. For even if one can assume, for the sake of an argument, that Hitler was morally depraved, and that he was so as a matter of necessity, these assumptions alone do not imply that such moral facts are epistemically accessible to cognizers - even supervening moral facts might simply be causally isolated epiphenomena. Given this, the assumptions provide no reason to suppose that anyone is ever in a position

alcohol. We might then plausibly suppose that if those who are vegetarians had adopted another diet, they would also have been much more inclined than they are to smoke and drink. If so, then while we could still suppose that the cancer rate among vegetarians would have been higher if they had eaten meat, we could not conclude that their vegetarian diet is relevant to the explanation of their incidence of cancer. The assessment, we may say, is skewed by the availability of a plausible alternative explanation of the different rates. To adequately assess the relevance of vegetarianism, we would have to consider scenarios in which different rates could not be explained by anything but the difference in diet.

⁷¹ The point here is that if, counterfactually, we did discover a difference in judgment, there would be reason to suspect the moral properties as a cause. Notice that this is not at odds with the claim that, since we don't understand how the moral property could affect Spike, we have no reason to expect a difference in moral judgment in the contrasted scenarios.

to dismiss the scenarios employed in Harman's version of the test because they are "epistemic."

Of course, Sturgeon also thinks "we have good reason" (78, n. 30) to suppose that Hitler was morally depraved, and this conjoined with his other assumptions together imply that he has good reason to dismiss Harman's test. But it is difficult to see how he can make this further assumption without ultimately either begging the question or making his argument otiose. For he clearly owes us an explanation of how good reasons can be obtained, and it seems that the answer must be either that they are gained through the explanatory inferences in question, or in some other way. Obviously he cannot simply assume that his reasons are furnished by inferences to the best causal explanation without assuming a position on the issue at hand. But nor can he suppose that his reasons are supplied in some other manner without making his account depend on our having solved in advance the problem of moral justification.

Perhaps Sturgeon thinks he can suppose that he has good reasons to believe that Hitler really was depraved without making any assumptions about the causal etiology of his judgment just because (as he points out) he can legitimately assume the approximate truth of what strikes him as being "the best overall theory of what we are like and of what the world is like," and because that theory includes convictions which themselves support his judgment about Hitler. The rationale for this move is supplied partly by his commitment to a naturalistic epistemology. He writes:

. . . we have in general no a priori way of knowing which strategies for forming and refining our beliefs are likely to take us closer to the truth. The only way we have of proceeding is to assume the approximate truth of what seems to us to be the best overall theory of what we are like and what the world is like, and to decide in the light of that what strategies of research and reasoning are likely to be reliable in producing a more nearly true overall theory. (1984: 77)

Since included among the commitments that are germane to the overall theory that Sturgeon accepts are some moral convictions, Sturgeon could apparently have a reason for

believing that Hitler was morally depraved quite apart from any assumption he may hold about the causal etiology of that particular belief. For suppose that part of the overall theory to which Sturgeon may legitimately be committed is some moral theory M. M might, it seems, be a theory that supports his judgment about Hitler, while itself being supported by considered moral judgments about particular cases not including that judgment. Of course, in order for those other considered judgments to do their justificatory work, at least some of them would themselves presumably have to be regarded as being accepted as a result of input from real moral facts, but it seems that this supposition might be warranted by certain second-order convictions about the reliability of considered moral judgments - convictions that also comprise the overall theory to which Sturgeon may legitimately appeal.

But surely this is a misapplication of the naturalistic approach to epistemology. That approach is marked by a refusal to attempt to satisfy the traditional demand for an a priori justification of empirical beliefs; instead, the naturalist insists on construing the problem of knowledge as a problem for the empirical disciplines.⁷² Naturalism thus "makes free use of the findings of neurophysiology, cognitive psychology, evolutionary biology,"⁷³ and of those disciplines that study the environment in which we operate. It is these findings that largely make up our overall theory - an overall theory that describes cognizers and the world in a way which explains how knowledge of various sorts is gained. Of course, the overall theory we in fact hold includes not just sophisticated scientific beliefs and theories, but also many commonsense beliefs about the world and about what we are like. For our commonsense view of things is often (though not always) vindicated and extended by scientific inquiry. But although the naturalist's approach allows us to assume the truth of all of the commonsense and scientific beliefs germane to our

⁷² See Maffie 1990: 282; Quine 1960: 275, 1969.

⁷³ Maffie 1990: 282.

overall antecedently-held theory of the world and of human cognizers, it is unclear why anyone taking this approach would include moral beliefs in their overall theory. For while the aforementioned empirical disciplines explain how cognizers interact with their environment in a way that makes perceptual knowledge possible, we have no comparable theory of how cognizers could apprehend real moral facts, such as the fact about Hitler. Clearly our moral theories, which just represent various moral views, do not seem to comprise (even in part) any such explanatory account; so it is difficult to see how a naturalist could make use of them.

All of this suggests that the appeal to a naturalistic epistemology does not help provide Sturgeon with a reason to suppose that Hitler really was morally depraved. And if it does not do this, it fails to provide him with a reason to regard any scenario in which Hitler was not morally depraved as "epistemic" and thus as irrelevant to any assessment of explanatory relevance.

4.4: Jobe's Application of the Test

The above arguments to the contrary notwithstanding, it is worth mentioning that even if applying the test Harman's way involves making plainly absurd or otherwise problematic assumptions, there may nonetheless be an alternative to Sturgeon's version of the test. Jobe (1990: 269-70) suggests that explanatory relevance can be assessed by comparing scenarios in which the actions, although extremely different in their natural (and hence moral) properties, appear exactly similar (e.g., cat burning, and burning a cushion in the shape of a cat). This test isn't skewed because, although the natural properties are adjusted along with the moral properties, the natural properties do not causally affect the observer differently, and hence could not conceivably explain a difference in judgment. A difference could then be explained by appealing to the difference in moral status. But, in fact, we don't expect there to be a difference in moral judgment. We expect Spike to judge

as he does whether the children are burning a live cat, or just a similarly appearing cat facsimile (such as a cat cushion or, perhaps more plausibly, a cat robot) and thus we expect moral properties to fail this test.

It may be objected that the case of the robotic cat is similar in a crucial respect to the example Sturgeon gives of the person who hates children. If Spike hates children, we could not plausibly explain his judgments about children as perspicacious moral "observations." But this does not make all moral explanations suspect any more than the possibility of an accidentally-infected petrie dish revealed a general problem germane to all in vitro culture tests; it is just that the person with a strong bias against children will presumably be inclined to condemn them no matter how they behave (just as conducting an in vitro test with an accidentally infected dish will produce a positive result whether or not the patient has a strep infection). Similarly, in the robotic cat case, it may be claimed that we have departed from the ideal causal scenario in perception (we are faced, perhaps, with something like an optical illusion). So, the realist might say, there really has been a change in moral properties (burning a cat robot is not wrong), but we have an explanation of why the ordinary moral observer does not see this that is both independently plausible and compatible with moral realism: i.e., the cat robot looks exactly like a cat! However, the proposed explanation seems more amenable to antirealism than to realism. A moral antirealist who thinks that moral sensibilities and (perceptual input from) nonmoral facts are the only factors relevant to explanations of moral judgments can perfectly well accept the explanation: he can say that there is no difference in judgment across these scenarios simply because in both cases the sensibility of the assessor and the perceptual input are the same. By contrast, the realist, who presumably thinks moral properties are an additional explanatorily relevant factor,⁷⁴ seems incapable of explaining why the cognizer does not

⁷⁴ Few if any moral realists will, of course, want to claim that moral properties are the only factors relevant to the explanation of moral judgments.

perceive the moral difference notwithstanding the sameness in nonmoral input.⁷⁵ If moral properties really are relevant, it seems that we should expect a difference in judgment across scenarios that are morally different even though the other explanatorily relevant factors remain the same. In this respect, the cat robot case differs from the case of the child hater: the fact that a person hates children accounts, in a way that is compatible with realism, for his apparent insensitivity to moral differences. But the fact that the cat robot looks exactly like a cat does not similarly help to explain how we expect Spike to judge.

⁷⁵ In general, if candidate explanatory factors A, B and C are relevant to the explanation of E (and other factors can reasonably be ruled out), it will be unreasonable to expect E to occur in only one of two scenarios that otherwise differ only with respect to C. Obtaining E in both scenarios would suggest that C is in fact irrelevant to the explanation of E.

CHAPTER FIVE

EXPLANATORY NECESSITY

5.1: Introduction

Chapter Four examined the AMR case for the thesis that moral facts are relevant to explanations of certain nonmoral facts. It will be noticed that the relevance thesis is rather weak: in particular, it provides no reason to suppose that any inquiry-independent moral facts actually exist. It says only that inquiry-independent moral facts, if they existed, would be relevant to explanations of nonmoral facts. But Brink and Sturgeon also make the much stronger claim that assumptions about moral properties are often required to explain certain nonmoral facts. If correct, this suggests that moral property ascriptions might be justified by their essential explanatory role.

AMR theorists present two main arguments for the explanatory necessity claim: these I shall call the "Generality Argument" and the "Argument from Moral Conversions." Section 5.2 examines the Generality Argument; 5.3 and 5.4 explicate and assess the Argument from Moral Conversions.

5.2: The Generality Argument

The Generality Argument, in contrast to the counterfactual argument, is an argument for the indispensability or necessity of explanations citing moral facts or properties. Briefly, the claim is that nonmoral explanations are frequently "too specific," given a constitution/nonidentity account of moral properties.⁷⁶ Recall that on a

⁷⁶ The claim, and the argument for it, are attributed to Sturgeon by Zimmerman (1984: 86), who cites Sturgeon's "correspondence with Harman" as the source. See also Brink 1989: 194-197.

constitution/non-identity account, moral properties are multiply realizable. According to Sturgeon and Brink, it is typically irrelevant which natural property in fact produced a given moral property, for typically the explanandum item would have occurred no matter how the moral property is realized. But since (a) it is often reasonable to suppose that the explanandum item would not have occurred if none of the possible natural bases underlying a moral property obtain, and (b) we may plausibly regard many moral properties as being realizable via bases that share few or no common natural properties, it seems that moral properties are sometimes not only relevant to, but required for, certain explanations. In short, often only moral facts occur at the level of generality required to account for the explanandum. Thus, to take the example Brink gives, injustice in South Africa is realized by the particular political and economic arrangements that presently exist there.⁷⁷ Yet in explaining the anti-government protests, it is better to cite the moral fact of injustice than to cite the particular political and economic policies, because, first, the protests would have occurred even if the government policies had differed in certain ways, and second, the only thing these actual and counterfactual policies have in common is that they realize injustice.

Brink regards moral explanations as a species of explanation which cites "structural" rather than "compositional" facts. Compositional facts are, it seems, the

⁷⁷ Though it is clear that Brink regards the injustice of South African society as explanatorily relevant, and often appeals to that example in arguing his case (see 1989: 190-192), the actual example Brink uses at 1989: 195 is that of South Africa's "racial oppression" rather than its "injustice." But it is not clear that the former example supports Brink's argument. For it is not clear that "racial oppression" is a moral, as opposed to a natural (i.e., social) fact. One can easily define it in a morally neutral way: say, as a policy (or set of institutions, political and economic arrangements, etc.) which always gives exclusive consideration to the interests of racial group W at the expense of other groups. Such facts are normally considered pertinent to moral assessment, but they are not necessarily "moral" facts (both proponents and opponents of apartheid might acknowledge that this is a policy of the South African government, without agreeing about whether there is anything morally amiss about it).

specific possible underlying bases of higher-level constituted properties; structural facts are constituted facts ranging over multifarious possible bases. According to Brink, the claim that some causal explanations must cite (structural) moral facts rather than (compositional) natural bases is plausible because in general "causation itself is often structural and not (simply) compositional" (1989: 195). He presents a tidy formal argument for the latter, more general, claim:

Assume that we begin by considering some world w^1 in which we explain some fact or event E by appeal to some higher-order fact or event H , and that we (correctly) think that H caused E (since we think this sort of explanation is causal). (H and E can be either types or tokens.) Now, H is constituted by, but not identical with, some set of lower-order facts or events L^1 ; this is at least in part because, although actually constituted by L^1 , H could have been realized (in other worlds) by any of a very large number of somewhat different configurations of lower-order facts $L^2, L^3, \dots L^n$. Now we might also think that we could explain E by appeal to L^1 . If this is right, we might ask if H is explanatorily eliminable in favour of L^1 . I think not. For in other worlds, otherwise as similar to w^1 as possible, in which H is constituted by one of these alternative L bases, H and E will also occur, and it will also be correct to say that H causes E and so explains E . This implies that H 's causal power is neither equivalent to, nor derived from, L^1 's causal power (contrast Kim 1984). It is also true that E would have occurred when and as it did even if L^1 had not occurred. This counterfactual test may encourage us to claim not just that H 's causal power cannot be reduced to that of L^1 , but that H and not L^1 is the cause of E - that the cause of E is the structural fact H , not the compositional fact L^1 . (Brink 1989: 195-196)

For instance, according to Brink, in order to explain China's adoption of the one-child policy (E), one would have to cite the structural fact that the rate of population growth prior to implementation of the policy was constraining capital investment and modernization (H), even though H consisted partly in a specific birthrate, which in turn consisted in actual births of a particular set of individual Wangs, Chows, and so forth (L^1). Why cite H and not L^1 ? Because H might very well have been realized via some other way (e.g., via another set of births), and because we would expect to have had both E and H regardless of which base realizes H (Brink 1989: 196).⁷⁸

⁷⁸ It is worth noting that Brink sometimes makes remarks that seem at odds with his generality argument. For example, he writes: "Because the natural facts, which Harman concedes help explain South African instability and protest against apartheid, are

However, even if this is plausible as an account of causality, the argument does not establish the indispensability of moral facts. There are two difficulties. First, there is what I shall call the "H but not E" problem. We have noticed that a general structural fact H will be necessary to explain something E only when E will obtain only if, but no matter how, H is realized, but the supposition that both E and H would have occurred even if H had been realized in another way seems implausible when H is a moral property and E is a nonmoral fact. This poses a problem because, in general, when not all the possible bases $\{L_1 \dots L_n\}$ underlying a general property H can be expected to give rise to an explanandum item E, it would seem positively unreasonable (and certainly unnecessary!) to appeal to H in order to explain E. Suppose, for instance, that a brush fire sweeps through an opulent Los Angeles hillside suburb and all of the wooden houses burn to the ground, while the brick houses remain standing. While an explanation of why only the wooden houses burned down should obviously cite the difference in composition among the houses (the houses that burned down were made of a more combustible material than the houses that did not) it would not, strictly speaking, be correct to say that the wooden houses burned down "because they were made of wood." Rather, since not every kind of wood burns under such conditions (redwood, for example, does not), the correct explanation must be a more specific one (by appealing to the type of wood involved).

Now, explanations that cite moral structural properties appear in this respect to be like the wood explanation, but unlike the proposed explanation of the one-child policy. Although we may reasonably suppose that China would have adopted the one-child policy in response to its relatively high birthrate no matter what particular geographical/temporal constellation of Wangs, Chows, and so forth, happened to be born during that period, there

what constitute the racial injustice of apartheid, their explanatory power ensures the explanatory power of the injustice that they realize" (Brink 1989: 191; emphasis added).

seems to be no general moral property that one could reasonably expect moral cognizers to recognize no matter how it is realized. Suppose that E (the explanandum item) is Spike's judgment that the children burning the cat are acting wrongly, that H is the constituted and hence multiply realizable property wrongness, and that L^1 is the act of cat burning (plus whatever else contributes to the production of the act's wrongness in w^1). Since H is constituted but nonidentical to L^1 , it might have been realized by base L^2 , and so on. But given the wide range of ways in which "wrongness" might be realized, and the intractability of many moral situations, it seems that one could not reasonably expect Spike to recognize every, or even most, alternative instances of wrongness. Suppose, for example, that the children in Harman's example are not burning a cat but instead are eating hamburgers, and that eating meat is (for various reasons) wrong. In that case, and probably many others, Spike might very well not think that the children are engaged in wrongdoing. In other words, given many of the alternative bases, E might very well not have occurred. But if so, one could not reasonably cite the general fact H - which ex hypothesi ranges over all of these bases - as the cause of Spike's judgment.⁷⁹

⁷⁹ It will be noticed that Brink restricts his discussion to worlds that are, apart from the fact that H is constituted by a base other than L^1 , "as similar to w^1 as possible" (call this restriction R). Brink imposes R, one may suppose, in order to contend with worlds in which H is realized by, say, L^2 , and in which E does not occur, but where this seems irrelevant to the question of whether we need to postulate H in order to explain E in w^1 . For instance, there may be worlds in which there is H and L^2 , but in which the physical laws have changed so that L^2 no longer causes E, or in which a cognizer fails to recognize L^2 as H because she is drunk (and thus fails to meet conditions of general cognitive reliability), or does not possess a requisite degree of moral training to be able to recognize that L^2 is H. Such possibilities are clearly irrelevant to the question of whether it is necessary to ascribe H in order to explain E in w^1 : for instance, in worlds where the physical laws have changed so that H (as composed by L^2) no longer causes E, E's absence may be attributed to the change in laws, not to the change in underlying bases, and it is simply irrelevant to the issue at hand that E would not have occurred for that sort of reason. It must be stressed that my claim that E would not obtain given certain realizations of H does not require violating restriction R. My claim is vindicated if one could plausibly expect that for some world w^n (a world similar to w^1 in every respect save for the fact that a moral property H is realized by a base other than L^1), E would not occur. This claim is

The "H but not E" difficulty also afflicts the other examples Brink and Sturgeon mention. Why did opposition to slavery occur in the United States and in Britain long before it developed in the rest of the New World, even though the practice was widespread in all of these places? According to Sturgeon, because it was morally worse in the United States than it was elsewhere.⁸⁰ Applying Brink's reasoning here, one would say that the appeal to its worseness is required because it could have been worse in many different ways - ways that are so different from each other in their natural and social properties that the only thing they have in common is the fact that they are worse - and because we may reasonably expect to see such opposition to badness (of this degree) no matter how the badness had been realized. But is the last supposition really tenable? I think not. For the general property of "being worse" or "bad relative to something else" is so general that it need not even involve slavery, or indeed even any sort of violence or political oppression. Perhaps as "bad" as a slave state is one that is ineffective and anarchic, or that is (grossly and unnecessarily) ecologically unsustainable, or that although democratic and egalitarian in

quite plausible. Given that cat burning and meat eating are possible bases (respectively L^1 and L^2) of wrongness (H), that w^1 is where Spike (who is not drunk or psychopathic) sees the children burning the cat and thinks "those children are engaged in wrongdoing" (E), and that w^2 is exactly similar to w^1 (the same physical laws obtain, here also Spike is a sober and morally perspicacious observer, etc.) except that Spike sees the children eating meat, Spike in w^2 may very well not make the moral judgment he made in w^1 . We certainly do not need to transport Spike to some w^3 where the relevant physical laws have changed, or where Spike is drunk or has had a bad moral education, or where his moral views are different from those he has in w^1 in order to make this supposition.

⁸⁰ Sturgeon apparently thinks that explanations that cite injustice are common in historical writing, but gives no examples. Perhaps he has in mind the genre that Railton (1986) mentions, and that is exemplified in Barrington Moore's Injustice: the Social Bases of Obedience and Revolt (1978). Moore's tactic, however, seems compatible with moral antirealism: according to Moore, the source of social unrest lies with the frustrated (but apparently universal) desire for "reciprocity" (which is the idea that citizens ought to receive some benefit [e.g., protection] in return for recognizing authority, etc.). Injustice is just the lack of reciprocity. But we can perfectly well explain uprisings by citing the frustrated desire without also assuming that there is anything morally amiss about having one's desire for reciprocity frustrated.

its internal politics is nonetheless aggressively imperialistic. But I see no reason to suppose that, *ceteris paribus*, all of these forms of "badness" should prompt the same political response. The same seems true of the South African "injustice" case: South Africa might have been "unjust" in a number of very different ways. It might have instituted a system of sexual or religious, rather than racial, apartheid, it might have been a communist dictatorship rather than a limited democracy, it might have imposed radical laissez-faire and social Darwinist economic policies and ignored any resulting class disparities, or it might have been a pleasure tyranny of the sort described by Aldous Huxley (1946).⁸¹ It seems implausible to suppose that the same sort of political response would have arisen no matter which of these injustices was actually realized. Thus, while it would appear irrelevant to a reasonable explanation of the adoption of the one-child policy that the policy was preceded by a particular constellation of births, it would not appear irrelevant to a reasonable explanation of the actual political opposition that arose to the actual policies of the South African government, that injustice in South Africa was actually realized in the way that it was.

An AMR theorist might respond to this objection by making the moral property in question more specific; perhaps, she might suggest, the moral property to which we must

⁸¹ One might wonder whether worlds containing pleasure tyrannies pass restriction R (see note #3 above), since they involve agents whose cognitive and perceptual abilities have been compromised. The citizens of Huxley's *Brave New World* don't resist because they have been drugged into submission, and it may be thought that this makes worlds containing pleasure tyrannies like worlds in which Spike has been drugged and in which he fails to recognize the wrongness of meat eating because he is drugged. But a world which is exactly similar to the actual one except that injustice is realized via a pleasure tyranny passes R; the world in which Spike fails to see the wrongness of eating meat because he is drugged does not. In the former world, the fact that the citizens are drugged into submission is included under H; it is one of the ways in which H might be realized; but the world in which an intoxicated Spike fails to recognize the wrongness of eating meat is not a world which is similar to the cat-burning scenario in every respect except for a difference in the way the moral property is realized. The fact that Spike is drugged is not a possible basis of the wrongness he encounters on the streetcorner.

appeal is just badness as it applies to the practice of slavery, or injustice as it applies to racial apartheid or segregation, or wrongness as it applies to cat burning.⁸² Such a move is not implausible. One might contend that slavery and apartheid have a unique moral standing, and that the act of burning a cat is a special kind of wrongdoing. And it may be claimed that the corresponding property-ascribing moral judgments are best understood as ascriptions of these more specific moral properties - for instance, that Spike's judgment about the children's behaviour is best construed as the judgment that "the children are wrong to burn the cat" (as opposed to just "the children are acting wrongly"). Making the relevant moral facts and judgments more specific in these ways seems to make it much less probable that in such cases we could have the moral fact without the associated moral judgment. For instance, let "narrow E" be Spike's judgment that "the children are acting wrongly in burning the cat," "narrow H" be the candidate explanatory moral fact that the children do wrong in that particular way; L¹ be the scenario in which Franky and Barney set Fluffy on fire and L² be that in which Jennie and June set Whiskers on fire. (Presumably, while narrow H essentially involves cat burning, and so forth, it is nonetheless multiply realizable via L bases involving different individuals.) We may reasonably expect Spike to make his judgment regardless of who does the burning and which cat is burned, so the construal seems to avoid the "H but not E" problem. And it seems to provide for an effective response to the "hamburger-eating" counterexample mentioned above. There is no question that narrow E would probably not have occurred if the children were eating hamburgers (L³) instead of burning a cat, for (obviously) Spike would not say "those children are acting wrongly in burning the cat" upon seeing children who are eating hamburgers or doing some other arguably wrong apart from cat burning. But this, a realist could say, is not a case in which there is narrow H but no narrow E;

⁸² This is apparently how Zimmerman (1984) interprets Sturgeon's claim.

rather, it is just a case in which both narrow E and narrow H are absent.. So the hamburger counterexample fails to show that the wrongness of cat burning is not indispensable to Spike's judgment that the children are wrong to burn the cat.

Parallel claims seem tenable regarding the slavery and South African injustice cases. It may be claimed that slavery in the United States might have been morally worse than elsewhere in many (actually unrealized) ways (e.g., different, but equally cruel, forms of torture might have been used to punish resisters), but one would expect the opposition to have arisen no matter which of these (equally worse) forms of slavery actually obtained. Similarly, one might claim that apartheid-style injustice in South Africa might have assumed a number of different forms, but one could reasonably expect unrest in South Africa to have arisen no matter which of these forms it had taken.

However, there are two problems with this approach. First, the claim that we should expect, on the revised accounts, to have the explanandum item no matter how the candidate explanatory fact might have been realized, seems plainly false in at least the slavery and apartheid-style injustice cases. On a not implausible moral view of the matter, it just seems false that the opposition to slavery in the United States would have arisen no matter how the practice there had been worse than elsewhere. There is arguably something very bad about enslaving someone via a means of control that serves to secure the compliance of the enslaved so that no force is needed; indeed this form of slavery seems worse in some respects (perhaps the most important respects) than the violent kind, for it corrupts those it enslaves more profoundly than does slavery of the former type, and is more easily perpetuated.⁸³ But there is no obvious reason to expect political opposition to have arisen sooner in the U.S. if the slavery practised there had been worse than elsewhere in this respect. In fact, one might tenably expect just the opposite. So the simple appeal to

⁸³ See again Huxley 1946.

moral worseness appears not to explain the event in question. The same point can be made about the "injustice of apartheid." Apartheid segregation might have been unjust in many other ways (depending, for instance, on how it is imposed, on whether a system of tribal homelands is employed, and so on), but it is just implausible to suppose that it would have provoked the sort of opposition (inside and outside South Africa) which the actual policy has no matter how it had been unjust. Suppose, for instance, that the South African government had relied much less on brute force and much more on persuasive propaganda, generous financial rewards for individual blacks willing to defend the social order and to drop their demands for the franchise, and so forth. In that case we would arguably have seen fewer demonstrations, fewer violent reactions by the State, and so on, but it is also arguable that a South Africa of this sort would be no less unjust than it actually was under apartheid.

Second, in all of the cases considered here, once the narrow construal is adopted the need for a "moral" explanation arguably evaporates. For recall that, on the generality view, moral facts must be postulated in making the proposed explanations because only they occur at an appropriate level of generality. But on the revised account of each of the cases, there would seem to be a nonmoral fact that is general enough to account for the item we want to explain. In the revised account of the cat burning case, it is the fact that children are burning a cat; in the slavery case, it is the fact that a form of slavery has been adopted; and in the apartheid case, it is the fact that a form of apartheid has been instituted.

All of this suggests that the AMR theorist cannot evade the "H but not E" problem by opting for a more specific construal of the Hs and Es in question. But even if the "H but not E" difficulty is dismissed, there is a more serious problem. Even if there are some E facts whose explanation requires an appeal to certain general H facts rather than to specific L bases, we need not, in pursuing these explanations, construe the H facts

realistically. Suppose Spike tastes chocolate ice cream and says "This is good," but that there is no perceivable feature common to all of the foods which Spike thinks are good. He says oranges and chocolate are good, but not orange and chocolate ice cream, he says curry is good, but only curried legumes; and he says legumes are good only when they appear in a curry, etc. How to explain Spike's judgment about chocolate ice cream? Perhaps some concern for generality is appropriate here. A strict particularist approach - one that refused to cite any feature common yet peculiar to all the foods Spike likes - would seem incapable of explaining what it is about any particular item that makes it agreeable to Spike; one could not, for instance, explain his judgment about chocolate ice cream by simply citing the fact that chocolate ice cream contains chocolate; one would have to cite some fact about Spike, i.e., about his inclinations to like certain foods, and so forth. But this suggests that one can be a generalist without being a realist; i.e., without postulating a real and causally efficacious structural property common to all the foods Spike likes. For even if there is no basic perceivable feature common and peculiar to all the foods Spike thinks are good, they still have something in common: viz., they all belong to the set of foods that Spike likes, given his personal experiences with various foods, and his culturally inculcated attitudes and tastes. One might employ Harman's terminology and include this complex of attitudes and tastes in Spike's "psychological set." A general but irrealist explanation of Spike's judgment about chocolate ice cream would then cite as its cause the fact that the item is highly rated by his psychological set. This, suitably adumbrated by an account of how the affection in question came to be included in that set, would make a tenable explanation of the judgment. Now, it is easy to see how this approach might be tailored to an antirealist account of Spike's judgment about the children burning a cat. The antirealist can explain that judgment on the assumption that cat burning bears a general property H (general in that it is shared by everything else Spike is inclined to condemn

morally), but insist that H is not a real moral property, but just the property of being rated in a certain way by Spike's psychological set. Thus construed, the antirealist explanation satisfies the demand that a general property be cited without postulating a real moral property.

If my case to this point is sound, I have established that there are certain cases in which we should not expect to have the item we want to explain given any possible realization of a putative moral fact, and that even if a generalist explanation is required in such cases, one can reasonably offer an antirealist one. An AMR theorist may still object that none of this is, strictly speaking, incompatible with the central point of the generality argument, which is presumably just that moral structural facts must be cited for at least some explanatory purposes. After all, Brink claims only that causation is often structural (1989: 193), and that many higher-order explanations are irreducible to lower-order ones (1989: 193). Given this, it may be contended that I have established at most that Brink's own paradigmatic illustrations of his thesis fail to support his view. This being so, there is nothing in what I have said thus far that could prevent an AMR theorist from citing other, more promising, examples.

Nonetheless, the AMR theorist faces two general difficulties which, I believe, show that citing moral structural facts is never explanatorily necessary. First, assuming that the antirealist/generalist explanation of Spike's judgment in the cat burning case is indeed viable, there is no obvious reason not to apply that scheme to every case where a moral explanation seems attractive.⁸⁴ Second, it is not clear at all how Brink's argument could rebut Harman's moral epiphenomenalism charge. For it provides no insight into how

⁸⁴ This should include explanations of historical events too. Explanations of historical events that cite putative moral facts are explanations that depend on the idea that historical figures have perceived moral facts and then acted on the basis of those perceptions. But if we can offer a reasonable antirealist generalist explanation of moral judgments, it would seem that we can avoid having to cite inquiry-independent moral facts in making these explanations.

constituted moral properties - construed realistically - could affect cognizers so as to induce them to make certain judgments and behave in certain ways. It is easy to understand how the fact that China's birthrate inhibited capital investment could lead its politicians to adopt the one-child policy: regardless of which set of Wangs, Chows, and so on, had actually lived in China, in sufficient numbers the result would have been an overpopulated and undercapitalized society. But what do the myriad forms of wrongdoing that Spike might witness on the street corner all have in common that could cause and enter into the explanation of his judgment, and which is only capturable by a moral ascription? The obvious answer is ruled out by Brink ex hypothesi: i.e., that there is some common (natural) element in virtue of which acts are instances of wrongdoing. On the proposed account, what makes moral explanations necessary is the fact that in many cases there are no common basic properties. The only thing that instances of wrongdoing have in common is their wrongness. But how could this constituted general feature come to affect Spike? And why should this generalist explanation be better than the antirealist one mentioned above?

It is noteworthy that the capacity of moral facts to cause certain historical events, such as the fact that opposition to slavery arose where it did during the eighteenth century, seems just as mysterious. Notice that the present view is not just that the basic (compositional) features that ostensibly give an institution, policy, or practice its peculiar moral standing can sometimes cause certain historical events (no one, not even a moral epiphenomenalist, will deny this much). It is rather that historical events are often only explained on assumptions about structural moral facts, such as the fact that certain features of nineteenth century U.S. slavery made the practice morally worse than it was elsewhere. Yet, it seems we have no clearer understanding of how a fact of this sort could cause there to be relatively more unrest in the U.S., than we do of how the constituted general moral

property "wrongness" could affect Spike. This is perhaps unsurprising, for it is natural to suppose that if structural moral facts do figure as causes of historical events, they would do so precisely in virtue of their capacity to cause people to make certain moral judgments and then to act on the basis of those judgments (e.g., it is natural to suppose that if the moral "worseness" of slavery was a cause of the aforementioned pattern of unrest, it was so just because people perceived how bad American slavery was in comparison to the way it was in other places, and acted on that perception). So the aforementioned difficulty with respect to moral judgments would seem to infect moral structural explanations of historical events as well.

An AMR theorist might respond that I have simply failed to appreciate an important feature of the theory of structural causality, viz., that on the structural view, there is no mystery about how structural facts figure as causes largely because the theory allows us to account for structural efficacy in terms of compositional efficacy. For instance, there is no mystery about how the structural fact in the one-child policy example could figure as a cause, largely because we can understand how any of the possible constellations of births that might have underlied that fact could have induced the Chinese Government to adopt the policy. The theory recognizes a simple point: that it is often irrelevant to a reasonable explanation of something which particular base realizes a structural fact, but nonetheless not irrelevant to such an explanation that one of the possible bases of a structural fact obtains.⁸⁵ So it can account for the causal power of structural moral facts in terms of the causal power of their underlying bases even while insisting that making certain explanations sometimes requires postulating more general structural facts as causes. For

⁸⁵ Suppose billiard ball A hits ball B and sends B into the corner pocket. In a significant sense ball A (that particular ball) caused B to behave as it did. But in another sense it is largely irrelevant (to an explanation of B's behaviour) that A hit B, for presumably B would have behaved that way even if another ball had hit it in the way A did; what matters (for explanatory purposes) is that a ball of a certain mass, shape, and so forth, struck B from a certain direction and with a certain velocity.

instance, it can say that although a specific constellation of births actually precipitated the one-child policy, explaining that policy demands that we postulate the more general structural fact as a cause. Accordingly, an AMR theorist might say, even though we must analyze the efficacy of moral structural facts in terms of the efficacy of an underlying base, we must nonetheless postulate certain moral structural facts as structural causes in order to make certain explanations. And she can insist that this stance enables us to avoid the aforementioned mystery precisely because it allows us to analyze the efficacy of moral structural facts in terms of the efficacy of the compositional bases that realize them.

However, I think the objection itself obscures a feature of the structural account that, once appreciated, should prevent us from applying this method of remedying the aforementioned mystery to moral cases. The neglected feature in question is this: while ascribing causal/explanatory efficacy to structural facts no doubt requires that we be able to ascribe efficacy (of sorts) to their compositional bases, the structural theorist will surely want to insist that there is also a respect in which particular bases derive whatever efficacy they might have from their status as members of more general structural categories. Indeed, it seems that, on such an account, a compositional base L_1 could only be regarded as efficacious (in a way that allows us to account for the efficacy of a structural fact H) if there is a viable non-question-begging account of why all instances of H should induce the event we want to explain (i.e., one that does not postulate L_1 as a cause). Now, it is easy to see how this requirement is met in the one-child policy example: without making any assumptions about what specific constellation of Wangs and Chows actually preceded the Chinese Government's enactment of the one-child policy, we can see how any constellation of births that realized growth of a certain extreme degree could prompt a government to enact such a policy. But suppose we had no reason to make this last supposition, i.e., that we just could not see how facts of that sort could so affect a government. In that case, even

if we could reasonably regard the particular constellation that arose as a cause of the policy, this could form no part of a defence of the explanatory power of the structural fact that growth of such and such a degree arose. Rather, we would have to regard the particular constellation that arose as being a cause in some other capacity, in virtue of realizing some other structural fact. (Perhaps the constellation served as a cause just because it was part of a militantly anti-Catholic cohort - one disposed to vehemently support whatever policies could be expected to antagonize Catholics). And given this, we could not, of course, appeal to the efficacy of that compositional base in order to account for the efficacy of the original structural fact.

Now, moral structural facts seem disadvantaged in precisely this respect. In the cat-burning case, even if it is admitted that the children's behaviour has a kind of causal and explanatory efficacy with respect to Spike's judgment, there is just no viable non-question-begging reason to suppose that it has this efficacy in virtue of its belonging to the class of wrong actions, or in virtue of constituting the moral structural fact that the children act wrongly in burning the cat. We lack any understanding of how an underlying compositional fact could acquire such efficacy in light of its standing as a moral structural fact. So we cannot just explain structural moral efficacy in terms of compositional efficacy. Of course, such compositional facts might acquire their efficacy in light of their membership in some other structural category. For instance, as we noticed above, one might construct a viable generalist but irrealist account of moral structural facts, one that sees them as resulting from cognizers' attitudes have toward compositionally disparate items. For, as we have noted, there is no evident mystery in how a compositional fact could acquire efficacy in terms of belonging to that class of items.

The moral case contrasts sharply in this regard with the psychofunctional and nonmoral macrophysical cases. In each of the latter cases, the candidate supervening

explanatory property can plausibly be regarded as specifying a general feature of the alternative underlying bases about which there is no causal mystery of the kind we see in the moral case. In the psychofunctional case, the general explanatory item is the functional state which can be, on such accounts, realized by disparate neurological states, or even by disparate states of nonbiological systems. In the macrophysical case, it is an observable macroproperty. In neither of these cases is there any mystery about how the general explanatory fact could interact with an observer so as to cause her to make a judgment.

5.3: The Argument from Moral Conversions

If the above arguments are accepted, moral explanations are not, as Sturgeon claims, "as illuminating as" nonmoral explanations. In fact, since they cite facts whose capacity to cause the explananda is itself unexplained, we have no reason to accept them as causal accounts, and hence they are not illuminating at all.

For this reason, moral explanations seem, one might say, "mysterious." And it is just this sort of mystery that the AMR theorist sought to avoid: on the AMR account, there should be no need to posit strange ontological entities in the world or attribute obscure apprehensive abilities to moral cognizers. It seems that an assimilative account which eschewed these claims but posited obscure causal connections between moral facts and nonmoral facts is scarcely an improvement over old-wave moral realism. However, the AMR theorist might reply that we do not need to know precisely how moral properties cause us to judge as we do in order to cite them in a plausible causal explanation for a judgment. For she can point to many examples of nonmoral explanations that seem plausible notwithstanding the fact that they do not provide an illuminating account of how the factors they cite cause the factors which they explain. For example, if the result of many empirical studies showed a strong association between the Santa Ana winds and elevated rates of depression and suicide in Southern California, psychologists might with

justification cite the Santa Ana in causal explanations of those increased levels of depression and suicide, and they would be justified in doing so even if the "underlying" reason for the regularity was unknown. Such explanations cite rules or laws connecting logically distinct but repeatedly conjoined events in accounting for certain explananda. But there may be unknown mediate factors involved, or the events may be conjoined in virtue of their belonging to certain more general, but unknown, classes or types. The Santa Ana may cause depression by causing drought, which in turn causes plants and animals to die, which in turn makes people depressed; or it may cause depression because it is composed of air containing unusually high levels of 'positive ions', and because positive ions adversely affect brain chemistry so as to produce depression. But psychologists do not need to know these further facts in order to know that the Santa Ana causes depression; they merely require statistical correlations that are strong enough to warrant postulating a law requiring the conjunction of the wind and increased depression. In this case, and many others, there is some mystery involved in making an explanation that merely cites a law (for the explanation stops well short of specifying why or how the item in question figures as a cause), but considerably more mystery involved in avoiding the explanation entirely. For instance, if we assume that there is in fact a strong correlation between the Santa Ana and depression, then in spite of the obscurities that are currently germane to the Santa Ana hypothesis, ruling it out leaves the depression unexplained.

Now, this may not seem true of moral realist explanations of moral judgments: since there is arguably an adequate antirealist alternative, moral judgments are not left unexplained when realist hypotheses are ruled out. But certain AMR theorists have cited what they believe is a special class of cases that are not plausibly explained on antirealist assumptions: viz., moral conversions. Moral conversions are changes in moral judgment ostensibly made in response to observations. It seems that, if it is true that antirealism

cannot explain moral conversions, ruling out moral explanations leaves them entirely unexplained. And leaving something unexplained is arguably more serious than explaining something by appealing to a regularity or conjunctive relationship whose nature is not entirely understood.

Richard Werner (1983) has argued that on Harman's account of moral judgments (Werner calls Harman's view "internalism"), it seems moral conversions are not possible.⁸⁶ Since, on an internalist account, how a person responds morally to a situation depends on (i.e., is determined by) his psychological set, one ought to be able to predict how a person would respond to any moral situation given only a knowledge of that person's moral psychology, her relevant nonmoral beliefs, and a nonmoral specification of the situation to which the person responds. Yet, Werner points out, it seems that certain experiences often inspire people to change their mind about a moral issue, and to judge differently. He gives the following example (1983: 657 ff.). Fred is a "pure" utilitarian. That is, he has read about, and has been suitably convinced by, descriptions of utilitarianism, and also finds that, in considering various moral situations, his judgments conform with the dictates of that theory. On Harman's account, Fred's judgments are explicable entirely on nonmoral assumptions about the world and on assumptions about Fred's nonmoral beliefs and utilitarian commitments. Since his "psychological set" determines how he will see and respond to moral situations, it seems we should be able to predict and explain his responses, without mentioning or knowing any moral facts. Werner invites us to suppose that Fred, after watching "Roots" on TV, decides to give up utilitarianism. Utilitarianism, he recognizes, holds that acts or policies are permissible if optimific, but he now sees that slavery is never permissible. Werner points out that such experiences are common, and compares them to instances in which one revises one's

⁸⁶ See also Kupperman 1987: 105-106; cf. Waller 1992.

judgment after considering relevant counterexamples (1983: 659). How is this common phenomenon to be explained?

According to Werner, the realist has a plausible explanation for Fred's conversion: Fred perceived the wrongness of slavery as a result of watching the movie; he had a moral observation that disconfirmed the theory he had held. By contrast, an internalist has no comparably plausible explanation. According to Werner, Harman could not explain the conversion without making at least one of a number of untenable assumptions. First, he could claim that Fred was never really committed to utilitarianism. But since Fred might very well insist that he was so committed, this position inevitably depends on the claim that Fred is mistaken about his own (past and present) commitments. Werner points out that normally we take a person's introspective reports of his own moral convictions as prima facie true. Nonetheless, as he sees it, one might challenge such a report on either of two grounds. First one might cite empirical evidence that Fred is mistaken. But Werner thinks that convincing empirical evidence against them is usually unavailable.⁸⁷ Second, one might suppose that it is analytically true that

PA: "X has a moral theory T, which truly reflects X's moral sensibilities and psychological set if and only if X will find every possible consequence of T morally acceptable."

If PA is accepted as an analytic truth, Fred's conversion becomes evidence that, his own reports to the contrary notwithstanding, he never was truly committed to utilitarianism. But, Werner argues, since PA (a) is not self-evident, and (b) fails the open-question test, it is not analytic. Consequently, it furnishes no reason to doubt Fred's, or anyone else's, sincere reports of their present or past convictions (1983: 660-661).

⁸⁷ Suppose Fred had made certain "nonutilitarian" judgments in the past. One might point to such behaviour as evidence that Fred was not a genuine utilitarian, even if Fred now sincerely believes that he was. But Werner seems correct in supposing that there are many cases in which such evidence is lacking.

Werner suggests that, alternatively, Harman could claim that Fred merely underwent a "profound psychological change" (1983: 661). Before he saw "Roots" his psychological set included only utilitarian elements; afterward, it was supplanted by, or at least came to include, nonutilitarian elements. Werner acknowledges that Fred, and others in relevantly similar situations, undergo profound psychological changes (before watching "Roots," he was a utilitarian, but afterwards he is not), but insists that the change should be interpreted in a way amenable to realism. As Werner sees it, the problem is again mainly posed by the probability that Fred himself interprets his conversion in a way amenable to realism: i.e., that he conceives his conversion as resulting from conscious reflection on the nature of slavery and from his perception of its wrongness (which the film made evident to him). Since first-personal interpretations are *prima facie* true, the burden of proof falls with the internalist to show that his conversion is due largely to "internal" factors (e.g., newly emerged, but hitherto subterranean, psychological inclinations) (1983: 661-662). But, Werner claims, there is no non-"ad hoc" way for the internalist to make such a case. He writes:

My point here is that internalism provides good reason to doubt such reports only if internalism has already been established as explaining all relevant moral observations and reports thereof. But to use the theory to rule out some recalcitrant observations and reports so as to render them nondefeating to the theory seems to constitute a paradigm of an ad hoc theory. Can the internalist convince us for reasons that are not ad hoc, for reasons other than the explanatory integrity of internalism, that observations and reports thereof like Fred's are merely instances where one is psychologically confused about one's own psychological state? If so, then internalism can escape the charge of ad hoc reasoning, of non-falsifiability. But it is just such arguments which are sorely lacking in the internalist's repertoire. Nor do I see, given the present state of philosophy and psychology, from whence such arguments are to come. Once again, the pendulum of the burden of proof seems to swing back in the internalist's direction. (Werner 1983: 662)

Now, if as Werner claims, moral conversions are not explicable on internalist assumptions, and if the fact that moral realist explanations of conversions appeal to unexplained regularities is not a deficiency (since many otherwise acceptable nonmoral explanations do

so as well), then internalism seems to compare unfavourably to moral realism in terms of explanatory power. The moral realist could claim that (a) the fact that moral realist explanations appeal to unexplained regularities is not a deficiency, since many otherwise acceptable nonmoral explanations do so as well, and that (b) moral realism, unlike internalism, can explain moral conversions.

5.4: Problems with Werner's Argument

Notwithstanding the apparent simplicity and eloquence of the explanation on which Werner's argument depends, a number of difficulties afflict his case. One general problem is that his argumentative strategy depends on an implausible view of the "problem with ethics" that Harman has described. It requires that we regard that problem as analogous to the problem of unexplained regularities in science. But they are fundamentally different. The latter problem is a scientific one, whose solution would be scientific. We may not know enough about the composition of the Santa Ana winds, or about the effects of positive ions on brain chemistry, to explain the conjunction of the two events, even though we may cite the winds to explain the depression. But further scientific investigation may uncover the answer. By contrast, the problem Harman cites is metaphysical (and epistemological), not scientific. We think we know how cognizers could perceive and be affected by unjust things, but we do not know how they could perceive and be affected by the fact that something is unjust. This is not because we have not yet discovered some empirically discernible feature common to unjust actions and policies that might be cited as the cause of a conversion, but because moral facts seem not to be the sort of thing that could be empirically discerned. Since the problem is not an empirical one, we could not expect to uncover a solution through scientific investigation, and so it cannot be dismissed as an empirically resolvable difficulty that is common to many nonmoral explanations. But this suggests that moral conversions cannot be explained on realist assumptions. For

the same arguments raised against moral explanations of moral judgments in general can and should be raised against moral explanations of moral conversions. And those arguments do their work even if, following Werner, we grant that most individuals who experience conversions are commonsense moral realists whose first-personal interpretations of their conversions must be accorded a presumption of truth. Thus, even if Fred thinks he has suddenly perceived the injustice of slavery while watching "Roots," and believes himself to have revised his moral views as a result of that perception, we cannot cite the fact that slavery is unjust as the cause of his conversion because we have no idea how a fact of that sort could causally affect him.

Another difficulty concerns the particular example Werner has chosen to support his thesis. As Werner sees it, Fred's conversion is best, and only, explained, on the assumption that there is an inquiry-independent moral fact imbedded in the film that Fred has perceived, and which is incompatible with the moral theory Fred had initially held. But if this were correct, one would expect the film to have quite specific effects upon other viewers. At least, one would expect it to cause most of the other utilitarians watching it to convert to "deontism" too, and to cause most of the "deontologically inclined" viewers to become even more firmly committed to their theory. This is, after all, what ostensibly happens in science: e.g., the discovery of the HIV virus confirmed the suspicions of those researchers who had all along thought AIDS was caused by a virus, and caused most of those medical researchers who, prior to that discovery, thought AIDS was caused by something other than a virus to convert to the viral theory. But one would not expect this kind of overall response in the "Roots" case. Even if one could expect "Roots" to cause many utilitarians to convert to deontism, there is no reason not to expect it to cause just as many other viewers to convert in reverse. Suppose for instance that Ted, unlike Fred, is a "pure" deontic theorist prior to seeing "Roots." Ted's main principle is that violent acts

against other human beings are never permissible. But upon seeing "Roots" he decides to give up his deontological theory. "Roots" convinces him that the slaves would have been fully justified in responding violently to their oppressors if that response would have brought the quickest possible end to the suffering inflicted upon them by the institution of slavery. This in turn causes Ted to embrace utilitarianism, for it now seems to him that maximal avoidance of suffering is more important than conforming to deontic principles. Seeing "Roots" has thus caused Ted to convert, but in the opposite direction.⁸⁸ Now, if there were an empirically-discernible moral fact imbedded in the film that could disconfirm utilitarianism in the way that the discovery of HIV disconfirmed non-viral theories of AIDS, one would expect the Ted scenario to be an extremely improbable one. (I doubt anyone responded to the discovery of HIV by converting from the viral theory to one of the alternatives!) Instead, one would expect the film, in addition to affecting Fred and similarly-inclined persons in the way Fred was affected, to make more firm the convictions of those who are initially like Ted. But in fact the Ted story seems just as probable as the Fred story; one would not expect "Roots" to have a uniform, or even remotely uniform, effect on which of these moral theories people are generally attracted to. Given this, Werner's appeal to the prima facie truth of each cognizer's own account of their conversions hardly helps his case. "Roots" causes both Fred and Ted to convert, but they will presumably give different and incompatible accounts of the cause of their own conversions. Fred will explain his conversion as the result of his perception of a "deontic"

⁸⁸ See Waller (1992: 135): "Objective moral facts are not necessary to account for changes in moral principles, and thus such real moral facts are an unnecessary ontological burden. But the burden is more than heavy, it is crushing. Any appeal to moral facts to account for moral belief changes will have to carry a double load. Changes in morality do not move in a single direction; in many cases they seem to move in opposite directions with almost equal alacrity. Given my own values, I would acknowledge happily that objective moral facts led many people in the late 1960s to reject militarism, jingoism and greed in favour of a more peaceful and egalitarian and noncompetitive society - were it not for the fact that during the same period there were children of leftists parents who rejected egalitarian principles to embrace Ayn Rand and rugged individualism."

moral fact that falsifies utilitarianism; Ted will see his conversion as resulting from his perception of a "utilitarian" moral fact that falsifies deontology.

Notice that the realist now assumes the burden of explaining why Ted responds in the way that he does, rather than just keeping his deontic approach (and perhaps just switching deontic principles), and more generally, that of explaining why "Roots" does not prompt a more uniform overall response. Now, I concede that Werner might explain this in a way that is generally compatible with moral realism. But it seems to me that he could do this only by discarding the original hypothesis that the moral fact of slavery's injustice was simply there to be perceived by ordinary cognizers such as Fred, in favour of a much more complex account. The alternative account would have to be more complex because it would have to explain why different inquirers respond differently notwithstanding the supposition that they are confronted by the same set of moral facts. It would thus presumably have to include either the assumption that some of the viewers had a perceptual or cognitive deficiency that prevented them from discerning the moral facts, or else the assumption that the moral facts contained in the film are just not easily discerned by all normal viewers. And each of these options would appear to present still further explanatory burdens. For instance, since we do not need to suppose that either Fred or Ted are impaired in the usual ways (i.e., by being drunk, stupid, uneducated, or in the grip of some distorting ideology) in order to expect them to disagree, it is not obvious just what sort of "deficiency" either of them could possess that could prevent him from discerning whatever moral fact is in the film. Nor is it obvious why, if moral facts are really the sorts of things that can at least sometimes be discerned by ordinary inquirers like Fred, they cannot be perceived by most other similarly-equipped inquirers at least most of the time.

But we hardly need to assume these further burdens in order to explain why people respond to films such as "Roots" in the ways that they do. A simpler and more plausible

explanation of their responses is suggested by Postow's recent explanation of Fred's conversion. Postow writes:

. . . before watching "Roots," the observer had not experienced strong emotional identification with the victims of slavery, but "Roots" induced that emotional identification through literary and cinematic devices. This is why the observer is only now disposed to judge slavery wrong. (Postow 1985: 287)

Ted's conversion could, it seems, be explained in this way too: prior to seeing the film, Ted had been incapable of identifying strongly with the victims of slavery, and this allowed him to hold steadfast to his injunction against violent opposition to oppression. But "Roots" induced that emotional identification in the same way it was induced in Fred, and this is why Ted gave up his original theory. Thus, what accounts for the difference between Fred's and Ted's conversions is not a difference in their cognitive or perceptual abilities, or the relative obscurity of the facts contained in the film. Rather, it is just that the film contains cinematic and dramatic elements that are designed to prompt empathy for the victims of slavery, and that each inquirer confronts the film bearing a different set of moral assumptions, and that consequently each arrives at his respective conclusion via a somewhat different line of reasoning. And this sort of explanation could account for any overall response to any film: whether a particular work prompted a uniform response or not, one could always cite deliberately-embedded cinematic and dramatic devices, and the fact that different cognizers begin with different (or similar) sets of moral assumptions, as the cause. Notice further that we do not have to assume that any particular set of moral assumptions are true in order to make these explanations.

An AMR theorist may still point out that we could not avail ourselves of this explanation for conversions that are made in response to actual (as opposed to cinematic or literary) situations. But here the realist encounters another problem. To be justified in converting from one scientific theory to another, someone must confront actual empirical evidence favouring one theory over another. But it seems that one can gain moral insights,

and be entirely justified in changing one's moral opinion about something, just by contemplating certain imaginary examples. And if cognizers might gain insights and convert in this way, explaining their conversions could hardly require the assumption that they have perceived empirically-discernible moral properties. In fact, the circumstances under which we make, and change, our moral judgments seem exactly as one would predict if internalism were true. Since internalism says that cognizers assess moral situations against their feelings or beliefs about actual and counterfactual situations and not against empirically-discerned moral facts, it should allow inquirers to arrive at moral conclusions by considering imaginary or hypothetical situations rather than by confronting actual ones. And in fact we typically do arrive at moral conclusions in this way (for instance, most of Harman's readers will probably agree that cat burning is wrong even though they have never seen a cat being burned); and we typically do not change our moral opinions in response to actual scenarios (e.g., presumably very few of those who do think that burning a cat is okay will change their minds after seeing an actual instance of cat burning, for we may plausibly suppose that anyone whose sympathies were so "jaded" that they could not see this just by considering the imagined scenario would be unlikely to change his mind when he encounters the actual event).

Of course, there is no question that first-time exposure to an actual situation sometimes does cause a person to change their moral view of something. I know, from sources I have every reason to think are reliable, that currently many countries have large populations of impoverished children, but I might still adopt a cavalier attitude toward the need for more foreign aid from Canada - until I tour Calcutta or Rio De Janeiro and witness the conditions directly. And perhaps there are some people who are inclined to be cavalier about cat burning or slavery until they see an actual instance of these things. But such conversions can normally be explained on a number of possible nonrealist

assumptions. Postow arguably hints at one of these explanations when she points to the ability of "Roots" to cause Fred to "identify with" the victims of slavery more strongly than he was able to manage before he saw the film. Clearly actual situations can have this kind of effect too. So, for example, my trip to Rio might cause me to change my mind about the need for Canadian aid by allowing me to identify more strongly with the victims of poverty in third world countries, and seeing an actual instance of cat burning may allow Spike to identify more strongly with the suffering of the cat than he could manage beforehand. Such conversions might also be explained on the assumption that witnessing the scenarios directly has allowed the person to achieve a fuller and more detailed understanding of the relevant object of assessment. My visit to Rio would leave me with many newly-acquired beliefs about actually existing poverty, e.g., beliefs about particular persons I have met, about the sort of lives they must endure, about their daily struggles, and so forth. My view of life in these places would have been brought into much sharper and more detailed relief as a result of my trip, and this would be sufficient to explain my conversion. The same point may be made about people who are inclined to say that cat burning or slavery are okay, but then change their minds upon seeing the actual actions: Spike might suddenly change his judgment about cat burning just because prior to seeing the actual event he had failed to appreciate just what burning a cat must involve. Perhaps he failed to appreciate the depth of suffering of which cats are capable, or believed that cats are incapable of experiencing pain, or that burning cats die quickly and painlessly; seeing the actual event convinces him that he was mistaken, and forces him to revise his moral judgment in light of his new understanding of the nonmoral facts.

In order to avoid this objection, it seems the AMR theorist would have to cite cases in which a cognizer - who from the start possesses maximum empathetic, identificatory, and imaginary abilities, and a full understanding of what an action would involve if it were

committed - forms a moral judgment about that action on the basis of his consideration of hypothetical scenarios, and then converts after witnessing the action actually being committed. Something like this can presumably happen in science: an AIDS researcher who is against the viral hypothesis might be equipped with the aforementioned affectations as well as a solid understanding of what would have to obtain if the viral theory (or one of the alternatives) were true, but we could expect him to convert after he or another scientist encounters empirical evidence for the viral theory. So, a parallel moral case would presumably be something like this: Spike is as empathetic and imaginative as anyone could possibly be. He also knows quite a lot about cats and about the actions of incendiary liquids, and on the basis of this knowledge he forms an opinion about what would happen to a cat that was set on fire. He expects, in particular, that it would experience a great deal of terror and pain. On the basis of this expectation (together with his imaginary and empathetic skills), he forms a moral opinion about the morality of cat burning. Then, he goes outside, rounds a corner, and sees children actually setting fire to a cat! Though he has never seen a cat being set on fire, it turns out that his expectation of what would happen was entirely accurate. Nonetheless, he changes his moral opinion after seeing the actual event. Ex hypothesi, a conversion of this sort could not be explained on either of the nonrealist assumptions mentioned above. But I doubt anyone who was like Spike in these respects would actually change their moral opinion under such circumstances.

What about Werner's claim that arguments from the "explanatory integrity" of internalism are ad hoc and nonfalsifiable? Presumably, the defence of internalism would be ad hoc if it depended on a claim, or set of claims, whose only rationale is that it salvages some theory, e.g., internalism. Thus Werner insists that the case for internalism must be ad hoc because it requires that we deny cognizers' own first personal interpretations of their conversion experiences as being the result of interaction with efficacious moral properties,

and because this view can only be refuted by assuming that internalism is true. But my arguments do not assume that internalism is true or that the requirements of that theory (whatever they may be) are met; they cite reasons for thinking that internalism is true (e.g., that how one would expect people to judge in various circumstances is contrary to how one would expect them to judge if the realist's account of conversions were true). The internalist can concede that first-personal interpretations of conversions are prima facie true, and even that most people are inclined to regard their conversions as resulting from the apprehension of hitherto unnoticed real moral properties, and then go on to argue, as I have, that consideration of certain counterfactual scenarios suggests that those first-personal interpretations are mistaken, without making any ad hoc assumptions. Nor is the internalist thesis "unfalsifiable." If, for instance, it were shown that cognizers often or even sometimes judge contrary to how I have suggested they would in the circumstances I have cited (i.e., that they sometimes judge in a way that is consistent with how one would expect them to judge if the realist's account of conversions were true), then my case for internalism would be defeated.

5.5: Richard Boyd's Response to the Moral Input Objection

If my case to this point is sound, Brink, Sturgeon, and Werner have established neither the relevance nor the indispensability of assumptions about moral facts to reasonable explanations of various sorts of moral judgments. It will be noticed that my case depends crucially on the assumption that moral facts, properly construed, cannot be causes - an assumption that is itself rationalized by the Harman-inspired observation that we simply have no understanding of how that kind of thing could be the cause of anything. Now, it might be thought that this problem has been adequately addressed in the literature by Richard Boyd (1988). Boyd defends an account of the meaning of moral property-ascribing terms that, if true, makes plausible the idea that our moral judgments are regulated by actual moral features

of the world. Boyd's argument is difficult and sophisticated, and an adequate assessment of it would require more space than can be afforded here. Fortunately, the task of giving a detailed assessment of this theory has been taken up by others. In this section, I briefly rehearse Boyd's account, and mention some reasons to reject it.

Boyd's theory is an account of the meaning and definition of moral property-ascribing terms. His theory is attractive for two reasons. First, it seems to address the concern about input by providing a means of understanding how moral reality "regulates" the application of moral property-ascribing terms in a way that is consistent with a plausible realist account of scientific terms. But it is also attractive for another reason: as an account of the definition of moral terms, it arguably remedies a deficiency in Brink's account, taken on its own. Recall that Brink seeks to avoid making moral judgments depend for justification on the semantic test of properties. Because, for Brink, moral properties or states are constituted by nonmoral properties or states, their ontic and epistemic status is purported to be independent of any semantic analysis of moral terms; rather, it results from the nature of the areas of reality those terms are used to describe. All of this is supposed to explain how a certain action or state of affairs can be, say, good, even if the semantic test of properties is rejected. However, avoiding the traditional objections in this way arguably comes at the cost of making the meaning of moral terms obscure. If moral properties are natural properties, but the meaning of the terms used to designate them escapes naturalistic definition, then moral terms appear indefinable in terms of the properties they putatively designate. Thus, Brink's account seems to obscure the connection between moral language and features of moral reality. What is required is a way of defining moral terms that makes their primary function (that of specifying natural/moral properties) clear, one that nevertheless makes the existence of certain moral properties compatible with failures of reduction, and one whose use is standard outside morality. If it could do all of this and also accommodate the notion that our use of moral terms is to a

significant extent "guided" or "regulated" by objective moral properties and states, it would, via a semantic theory, represent an adequate response to the moral input objection.

This is exactly the type of semantic theory recommended by Boyd (1988: 194-206). Following Locke, Boyd claims that sound scientific explanations are possible only if the natural kinds that figure within them are definable in terms of "corpuscular real essences" rather than the "nominal essences" which refer exclusively to sensible qualities (1988: 194). Sound scientific explanations require, according to Boyd, a rejection of the traditional empiricist mode of defining natural kinds. The alternative "real essence" approach he recommends itself requires that traditional empiricist semantics be replaced by a naturalistic conception of reference. Boyd finds such a conception in the causal accounts of reference defended by Putnam and Kripke. Boyd combines a (Putnam-inspired) causal theory of the reference and meaning of natural kind terms with a "homeostatic" account of their definition. On Putnam's account, the meaning of natural kind terms like 'water' is determined not by any set of descriptions conventionally associated with the items to which the terms refer (e.g., the ability of a liquid to extinguish fires), but partly by the existence of certain causal connections between their employment by inquirers and real essences of the kinds in question (e.g., having the microcomposition H₂O), and partly by the intentions with which cognizers employ the term in question (Putnam 1975). Users of the term 'water' intend to refer to the underlying real essence of water - an essence that is discoverable through scientific inquiry. Since, Boyd argues, the causal connections relevant to reference are exactly those involved in the reliable formation of belief, such causal theories account for the idea that reality itself "regulates" beliefs and the application of predicates (1988: 195). For Boyd, this regulation is indirect and "socially coordinated," but consistent with the causal accounts (1988: 195). When it obtains, it provides us with "epistemic access" to the properties to which the terms contained in scientific statements refer. Moreover, while traditional "real essence" definitions specify necessary and

sufficient conditions for the application of a term (e.g., water = H₂O), Boyd thinks that natural kind definitions can be modeled on forms of definition championed by "property-cluster" and "criterial attribute" theories in the ordinary language tradition (1988: 196). On such accounts, factors guiding correct predicate application are more flexible: an item falls under the extension of a term when a sufficient number of properties can be ascribed to it, and it is a conceptual matter how many properties are sufficient. But since the concepts themselves are vague, there is a degree of indeterminacy legitimately associated with the definitions (1988: 196). Boyd thinks that there are natural definitions in science that similarly involve conceptual vagueness and extensional indeterminacy. But in science, such definitions are justified by the task of identifying the elusive realm of essences which underlie appearances, and which could not proceed as effectively were an artificial rigidity imposed on the definition of natural kinds. Boyd calls these scientific definitions "homeostatic property-cluster definitions" (1988: 196). Thus, for Boyd, users of at least some natural kind terms, i.e., those for which homeostatic definitions are appropriate, intend to refer to certain underlying constellations of properties existing in homeostasis; and their employment of such terms can itself be seen as "guided" or "regulated" by those underlying homeostatic essences, in the way that their use of 'water' is guided by H₂O essences.

Now, one area of difficulty for the moral realist is the putative "essential contestability" of moral concepts, and the associated disagreement about what natural and social properties fall under the extension of value-ascribing moral terms. This disagreement makes it practically impossible to define moral terms by specifying necessary and sufficient conditions for their application, and this practical impossibility is widely taken as evidence against the claim that moral terms refer to features of moral reality. This is because widespread, seemingly unavoidable, contestability can plausibly be explained by the hypothesis that such terms are expressive, not descriptive. But, according to Boyd, the fact that extensional indeterminacy

(and an associated conceptual vagueness) is tolerated, and in fact, required, in science, makes it unreasonable to demand determinacy of moral terms (1988: 200ff.). Rather, the terms germane to property ascription in both ethics and science are definable homeostatically in terms of contingently-clustered properties. Boyd formulates a conception of goodness ("homeostatic consequentialism") that (he supposes) conforms structurally to the scientific definitions (1988: 203ff.). Thus moral property-ascribing terms are definable using a mode of definition that is standardly employed in the natural sciences ("homeostatic property-cluster definitions"). Moreover, according to Boyd, we can and should combine this account of the definition of moral terms and of the associated moral homeostatic essences with a Putnam-inspired causal theory of the meaning of moral terms (1988: 188, 191-192). On Boyd's account, moral terms like "good" function like natural kind terms in that users of those terms intend to refer to a certain single objective natural property or property cluster (homeostatically construed, of course) whose real moral essence is discoverable through moral inquiry. The result is a naturalistic moral realism that thoroughly assimilates both moral properties and terms to nonmoral counterparts, in a way that is amenable to realism in both areas.

It is significant that Boyd's theory has been rejected, in a series of articles by Mark Timmons and Terrance Horgan, on the following grounds.⁸⁹ Horgan and Timmons point out that the essentialist view of natural kind terms seems tenable largely in light of how we respond to Putnam's familiar Twin Earth parable. For instance, we think that a possible world that was exactly like Earth in every respect except that chemical XYZ is found in place of H₂O would not contain water, even if it contained people who used the term 'water' to refer to the stuff in their oceans, etc. On the contrary, we would say that such a world contains something other than water (i.e., XYZ), and that the occupants of that world employ the word 'water' with a different meaning than we do (they use it to refer to XYZ, we use it

⁸⁹ Timmons 1990; Horgan and Timmons 1990-1991; 1992.

to refer to H₂O). According to Horgan and Timmons, all of this points to a test of Boyd's theory: if moral terms derive their meaning in the way natural kinds do on Putnam's account, then we should expect to have parallel intuitive responses to a moral analogue of the Twin Earth parable, viz., the Moral Twin Earth parable.⁹⁰ That parable runs as follows: suppose that inquirers on Earth manage to get their moral beliefs into reflective equilibrium of the sort envisioned by the Brink-BonJour account, and that the beliefs in that state together imply some substantive moral theory, say, some form of consequentialism. On the Boyd account of the meaning of moral terms, the moral terms employed by these inquirers would be used to refer to natural properties whose real moral essence is described by the consequentialist theory they hold. The real essence postulated by a theory of this sort would be, of course, a "consequentialist" one. But, Horgan and Timmons suggest, suppose that there is a world, Moral Twin Earth, that is exactly like Earth in every respect except that its occupants accept a deontological theory when they get their convictions into reflective equilibrium. The occupants of Moral Twin Earth use the terms 'good' and 'right' to refer to natural properties whose real moral essence is "deontological" rather than consequentialist. As Horgan and Timmons see it, there is an important difference in the way we would naturally characterize the interplanetary disagreements in the Twin Earth story and the Moral Twin Earth story. It is natural to characterize the difference between Earthlings and Twin Earthlings as a difference in meaning, not belief: Earthlings and Twin Earthlings do not hold different theories about the composition of water, they just attach a different meaning to the term 'water'. But we would not characterize the difference between Earthlings and Moral Twin Earthlings in this way. It would be more natural to characterize that difference as a difference in moral belief rather than meaning. Earthlings and Moral Twin Earthlings do not employ the term 'good' in different senses, they just hold different and incompatible theories and

⁹⁰ See Timmons 1990: 118-121; Horgan and Timmons 1990-1991: 457-461; 1992: 244-248.

beliefs about goodness. And this is a reason to doubt that the causal account of meaning Putnam applies to natural kind terms can be applied to moral terms.

To this we may add another problem with Boyd's semantic theory. The difficulty becomes apparent when we consider the justification for the use of homeostatic definitions. It is reasonable to think that we should opt, if possible, for modes of definition that include more determinate and easily specifiable conditions for the application of property-ascribing terms. In general then, homeostatic definitions are something to be avoided, and are a mode of last resort. They are appropriate when we are confronted with an area of discourse involving terms definable in no other way, especially in areas which involve concepts whose contestability seems unavoidable. But, crucially, their use is not sufficiently justified by the mere fact of contestability. If it were, we could, a priori, rule out contestability in any area as evidence that the concepts involved lack substantive reference to areas of reality. Their employment thus requires further justification. According to Boyd, their use in science is justified instrumentally. Were the concepts employed in scientific explanations constrained by an artificial rigidity, their explanatory power would be compromised:

There are natural kinds, properties, etc. whose natural definitions involve a kind of property cluster together with an associated indeterminacy in extension. Both the property-cluster form of such definitions and the associated indeterminacy are dictated by the scientific task of employing categories which correspond to inductively and explanatorily relevant causal structures. In particular, the indeterminacy in extension of such natural definitions could not be remedied without rendering the definitions unnatural in the sense of being scientifically misleading. (1988: 196)

But what, we may ask, rationalizes their use in ethics? If, as Boyd appears to think, it is their role in facilitating causal moral explanations, his project comes to depend on the thesis that assumptions about moral facts can explain. So although it is ostensibly designed to explain the efficacy of moral facts, it does not in fact do this. To appreciate this more fully, consider that Harman's problem remains even if we assume that moral facts are facts about homeostatically-arranged clusters of natural and social properties: the capacity of the fact that

a given cluster comprises a moral property to cause and causally explain things seems just as mysterious as seems the efficacy of moral facts more simply construed.

CHAPTER SIX

ASSIMILATIVE MORAL REALISM AND SUPERVENIENCE⁹¹

6.1: Introduction

AMR is a form of ethical naturalism. It claims that moral properties "are" natural⁹² properties, just as nonmoral macrophysical properties "are" the microphysical states that compose them. In the course of defending this theory against the traditional view that the lack of any naturalistic reduction for moral terms somehow implies that naturalism must be false (what Brink calls the "is/ought" thesis⁹³), Brink claims that both moral and nonmoral macrophysical properties share certain important features: specifically, both may be (and typically are) constituted, supervening, and synthetically necessitated. The idea, stated simply, is that moral properties can reasonably be held to "be" more basic complexes of natural properties without being reducible to those basic properties, because moral properties can be regarded as constituted by, and supervening upon, their naturalistic bases, and because properties so construed can in general be held to "be" the lower-level

⁹¹ This chapter contains revised versions of arguments that appear in Yasenchuk 1995.

⁹² Actually, Brink also includes among types of "base" properties, "social scientific" properties (see, e.g., 1989: 9), and even "psychological" properties (e.g., 1989: 161). He provides no explanation or examples. Though questions might be raised about whether Brink too uncritically accepts such properties within the scope of his theory, I shall not raise them here. Unless otherwise stipulated, I shall abbreviate Brink's cumbersome phrase "natural and social scientific properties" to simply "natural properties" and intend the latter to include natural, social, and psychological properties.

⁹³ The source of the view is thought to be Hume 1739: III, i, 1/469-470. For a representative set of discussions, see Hudson 1969.

properties that realize them even though ascriptions of the upper-level constituted property are not reducible to ascriptions of any particular set of basic properties (Brink 1989: 156-163).

It seems to me that Brink has succeeded in showing that if moral properties are like he thinks they are, then this traditional objection to ethical naturalism must fail. That is, ethical naturalism, and thus his naturalistic version of moral realism, is not undermined by the existence of an "is/ought" gap. However, I argue here that his view of moral properties is untenable. Since that account implies a corresponding view of moral claims, viz., that ascribing a moral property is to ascribe a property that possesses the features he mentions, one way of assessing it is via an analysis of moral and nonmoral macrophysical property ascriptions. I engage in this endeavour here. I argue that even if moral property ascriptions can tenably be construed as ascriptions of properties that bear the features mentioned by Brink (features shared by properties that are commonly ascribed by science), there remain significant modal differences in the way the two types of properties are assigned to the world. The differences are significant because the particular view of moral supervenience that I will defend seems incompatible with AMR's account of the epistemic justification of moral claims.

To be viable, AMR must satisfy two constraints. First, it must be compatible with a plausible account of what I shall call our "moral thinking," i.e., of the ways in which we ordinarily ascribe moral value to the world, including the strength of the necessity with which such ascriptions are thought to obtain (whether logical, conceptual, physical, etc.). Recent work has shown that the modal logic of putative moral property ascriptions is directly relevant to how they are to be understood from a metaethical point of view (whether they be construed as conceptual claims, metaphysical claims, etc.).⁹⁴ A realist

⁹⁴ See Blackburn 1985; Klagge 1984: 370-380; 1987: 312-315; 1988: 461-470; McFetridge 1985: 245-258; Dreier 1992.

metaethical theory cannot therefore neglect these features of property ascription, at least not without explanation. Second, because it is an "assimilative" account, it must construe the relation between moral and natural properties as similar to that which holds between types of nonmoral properties for which a realist explanation is plausible.

Section 6.2 contains an exegetical account of Brink's views on constitution, supervenience, and necessity; 6.3 draws further distinctions that I argue need to be drawn in order to assess the force of the constitution account and 6.4 reassesses Brink's position in light of these distinctions; 6.5 explores the type of nomological principle appropriate to nonmoral macrophysical property ascriptions; 6.6 seeks such a principle for moral property ascriptions; 6.7 examines the case of psychofunctional property ascriptions; 6.8 assesses the implications of the results of those inquiries for the assimilative project.

6.2: Brink on Moral Constitution and Supervenience

Can moral properties be understood in such a way as to satisfy the above constraints? Brink characterizes the relation between moral and natural properties in terms of two concepts, that of constitution, and that of supervenience.

While he does not rule out the possibility that some moral properties may be identical to certain natural properties, Brink thinks that moral properties are often "constituted by, though nonidentical to," specific natural property bases. His concept of constitution can be compared to that of identity. Identity is a type of constitution relation (1989: 157). But while properties that are identical to their bases could not possibly be realized via some other base, constituted but nonidentical properties are, to use a familiar philosophical expression, "multiply realizable."⁹⁵ He writes:

⁹⁵ He construes the claim that moral properties are identical to their bases as the claim that "moral terms and certain natural and social scientific terms express or designate the same properties (i.e., refer to the same properties in all possible worlds" (1989: 157). Multiple realizability, by contrast, presumably allows moral terms to designate different natural properties across (and within) possible worlds. That the identity of an effected

F can be G even if the property (or properties) designated by 'F' is not (or are not) the same as that (or those) designated by 'G'. If G actually composes or realizes F, but F can be, or could have been, realized differently, then G constitutes, but is not identical with, F. (1989: 157)

Brink suggests two ways in which compatibility with multiple realizability benefits a constitution account of moral properties. First, it (unlike the identity account) is amenable to our moral thinking. We often ascribe the same moral properties to different objects sharing few or no common base properties (1989: 157). Second, it facilitates Brink's assimilation of moral and nonmoral properties. He points out that we ordinarily make constitution claims in many cases involving high level nonmoral properties for which a realist account is plausible: for example, in characterizing the relation between medium-sized objects such as tables and their microphysical makeup (1989: 158). According to Brink, the fact that such nonmoral properties are constituted by, though nonidentical to, their base properties is not typically regarded as warranting antirealism about them; hence it could not plausibly warrant it in the case of moral properties. To Brink's two points we may add a third: moral property-ascribing terms do not lend themselves to easy linguistic analysis, and certainly not to an analysis that permits their extensions (i.e., the natural properties to which they putatively refer) to be precisely determined (Moore 1903: 15-17). Multiple realizability helps to explain this lack of success in a way that is amenable to naturalism. If moral properties are multiply realizable, the extensions of the terms that denote them may be multifarious, and so resist precise elucidation.

property with its base precludes its multiple realization is now generally recognized in the literature in the philosophy of psychology. There, identity is associated with reductionism and the latter with the exclusion of multiple realization. See Putnam 1967; Kim 1972; 1992.

According to Brink, on a naturalist account, the supervenience of a property F on a property or set of properties G follows from the fact that F is constituted by G.⁹⁶

Supervenience is

. . . a nomological or lawlike relation between, say, properties such that one property F (the supervening property) supervenes on another property or configuration of properties G (the base property or properties) just in case it is a law that if something is G, it is F. One consequence of supervenience is that two things cannot differ in their supervening properties without differing also in their base properties. (1989: 160)

Brink's idea appears to be that constituted properties supervene upon the properties that constitute them because constitution is an account of property realization that bears nomological implications. Constituted properties are not thought to be realized by their bases arbitrarily, but as a matter of law.

Following other writers on the topic, Brink distinguishes two types of supervenience: "strong" and "weak." Strong supervenience specifies a nomological relation that holds "necessarily," weak supervenience, one that holds only "contingently":

With strong supervenience, base properties necessitate their supervening properties. Once the base properties are fixed, the supervening properties are thereby determined A set of properties F strongly supervenes on a set of properties G just in case G necessitates F Weak supervenience states a contingent relation (If, even though G is sufficient in actual circumstances to produce F and would also produce F in relevantly similar counterfactual circumstances, the properties causally connected with G could have been different enough for G to produce H rather than F, then F only weakly supervenes upon G. (1989: 160-161)

Brink's sketch of constitution appears to be neutral regarding which of these laws holds in specific cases. Accordingly, Brink claims that ethical naturalism "implies both strong and weak supervenience" (1989: 160). However, he apparently also thinks that only strong supervenience adequately describes the actual nomological relations that hold between

⁹⁶ For Brink, constitution and identity are "construals of ethical naturalism"; he states further that "The supervenience of moral facts and properties follows from, but does not establish, ethical naturalism" and that "Naturalists claim that moral properties supervene on natural properties because moral properties are constituted by natural properties" (1989: 160).

moral properties and their bases.⁹⁷ Although moral properties always strongly supervene on some base natural property or set of properties, we often encounter difficulties identifying the properties required to literally necessitate a moral property; in these cases we "settle for identifying the weak supervenience base" (1989: 161). The difference between the two principles is thus purely epistemological, reflecting the different epistemic situations within which cognizers may find themselves. Brink sees no difference between moral and nonmoral properties in this regard. He notes that both functionalists in the philosophy of psychology who examine the nature of pain and moral theorists who examine the nature of justice typically make weak supervenience claims. The rationale in each case is presumably the same, i.e., that the theorist deals with a very complex subject matter about which much is still unknown. It follows that a critique of Brink's view of the real relations holding between real properties should focus on his account of strong supervenience.

With strong supervenience, he claims, base properties "necessitate" supervening properties. Brink follows Kripke (1980) in construing necessity as a "metaphysical" notion. As such it contrasts with analyticity, which is "linguistic or semantic" (1989: 166). It is this dichotomy that permits Brink to resist an antinaturalism based purely on semantic considerations (roughly, one that infers that moral properties are not natural properties from the fact that moral property-ascribing statements are not analytic truths). Brink claims that analyticity is no test of property identity or constitution. Base properties may thus necessitate higher level properties even though statements affirming this necessity are not analytic truths. In the moral case, Brink claims, "the naturalist's identity or constitution claims can be construed as expressing synthetic moral necessities" (1989: 166). So, Brink

⁹⁷ Cf. McGinn, 1981: 176, who takes supervenience relations to be brute metaphysical facts.

holds that moral and nonmoral macrophysical properties have this much in common. They are constituted by, and so (strongly) supervene upon, their respective base properties; moral and nonmoral property ascriptions state synthetic necessities.

6.3: Necessity, Supervenience, Constitution: Some Distinctions

Now, it seems to me that if moral properties and their nonmoral counterparts were alike in the respects described above, then any objection to ethical naturalism that is based entirely on antireductionistic assumptions must be false. This is because the constitution account of properties explains handily how a set of higher level properties could be realized by a complex of lower level properties even though upper level property ascriptions are not reducible to ascriptions of lower level properties. So if moral properties are just like ordinary constituted supervening nonmoral macrophysical properties, then ethical naturalism might be true even if it is impossible to reduce ascriptions of moral properties to ascriptions of natural properties.

However, this is not sufficient to establish that moral and nonmoral macrophysical properties are assigned to their respective bases in the same way, and so does not establish the desired assimilation of moral and nonmoral macrophysical properties. To appreciate why, we need to draw some distinctions.

Constituted properties have two central features. First, since they are produced by their respective base properties as a matter of necessity, and since those necessary connections may be defined nomologically, they are supervenient. Second, since constituted properties may be produced via base properties other than those that figure in any given supervenience law, they are multiply realizable. In the discussion that follows it will be useful to keep in mind these features of constituted properties. Thus construed, the concept of a constituted property is essentially modal: it involves claims of necessity (and their associated nomological principles), and possibility affirmations. These modal

concepts can be disambiguated, and this will complicate our conception of constitution considerably.

That they can be disambiguated is clear from recent work on the topic of supervenience. A good framework for a more detailed examination of constituted properties is provided by Simon Blackburn (1985). Blackburn's paper discusses supervenience, not constitution. Nonetheless, like Brink, Blackburn construes supervenience as nomological necessitation, and since neither writer takes the supervenience of effected properties to be incompatible with their multiple realizability,⁹⁸ much if not all of what Blackburn says about supervening properties may be applied to constituted properties. Blackburn has, of course, constructed an argument against moral realism of his own, and I hope to establish a rather similar thesis here, but I wish to make use not of Blackburn's argument, but of certain elements of the framework he constructs in elaborating his claims.

In his paper, Blackburn recognizes a much greater diversity of supervenience principles than Brink does. The diversity has two sources. First, he points out that necessity claims differ in strength; he recognizes "physical," "metaphysical," and "conceptual" necessities (and possibilities). The differences between them become apparent when we conceive of necessity claims in terms of possible worlds. A necessity claim, expressed in that idiom, can be formulated as something approximating

□: P (some state of affairs) necessarily obtains if and only if it obtains in every possible world.

But what worlds are possible? The answer depends both on the type and content of the theory that motivates the necessity claim. Varying types of theory suggest correspondingly

⁹⁸ We have already noted Brink's views; see Blackburn 1985: 60-61.

different construals of '□'. Physical or scientific theories make claims that hold with the following type of necessity:

□_p: P (some state of affairs) necessarily obtains if and only if it obtains in every physically possible world.

Since what is considered physically possible depends on what scientific theories one accepts, whether something obtains with this type of necessity depends on the content of those theories. For example, currently accepted scientific theories make certain claims about the behaviour of magnetism, gravity, etc., that are held to obtain in every world that is possible as far as those and other currently accepted theories are concerned. While it is conceivable that they not obtain, it is physically necessary that they do. Similarly, metaphysical theories motivate claims that hold with the following type of necessity:

□_m: P (some state of affairs) necessarily obtains if and only if it obtains in every metaphysically possible world.

Since the concept of metaphysical modality falls somewhere between physical and conceptual modalities, it implies that something can exist in some world that is physically impossible but conceptually and logically possible without existing in every conceivable and logically possible world. Just what kind of world this is held to be will depend on the content of the metaphysical theory that motivates the claim. Different theories suggest different views about what is metaphysically possible. For example, a metaphysical supernaturalist will regard worlds containing God, spirits, supernatural properties, etc., as metaphysically possible, while a materialist presumably would not. The materialist just insists that only material things can exist. The materialist, qua materialist, can thus regard many physically impossible worlds as metaphysically possible (e.g., worlds containing partly gaseous black holes). On the other hand, although materialists must insist that worlds containing nonmaterial elements are metaphysically impossible, they may regard such worlds as conceptually and logically possible.

There are still further types of necessity that one might distinguish. Blackburn's scheme includes what he calls "conceptual or analytic" necessity. He gives this various interpretations: e.g., it is the necessity with which a claim obtains when making it is "constitutive of competence" with the concepts or vocabulary germane to it (1985: 66); or something that obtains in every world that is possible as far as "conceptual constraints" go; or something that obtains in every "analytically possible" world (1985: 67); or that is a "logical necessity" (1984: 184). These suggest further distinct forms:

\Box_c : P (some state of affairs) necessarily obtains if and only if it obtains in every conceivable world;

\Box_a : P (some state of affairs) necessarily obtains if and only if it obtains in every analytically possible world;

Blackburn's view may be that these are equivalent notions, and indeed they might be.⁹⁹ But it seems more plausible to assume that they are not. For it seems that something might be conceptually necessary, in the sense that one could not conceive of it not obtaining, even if its obtaining is not assured by the semantic analysis of any word or concept, or entailed by any logical principles.¹⁰⁰ So, I shall assume that \Box_c means "necessary as far as our conceptual abilities are concerned," and that \Box_a denotes analytic truth (the kind revealed by semantic analysis).

⁹⁹ For example, one might argue that (a) it is inconceivable that something not obtain only if it is required by a correct analysis of the concepts one holds, and that (b) it could be required by such an analysis only if its denial involves a breach of the formal rules of, say, first order logic. Cf. Putnam 1975.

¹⁰⁰ The literature provides some convincing examples of conceptual but nonanalytic necessities. For example, Alvin Plantinga, thinks that so-called "essential" properties obtain with such necessity. According to Plantinga, Socrates bears the property of "being snubnosed in α " (where " α " = "the actual world") (1974: 63). This property is nonanalytic ("nontrivial" such as the property of being unmarried if a bachelor), but conceptually necessary because it obtains in every conceivable world (including worlds in which Socrates is not snubnosed). See also Kaplan 1989 and Putnam 1975.

Since supervenience principles state nomological necessities, they are accordingly multifarious. They can state that Gs must produce Fs as a matter of physical, metaphysical, or conceptual necessity. These differences can be expressed informally in the following way, using possible worlds language and the idea of consistency of distribution (of supervening over subvening properties). We can say that supervenience principles demand a certain consistency of distribution of effected properties over base properties, but that the stringency of this demand can vary depending on whether it holds in every logically possible world, or only every metaphysically possible world, or only every physically possible world.

Second, supervenience principles differ in logical form. Blackburn distinguishes four nomological principles as candidate forms of supervenience. One necessitates a conditional in which a universal claim is consequent upon some existential generalization (about how properties are, in fact, realized):

$$(W) \Box[(\exists x)(Fx \ \& \ Gx \ \& \ (Gx \ U \ Fx)) \rightarrow (\forall y)(Gy \rightarrow Fy)] \text{ (where "U" = "underlies")}$$

A second reformulates W by making the consequent itself a necessity claim.

$$(I) \Box[(\exists x)(Fx \ \& \ Gx \ \& \ (Gx \ U \ Fx)) \rightarrow \Box(\forall y)(Gy \rightarrow Fy)].$$

A third necessitates a universal claim simpliciter (e.g., "all Gs produce Fs"):

$$(S) \Box(\forall x)(Gx \rightarrow Fx)$$

Now, all sorts of formal objections may be raised against Blackburn's formulae,¹⁰¹ but since I am not concerned with the difficulties raised by his choice of notation, I shall set

¹⁰¹ For example, Dreier (1992) suggests that use of 'U' in Blackburn's formulae is problematic since it is naturally construed as the converse of the supervenience relation, and in that case there is no need to write out an entire formula to express weak supervenience. The fact that Blackburn states that S "characterizes the relation 'U' that holds when one (property) 'underlies' the other" (1985: 60), when S contains 'U' does appear vacuous. But it is not. Motivating the "Gx U Fx" phrase is clearly the idea that an item x bears F in virtue of the fact that it bears G, i.e., that x's G-ness realizes or produces its F-ness, and this may be logically prior to the nomological principle expressed by S. The production or

those aside and assume that the formulae capably express three interesting concepts. Expressed informally in the aforementioned idiom, W, I, and S are related like this: S requires a consistent distribution of effected properties both within and across each possible world. It says that, as a matter of (some type of) necessity, all Gs realize Fs: i.e., every G in every possible world realizes an F property. It thus affirms that, where Fs and Gs are concerned, possible worlds are both internally consistent, and consistent with each other. W, by contrast, demands only a consistent distribution of Fs over Gs within possible worlds. It affirms that, in any world w containing at least one thing that is both F and G and whose G-ness underlies its F-ness, every other G thing in w will similarly realize F.¹⁰² I says that Gs produce Fs in every possible world, provided that there exists some F-producing G thing.¹⁰³

Given the pluralistic account of necessity, and the formal differences in possible supervenience laws just introduced, we can distinguish twelve distinct supervenience principles. Arranged in order of ascending strength:

realization of one property by another does not by itself amount to supervenience since supervenience is a nomological relation but production and realization need not be.

¹⁰² Blackburn's own informal translation is imprecise. The formula, he writes, "says that as a matter of necessity, if something x is F, and G underlies this, then anything else in the physical or natural or whatever state G is F as well" (1985: 60). Blackburn clearly interprets the formula such that the satisfaction of the specified condition determines the relations between Fs and Gs only within that possible world in which the condition is satisfied, and not for Fs and Gs in every possible world. The latter concept is presumably expressed by I.

¹⁰³ Blackburn apparently thinks S implies "reductionism" (1985: 63), and that I risks reductionism (since the satisfaction of the antecedent in I allows us to derive S). But his concern about the connection between S and reductionism seems misplaced. S does not, as Blackburn thinks, represent the reduction of Fs to Gs since it does not exclude alternative subvenient bases for F, i.e., the possibility that F is multiply realizable. Note that it is the multiple realizability of effected properties that is standardly thought to preclude their reduction (see e.g., Block 1990: 146; Lepore and Loewer 1989: 180), and that it is because supervenience laws typically allow multiple realizability that supervenience is thought to be an alternative to reduction.

1. Weak physical supervenience (WPS): Fs distribute consistently across Gs within each physically possible world, i.e.,
 $\Box_p [(\exists x)(Fx \ \& \ Gx \ \& \ (Gx \ U \ Fx)) \rightarrow (\forall y)(Gy \rightarrow Fy)];$
2. Weak metaphysical supervenience (WMS): Fs distribute consistently across Gs within each metaphysically possible world, i.e.,
 $\Box_m [(\exists x)(Fx \ \& \ Gx \ \& \ (Gx \ U \ Fx)) \rightarrow (\forall y)(Gy \rightarrow Fy)];$
3. Weak conceptual supervenience (WCS): Fs distribute consistently across Gs within each conceivable world, i.e.,
 $\Box_c [(\exists x)(Fx \ \& \ Gx \ \& \ (Gx \ U \ Fx)) \rightarrow (\forall y)(Gy \rightarrow Fy)];$
4. Weak analytic supervenience (WAS): Fs distribute consistently across Gs within each analytically possible world, i.e.,
 $\Box_a [(\exists x)(Fx \ \& \ Gx \ \& \ (Gx \ U \ Fx)) \rightarrow (\forall y)(Gy \rightarrow Fy)];$
5. Intermediate physical supervenience (IPS): Fs distribute consistently across Gs within and across every physically possible world if some G produces F, i.e.,
 $\Box_p [(\exists x)(Fx \ \& \ Gx \ \& \ (Gx \ U \ Fx)) \rightarrow \Box_p(\forall y)(Gy \rightarrow Fy)];$
6. Intermediate metaphysical supervenience (IMS): Fs distribute consistently across Gs within and across every physically possible world if some G produces F, i.e.,
 $\Box_m [(\exists x)(Fx \ \& \ Gx \ \& \ (Gx \ U \ Fx)) \rightarrow \Box_m(\forall y)(Gy \rightarrow Fy)];$
7. Intermediate conceptual supervenience (ICS): Fs distribute consistently across Gs within and across every conceptually possible world if some G produces F, i.e.,
 $\Box_c [(\exists x)(Fx \ \& \ Gx \ \& \ (Gx \ U \ Fx)) \rightarrow \Box_c(\forall y)(Gy \rightarrow Fy)];$
8. Intermediate analytical supervenience (IAS): Fs distribute consistently across Gs within and across every analytically possible world if some G produces F, i.e.,
 $\Box_a [(\exists x)(Fx \ \& \ Gx \ \& \ (Gx \ U \ Fx)) \rightarrow \Box_a(\forall y)(Gy \rightarrow Fy)];$
9. Strong physical supervenience (SPS): Fs distribute consistently across Gs within and across each physically possible world, i.e., $\Box_p (\forall x)(Gx \rightarrow Fx);$
10. Strong metaphysical supervenience (SMS): Fs distribute consistently across Gs within and across each metaphysically possible world, i.e., $\Box_m (\forall x)(Gx \rightarrow Fx);$
11. Strong conceptual supervenience (SCS): Fs distribute consistently across Gs within and across each conceivable world, i.e., $\Box_c (\forall x)(Gx \rightarrow Fx);$
12. Strong analytical supervenience (SAS): Fs distribute consistently across Gs within and across each conceivable world, i.e., $\Box_a (\forall x)(Gx \rightarrow Fx);$

Finally, we can amend this Blackburn-inspired list of disambiguated supervenience laws to explicitly suit constituted properties by conjoining them with statements affirming their multiple realizability. The amendments are needed because constituted properties are not merely supervening properties. F could supervene on G even if G and F are identical.¹⁰⁴ But properties that are constituted by, though not identical to, their bases are both supervening and multiply realizable.¹⁰⁵ F is constituted by G if G necessitates F while, as Brink writes, "F could have been realized differently" (1989: 157). It follows that 1-12 are not explicitly constitution claims; as they stand they may be affirmed of both properties that are identical to their bases and those that are constituted by them. In order to stand as explicit constitution claims, Blackburn's supervenience principles need to be conjoined with claims affirming the multiple realizability of effected properties.

The multiple realizability (MR) of constituted properties is, like their supervenience, a modal feature, and hence it too can be disambiguated. We can represent MR as

$$\diamond(\exists x)(Fx \ \& \ -Gx),$$

and regard the operator ' \diamond ' as we have regarded ' \square ' in the discussion above, i.e., as denoting possibility as a generic concept, from which a number of notions might be differentiated:

¹⁰⁴ Suppose it is a logical truth that the properties F and G are identical:

$$AI: \square_c(\forall x)(Gx \leftrightarrow Fx)$$

From this we can derive the supervenience law SAS

$$\square_a(\forall x)(Gx \rightarrow Fx)$$

and for that matter any weaker supervenience claim. Note that this in no way tells against the claim that supervenience is generally compatible with the multiple realizability of effected properties. It is the identity of effected properties that precludes their multiple realizability, not the fact that they supervene upon their bases. Accordingly, it is AI which rules out the multiple realizability claim

$$\diamond_a(\exists x)(Fx \ \& \ -Gx)$$

and not SAS (or any other supervenience claim), and we cannot derive AI from SAS (or any other supervenience claim).

¹⁰⁵ Notwithstanding Brink's claim that "identity implies constitution, but not vice versa" (1989: 157).

viz., \diamond_a , \diamond_c , \diamond_m , \diamond_p . Accordingly, we can differentiate different strengths of MR claims, beginning with:

Logical Multiple Realizability (LMR): It is logically possible that something is F but not G, i.e.,
 $\diamond_l(\exists x)(Fx \ \& \ -Gx)$;

and so on so as to yield analytic, conceptual, metaphysical, and physical counterparts. We then can yield constitution claims by conjoining them with some supervenience law. For example, a scientific theory about rainbows may include the claim that certain physical states (e.g., certain arrangements of water molecules existing under certain conditions) physically necessitate rainbows, but allow that rainbows may (as far as our physical theories are concerned) be realized via some alternative base (e.g., deuterium vapour); assuming that SPS is an adequate expression of the form of that necessity claim, we may represent this symbolically by:

$$\Box_p(\forall x)(Gx \rightarrow Fx) \ \& \ \diamond_p(\exists x)(Fx \ \& \ -Gx).$$

Since MR claims and supervenience laws are logically independent,¹⁰⁶ there are no logical constraints on which (modally disambiguated) supervenience law should be conjoined with a given MR claim. Instead, such constraints are supplied by the theories which motivate constitution claims. Since affirmations of the supervenience and of the multiple realizability of effected properties are each typically made as constituent elements of a unitary account of those properties (viz., a constitution account), their modalities are normally uniform, reflecting the modal commitments of the motivating theory. Hence, just as the constituent claims in scientific constitution accounts of rainbows are modally uniform, we should expect a metaphysical constitution theory of substance to be so as well; i.e., to hold that substance is, as a matter of metaphysical necessity, realized by, say, natural

¹⁰⁶ No version of supervenience rules logically rules out any MR claim.

elements, but allow that it is metaphysically possible that substance be composed of, say, supernatural elements.

6.4: Modal Diversity and Brink

We have seen that, on Brink's view, both moral and nonmoral macrophysical properties are constituted by and strongly supervenient upon their respective bases. They are constituted in virtue of their multiple realizability, and they are strongly supervenient in virtue of the fact that they are necessitated by their bases. However, given the plurality of possible ways in which supervenience and constitution claims may be construed, the aforementioned distinctions appear to leave room for considerable modal differences between the moral and nonmoral cases. Brink's explication of the modal implications germane to his concept of strong supervenience suggests a claim formally equivalent to S:

Once the base properties are fixed, the supervening properties are thereby determined. Any world qualitatively identical to the actual world in respect of base properties must instantiate qualitatively identical supervening properties. (1989: 161)¹⁰⁷

Nonetheless, different effected properties could be constituted and strongly supervening while conforming to different S-type principles; say, SPS/PMR on the one hand, and SCS/CMR on the other. Nor does Brink's construal of necessity in general as "synthetic" indicate the modal strength of strong supervenience in any precise way. Statements affirming physical, metaphysical, and even conceptual, necessities (and their corresponding supervenience laws) may all be synthetic. His expressed construal of necessity in general as "metaphysical" is also unhelpful in favouring any particular conception of necessity.

¹⁰⁷ Brink's method in 1989 does not approach Blackburn's rigour, and so it is not entirely clear that Brink does intend an S-type claim by strong supervenience. One might read the passage quoted such that the possible worlds claims (beginning with the second sentence in the passage) explicate the idea that strongly supervening properties are "necessitated" and "determined." That interpretation suggests that I represents Brink's strong supervenience. Of course, an I-account of moral and nonmoral properties would also leave room for modal differences.

Given his contrast of necessity and analyticity (1989: 165-166), it may be taken as the claim that ascriptions of necessity designate or refer to features of the world (as opposed to language), but this does not constrain us toward any single type of necessity, since it is true of both metaphysical claims, which involve \Box_m , and scientific claims, which involve \Box_p .

This considerably complicates the quest for assimilation. If neither constitution accounts simpliciter, nor constitution/strong supervenience accounts, of effected properties do not entail any particular form of necessity or supervenience, neither is sufficient to establish that moral and nonmoral macrophysical properties are each realized by, and supervene upon, their respective bases in the same way. Pace Brink, each may be "necessitated," "constituted," and "strongly supervening" properties even if it turns out that moral property ascriptions are conceptual truths, and nonmoral macrophysical property ascriptions are just physical claims. And this seems at odds with the spirit of the AMR account. In order to assess the viability of that account, we need to go beyond the argument to the constitution/strong supervenience of moral and nonmoral macrophysical properties.

6.5: Nonmoral Macrophysical Properties

How should we construe claims of supervenience as they apply to nonmoral macrophysical properties? It seems clear that such properties supervene more weakly than SAS, or SCS implies. 'Water = H₂O' is not an analytic truth; nor does its denial violate any principles of first-order logic. It is also conceivable that H₂O realize macrophysical properties other than those that it in fact realizes.

This leaves versions 1 through 10 as candidate principles.¹⁰⁸ Let us consider the strongest of these: SMS. SMS affirms an \Box_m -type necessity. Applied to our example, it

¹⁰⁸ Note that the fact that it is not a logical or analytic truth that water is H₂O, while counting against the application of strong analytic and logical supervenience to ordinary nonmoral macrophysical properties, does not rule out their supervening in the sense of

states that it is metaphysically necessary that H₂O realize the macrophysical properties associated with water; that there are no metaphysically possible worlds in which H₂O does not realize those properties. Are there "metaphysically possible" worlds in which H₂O does not realize the macrophysical properties associated with water? Since what is metaphysically possible depends on what metaphysical theories one accepts, and since metaphysical claims (contrasted with logical claims) are relatively contestable, just what is metaphysically possible is not obvious.

Nonetheless, microphysical states such as H₂O do not appear to metaphysically necessitate their corresponding macrophysical properties. Presumably, what macrophysical properties a microphysical base realizes depends on what specific physical laws are actually in effect. But it would be odd for a metaphysical theory to say that a slightly different set of physical laws is impossible. Metaphysical theories issue claims about what kinds of things can exist rather than claims about the specific macrophysical properties any micro-base must realize. Consider as examples the metaphysical perspectives Brink mentions in his book. Naturalism rules out "nonnatural" (sui generis) and supernatural properties; Moorean "nonnaturalism" admits natural and "nonnatural" properties, but not supernatural properties; and supernaturalism either allows them all, or insists that only supernatural items exist. These theories, at most, suggest limits on what sorts of different properties should be allowed in a metaphysically correct description of the world. But none rules out the possibility that H₂O realize macroproperties other than the

weak analytic or logical supervenience. Ordinarily the denial that it is, for example, a logical truth that G is F, implies

$$\diamond_1 (\exists x)(Gx \ \& \ -Fx)$$

which affirms that in some logically possible world there exists a G thing which is not F. This is quite compatible with the claim made by weak logical supervenience

$$\Box_1 [(\exists x)(Fx \ \& \ Gx \ \& \ (Gx \ U \ Fx)) \rightarrow (\forall y)(Gy \rightarrow Fy)]$$

which just affirms that there is no logically possible world in which some Gs realize Fs while others do not.

ones that are, in fact, effected by it; hence none makes it metaphysically necessary that H₂O realize those particular properties. Since SMS, applied to this example, would state a nomological claim that involves this type of necessity, it follows that SMS is not the form of supervenience appropriate to such properties.¹⁰⁹

By contrast, SPS fares much better as an account of the nomological relation between macrophysical properties and their bases. Assignments of macrophysical properties are typically scientific, and hence physical, claims: they are typically conceived as holding in every physically possible world, where what is physically possible is itself determined by the scientific theories we accept. Since our best theories affirm that H₂O realizes "water" properties, and not, say, "gasoline" properties, we regard the latter as physically impossible, although not metaphysically, conceptually, and logically impossible.¹¹⁰ Since SPS states a nomological principle that holds with this type of necessity, it seems the appropriate concept to use in expressing the sense in which ordinary nonmoral macrophysical supervening properties distribute across subvening properties.

This view of the role of scientific theory permits us to rule out the intermediate and weak formulations of supervenience. IAS and ICS make it (respectively) analytically and conceptually necessary that H₂O realize water properties provided that some instances of H₂O actually do realize those properties; but our physical theories do not state such necessities even given that existential assumption. IMS makes it metaphysically necessary on that assumption, but our physical theories do not. As far as our physical theories are

¹⁰⁹ This is so notwithstanding Brink's claim that necessity simpliciter is a "metaphysical" notion (1989: 165).

¹¹⁰ This is not to say that scientific theories, or perhaps more accurately, the scientific view of the world, contains no implicit or explicit metaphysical commitments. On the contrary, it is natural and appropriate to associate the scientific view with materialism, given the history of science. But such metaphysical commitments, if they are there, do not rule out alternative supervening properties for H₂O such as gasoline properties.

concerned, even if H₂O actually realizes water, it is not logically, analytically, conceptually or metaphysically necessary that it do so. IPS obtains in the case of H₂O, but only given SPS: it is strictly in virtue of the fact that H₂O produces water properties in every physically possible world that it does so on the assumption that H₂O actually produces water.¹¹¹ IPS does not, therefore, express the nomological principle appropriate for such cases. Likewise with the weak concepts WPS, WMS, and WCS.¹¹² Our water property ascriptions are held to obtain in every physically possible world, but none of the weak concepts require consistency of distribution across worlds. At most, they require consistency within worlds - they issue bans on "mixed worlds." These bans come in various strengths. The strongest, issued by WAS, bans all mixed worlds. According to it, a world in which some Gs realize Fs, and some do not, is analytically impossible, and WCS makes it inconceivable. WMS rules out all metaphysically possible mixed worlds. It demands consistency within every metaphysically possible world. However, since it makes this demand on metaphysical, and not conceptual or logical, grounds, it allows

¹¹¹ Many of the concepts on the Blackburnian list logically imply weaker counterparts, e.g., SCS implies SMS, SPS, ICS, IMS, IPS, etc.

¹¹² The injunction to focus on strong supervenience in critiquing Brink may make the arguments against "weak" conceptions of supervenience seem otiose. But given Brink's lack of rigour, we cannot be confident that his weak and strong concepts are equivalent to Blackburn's respective generic laws W and S, and therefore that they specify the formal features germane to the "weak" and "strong" laws in the Blackburnian list. Brink might have intended I not S by strong supervenience (see note 18). As well, it is not entirely clear that Brink's weak supervenience prohibits mixed worlds (as do the W-type laws). Notice that Brink (1989: 161) refers to "circumstances" rather than "worlds;" perhaps then possible worlds may contain Gs in varying circumstances, some producing Fs, and some not. Notice also that Brink conceives weak supervenience as accommodating the idea, perhaps germane to our moral thinking, that certain social and economic conditions might realize injustice in some "possible" societies but not in others (1989: 161). Presumably he has in mind judgments such as the claim that, say, class oppression is not unjust when it is made inevitable by conditions of scarcity, so that it would be incorrect to say that feudal societies are unjust in virtue of that oppression. But if Brink's weak supervenience is equivalent to W and does ban mixed worlds, then it presumably precludes such judgments (at least when they are made about different societies within the same possible world).

mixed worlds as logical or analytical or conceptual possibilities. WPS, the weakest concept of all, proscribes only physically possible mixed worlds; it allows them as metaphysical, and conceptual or logical possibilities. None of these concepts sufficiently expresses the idea that our best scientific theories rule out as physical impossibilities the distributions of supervening properties incompatible with those theories.

SPS does not, however, represent the modal form of every scientific supervenience claim. Claims germane to a scientific theory that is still in its nascent stages of development, or more speculative theoretical claims, might be held with relative caution, and thus might not typically be regarded as ruling out competing postulations as (physically) "impossible." As Brink suggests, those cases involve a weaker nomological relation: one that does not make the production of the effected property physically necessary. Suppose, for instance, that scientists do not know how chemical *x* will behave at certain temperature *t*. They could not assert the SPS claim that, as a matter of physical necessity, anything that is a sample of *x* will boil at *t*. Nonetheless, presumably, it could still be claimed that whatever macroproperties chemical *x* does possess, it will possess as a matter of some sort of necessity, and in particular that if chemical *x* did possess the macroproperty "will boil at *t*," then it would possess that property as a matter of some sort of necessity. In short, even if they do not know that chemical *x* boils at *t*, they could still affirm the existence of a weaker, conditional, supervenience relation between the property of being chemical *x* and the property of boiling at *t*. Which principle is a suitable expression of the judgments we make under these circumstances? Scientific inquiry into property realization is restricted to the physical modal dimension. It inquires into the distribution of effected properties in worlds that are determined to be physically possible by our best physical theories (it has no interest in property realizations that are conceivable, logically possible, or metaphysically possible, but ruled out by our physical theories). This

leaves WPS and IPS as possible forms of epistemically-compromised supervenience. Of these IPS is the more plausible candidate. WPS is simply too weak. It requires only uniformity of property production within physically possible worlds. But even when we do not know what macrophysical properties a given kind might bear, the supervenience claims we make about that kind will involve a stronger necessity than this. Having no evidence for the claim that a substance bears a particular macroproperty warrants abandoning the claim that the macroproperty is (physically) necessitated by the specific base, but it does not warrant a rejection of the idea that, whatever macroproperty is produced by a specific base, is produced by that base as a matter of physical necessity. IPS states precisely that, as a matter of physical necessity, if a particular base produces a property, then it must do so in every physically possible world.

6.6: Moral Properties

If the above argument is correct, the nonmoral properties referred to by natural kind terms like 'water' supervene on their *microphysical* bases in the sense of what I have here called "strong physical supervenience." That principle demands a consistent distribution of properties in every world that is not ruled out by our best scientific theories. It prohibits both mixed physical worlds, and physical worlds whose distributions of supervening properties are inconsistent with each other.

SPS is not, however, the most plausible candidate where moral properties are concerned. A good place to begin our inquiry into the nature of moral supervenience is by turning again to Blackburn. We have seen that nonmoral macrophysical property ascriptions are not logical or analytic claims (6.5). Blackburn points out that the same is true of moral property ascriptions. He states that although we may be constrained, a priori, to assess naturalistically similar things similarly, no such constraints direct us, ab initio, to assess things in any particular way:

In particular in the moral case it seems conceptually or logically necessary that if two things share a total basis of natural properties, then they have the same moral properties. But it does not seem a matter of conceptual or logical necessity that any given total natural state of a thing gives it some particular moral quality. (Blackburn 1984: 184)

(Blackburn apparently equates "analytical" and "conceptual" necessity, so his point may be taken as the denial that anything assumes a moral quality as a matter of logical or analytical necessity.) This point is plausible at a more theoretical level too: it does not seem a matter of logical or analytical necessity that any particular natural or social property carry the evaluative significance that it happens to carry in our moral thinking. If, for example, pleasantness is a good-making property, it does not seem so on logical or analytical grounds. Given this, the application of moral predicates to subjects referring to various candidate bases could not be "necessitated" analytically ab initio. (Moral properties could only be necessitated by their bases ab initio if we could settle on an analysis of moral property-ascribing terms, including their extensions.) Paradoxically, moral supervenience must nonetheless involve the claim that, once fixed by their respective bases, moral properties must distribute in ways consistent with their ab initio distributions. On these grounds, Blackburn suggests his generic W as the type of supervenience involved in moral property ascription. (Given his failure to distinguish analytical and conceptual necessity, W presumably amounts to WAS/WCS). And he goes on to argue that W poses a problem for moral realism. W bans mixed worlds; the best explanation of that ban, in light of the lack of any formal or analytic constraints, favours metaethical projectivism. Naturalistically-similar objects of moral assessment just provoke similar "projective" responses (Blackburn 1971; 1984: chapter 6; 1985).

Now, I think that Blackburn is correct in thinking that moral supervenience lends itself to an antirealist explanation. I also think that he is correct in assuming that moral property ascriptions are not analytic claims, and that this fact bears on how such ascriptions

should be formally construed. However, it seems to me that, as things stand, AMR theorists do have an answer to Blackburn's worry about mixed worlds (albeit one whose viability remains to be assessed). Stated succinctly, Blackburn's argument looks like this:

1. Moral property ascriptions are not analytic.
2. If moral property ascriptions are not analytic, they must be W-type claims.
3. If they are W-type claims, then they proscribe against "mixed worlds."
4. The best explanation of the proscription against mixed worlds, given the lack of analytic or logical proscriptions, is that "projectivism" is true.
- ∴ 5. "Projectivism" is true.

The main problem here is that (2) is false. Moral property ascriptions might be nonanalytic without being W-type claims. They could be IPS, IMS, ICS, SPS, SMS, or SCS claims, for none of these principles represent claims of analytic necessity. Notice also that Blackburn's rationale for retreating to the conditional form of supervenience represented by W is, like the one that I offered for retreating to IPS in the scientific case under certain circumstances, ostensibly epistemic. His idea is apparently that, since the extension of moral property-ascribing terms is not discovered via an analysis of the meaning of those terms, the strongest claim we can make in the moral case is: "if a set of base properties produces a certain moral property, then any similar basket of basic properties must also produce that moral property." But this conclusion only follows given the further supposition that there is no other available means of determining the extensions of moral terms. And this further assumption is obviously one that no AMR theorist would accept, for the AMR approach is marked by the conviction that there can be moral knowledge even without the availability of extension-fixing analyses. For instance, Brink clearly thinks that moral claims can be epistemically justified in much the same way as are scientific claims: he thinks, for instance, that there can be observational evidence for moral claims, that such claims can be tested empirically, and that they can pass such tests (1989: 168, 182-197).

But if there can be moral knowledge of this sort, then we need not retreat to the conditional account of moral supervenience claims, and the worry about mixed worlds imposed by Blackburn's account can be avoided.

This response to the mixed worlds objection assumes, then, that moral claims need not be conditional in form, that we are sometimes justified in making "unconditional" moral supervenience claims. However, this assumption obviously leaves a plurality of candidate accounts of moral supervenience claims (i.e., SPS, SCS, SMS accounts, etc.), and it remains to be seen whether whatever account turns out to be correct is amenable to the particular brand of moral realism advocated by the AMR theorist. In the remainder of this section, I examine the nature of moral supervenience claims in more detail, and in section 6.8, I discuss whether my view of moral supervenience claims is compatible with AMR.

Brink never insists that there must be exact modal parity between moral and other forms of supervenience claims; nor is it obvious that the mere absence of such parity, if it obtained, would be incompatible with AMR. Nonetheless, AMR would clearly benefit from the presence of modal parity, for if parity were established, then there could be no special problem about the moral case rooted in a plausible account of moral supervenience claims. By contrast, the absence of parity admits the possibility of interpreting moral supervenience claims in a way that is incompatible with AMR. For this reason, it would be unfortunate for the AMR theorist if SPS did not emerge as the proper form of moral supervenience.

In fact, SPS is not a viable account of moral property ascriptions. Consider our assessments of personal character: in particular the judgment that an individual is bad. Underlying the idea that there are clear instances of moral badness is the notion that although the extension of the term 'bad', as applied to persons, may not have been precisely

determined, it is nonetheless true that certain combinations of base properties are undeniably bad. Consider as an example the protagonist in Shelley's The Cenci, who apparently exhibited all of the hallmark traits and dispositions associated with being a bad person: a father who, in addition to committing many other atrocities, over many years rapes and tortures his wife and children, until finally one of his daughters conspires with his wife to murder him so that they both might escape the abuse. Here, one wants to say, is a clear example of a bad person. Now, in looking at the nonmoral case, we saw that an inquirer who accepts the theory that H₂O must realize a particular macroproperty or set of macroproperties will make an SPS supervenience claim - just because their theory, as a physical theory, rules out alternative physical supervening properties for H₂O, but does not itself impose any metaphysical, conceptual, or logical constraints on what macroproperty H₂O realizes (it admits as metaphysical, conceptual, and logical possibilities that H₂O realize alternative properties). In the moral case, the situation is different. The inquirer who supposes that Count Cenci is a bad person in light of what he has done should, I think, find it inconceivable that the Count could commit the crimes that he committed, and possess the traits of character that he possessed, but not be a bad person. Of course, I am not claiming here that there really is a moral fact out there, i.e., that Cenci really is bad in light of his character and behaviour, and that this fact obtains as a matter of conceptual necessity. I am suggesting, rather, that the person who claims that the Count is bad in light of his character and behaviour should be interpreted as making this sort of modal claim; that the correct analysis of this and other moral property ascriptions is that they are (at least implicit) SCS claims. Two pieces of evidence support this thesis. First, we ordinarily interpret the moral judgments of others as SCS claims. To appreciate this, consider the stance of a person who claims that the Count is bad in light of his personality and behaviour yet who insists that it is conceivable that all of the basic facts or properties that

are relevant to the assessment of Cenci, and that point to his evilness, might have existed, just as they are, without making him evil. I expect that most would find this stance odd, even (internally) incoherent - akin to that of a scientist who affirms that certain conditions necessitate that a substance exhibit a particular macrophysical effect, but who also insists that it is physically possible for the substance to exist under exactly those conditions without exhibiting that effect. Now, a natural and plausible explanation of the apparent incoherence of the scientist's stance is that we assumed that his initial claim conformed to SPS and that, as such, it ruled out his possibility affirmation. Likewise with the aforementioned moral stance: a plausible explanation of the apparent incoherence of the moral stance is that we have interpreted the initial claim as an SCS ascription, and as such it rules out the associated conceivability claim. Secondly, our own judgments seem to us to have this form. To appreciate this, reflect on the matter yourself: do the Count's character and actions make him a bad person? I submit that anyone who agrees that they do will also find it inconceivable that he could have had exactly that character, behaved in exactly that way, under exactly those circumstances, and not have been a bad person. All of this suggests to me that this and similar moral property ascriptions are (at least implicitly) SCS claims, and not SPS claims.

This construal of moral supervenience claims invites objections. First, it might be supposed that making an SCS claim requires assuming a particular view of how much, and what kind of, mutation individuals can sustain without losing their identity; in short, that it demands assuming a position on the question of transworld identity. But it may be thought that since the problem of transworld identity is a difficult, perhaps intractable, metaphysical problem, about whose solution there is no consensus, we can hardly regard our ordinary moral property ascriptions as implicit SCS claims. Second, counterexamples might be proposed. For example, suppose a person located the source of Cenci's badness

in the misery he creates for others, and in the disregard and even hostility he expresses toward their welfare. Couldn't such a person nonetheless conceive of a world in which, say, misery is viewed as a good thing,¹¹³ or a world in which no one can be made better or worse off (where people do not have a welfare in the relevant respect); and wouldn't it seem to such a person that Cenci is morally unremarkable, or even good, in those worlds? I believe that these objections rest on a failure to recognize the assumptions that govern moral supervenience claims.

Suppose that

- (i) α is the actual world,
- (ii) G refers to the set of bad-making natural, social, and psychological properties $\{N_1...N_n\}$ that produce F in α (and perhaps elsewhere), and
- (iii) F refers to the moral property of badness which, putatively, supervenes on $\{N_1...N_n\}$;

The following assumptions then govern moral supervenience judgments.

First, we assume that

1. G includes all of the factors that are relevant in α to the assessment of Count Cenci (including his intentions, motives, dispositions, and actions).

¹¹³ Of course, if the claim at issue is correct (viz., that moral property ascriptions are conceptual claims) then anyone who thinks that misery is bad (perhaps most people) may at this point be tempted to object that there are no conceivable worlds in which misery is regarded as a good thing. But we can and should distinguish between the (moral) judgment that (a) "misery is a good thing" and the (sociological or psychological) claim that (b) "misery is viewed as a good thing." I believe that the temptation to regard (b) as inconceivable is explained by a simple fact about our moral psychology: many moral judgments seem so obviously true that we not only find it impossible to conceive of its not obtaining, but we also sometimes have difficulty imagining that anyone else might judge differently. But in fact, people often do make moral judgments that are at odds with our most confidently held moral assumptions. And (b) is no different: some Christians (e.g., Huxley 1946) seem to suggest that suffering is at least extrinsically good (that more value is realized in a world in which individuals suffer and struggle than one in which everyone is happy). It seems perfectly conceivable that there might be worlds in which everyone (or at least nearly everyone) agrees with Huxley about this.

Otherwise (if there were factors that are relevant but excluded), we could not say that G is the complete basis for the moral property correctly assigned to Cenci; instead, it would be G and the other factors that would form that basis.¹¹⁴

Second, since we are inquiring into the strength of the necessity with which G (not some other base) produces F, we assume that

2. G in every world contains the same properties as G in α ;

Finally, since it is we in α issuing the moral claim, employing our language, concepts and moral principles, whether something is "conceivable" depends entirely on what we, in describing α , are able to conceive. So we assume that

3. there are no worlds containing G ("G-worlds") in which agents can correctly answer conceptual questions (e.g., "are an agent's intentions ever relevant to the assessment of her character?") in ways that are incompatible with the answers that are correct in α .¹¹⁵

Thus, suppose that we, in describing α , are unable to conceive of a world in which G things are not F. Such worlds would be eo ipso inconceivable. Nonetheless, we might still be able to conceive of a world populated by cognizers whose conceptual abilities allow them to conceive of a world in which G things are not F. The point here is that even if we could conceive of this second sort of world, this would not make the first sort of world conceivable.

The present suggestion is that ordinary moral property ascriptions of the form "G things are F" imply this account of subvening bases, and also imply that G things are F as a matter of conceptual necessity. We can also meet the aforementioned objections in the following way:

(a) G-worlds are, given 2, those that contain a Cenci whose behaviour and psychological bearing are the same as in α . The identity of the individual (i.e., the putative

¹¹⁴ See also Blackburn 1985: 61-62.

¹¹⁵ See Kripke 1980.

bearer of the moral property F) in those worlds is thus assumed when we make moral supervenience judgments; hence the transworld identity of individuals is assumed to be possible. We do not have to settle the metaphysical problem of transworld identity in order to accept this view of moral property ascriptions. Moreover, it is worth noting that since many commonsense and philosophical positions do assume transworld identity, any objection to specific positions on the modality of moral property ascriptions that is rooted in doubts about transworld identity could be accepted only at great metaphysical cost.¹¹⁶

(b) G-worlds are also those in which individuals have their welfare violated, and in which persons are made miserable and unhappy. Since G in α contains a Cenci who violates the welfare of others and makes them miserable (by raping and torturing them, etc.), G in other worlds contains those elements as well. Hence the moral ascription implies that there are no conceivable G-worlds in which individuals do not possess a welfare, or cannot be made unhappy, and no G-worlds in which, for those reasons, F is not produced. To this, there is an important objection that needs to be addressed here. Many of the properties relevant to assessing Cenci may plausibly be construed as nonmoral natural properties because their presence is determined on the basis of conceptual criteria that are themselves evaluatively neutral. The criteria can be considered neutral because although making such judgments may often, or even typically, occur in concert with some evaluative stance toward the item(s) in question, an evaluative stance is not essential to making such judgments. For instance, a rough, but plausible, definition of torture is: the infliction of severe physical pain or emotional suffering on a person against her will. But torture is not by definition bad (it is likewise for most of the other relevant properties, such

¹¹⁶ One might plausibly maintain that, even if such judgments do not assume the persistence of identical agents, this has no bearing on the assignment of moral properties since those assignments are normally thought to be universalizable, and universalizability entails essentially that the identity of individual agents is irrelevant to what moral properties they bear (Rabinowicz 1979: 11).

as 'makes his daughter unhappy' and 'makes his daughter miserable'). However, the property of inhibiting another person's welfare seems partly evaluative: whether one agrees that Cenci's actions inhibit his daughter's "welfare" seems to depend in part on what psychological and bodily states one regards as valuable. But if G contains such properties it cannot be considered a purely natural base. Hence, even if it is in all other respects sound, the argument would establish merely that moral property ascriptions imply that there are conceptual connections between (certain) normative or evaluative concepts (viz., between moral badness and inhibiting welfare). Now, it is true that moral judgments are often justified by appeal to certain normative concepts; thus, the judgment that Cenci is a bad person seems warranted (in part) in light of the way he affects the welfare of those around him. But it does not follow that in making such judgments we are building the evaluative properties into the subvening base G. For the idea that claims about welfare can assume a justificatory role can be accommodated by an account that makes the property 'violates his daughter's welfare' a supervening moral effect. The present suggestion is just that once a complete description of all of the aforementioned morally relevant nonmoral properties (i.e., 'makes his daughter miserable', etc.) are incorporated into G, Cenci's badness will appear to follow as a matter of conceptual necessity. This is perfectly compatible with the claim that Cenci is held to be bad largely in virtue of certain partly evaluative properties such as 'violates his daughter's welfare'. For it may be claimed that the partly evaluative properties themselves are conceptually necessitated by the (nonmorally construed) G-base, and that there are conceptual (or perhaps even stronger) connections between the partly evaluative property and the moral property F. Thus the judgment that Cenci is bad in light of his character and behaviour implies that there is no conceivable world in which Cenci rapes and tortures his daughter and wife, etc., where he does not violate their welfare, and none in which he is not, in part thereby, a bad person.

(c) G-worlds are assumed not to be worlds in which the morally relevant emotions, attitudes, and beliefs of the agents included in G (i.e., Cenci, the individuals affected by his actions) are different from those that comprise G in α . For example, it is assumed that there is no G-world in which Cenci's daughter prefers or desires to be abused by her father, and no G-world in which, for that reason, G does not produce F.

(d) Assumption 3 rules out the following sort of objection. The judgment that Cenci is bad is based on the concepts employed by assessors at α (such as the conviction that the intentions that determine a person's behaviour are relevant to the assessment of that person's moral character). But could not a person conceive of a world w where those concepts do not obtain (e.g., where intentions are not relevant to the assessment of agents), and thus in which G might obtain while F does not (since the factors that comprise G may not be relevant to the assessment of Cenci (or anyone else) in that world)? On these grounds, one might claim that the ascription of badness to Cenci does not imply that he is bad as a matter of conceptual necessity. However, given 3, worlds in which conditions of moral relevance change are ruled out ex hypothesi by the same conceptual constraints that determine the modality of the judgment. Of course, those constraints do not rule out G-worlds containing moral assessors who employ concepts that are at odds with those employed by the assessor who makes the judgment in α . For example, they do not rule out G-worlds containing assessors who believe that consequences are irrelevant (or, for that matter, assessors for whom happiness is conceptually tied to pain, or for whom misery is good); they merely rule out worlds in which those assessors can correctly make those conceptual judgments. When one affirms that Cenci is bad in light of his raping and abusing, one affirms that he is bad in light of that behaviour even if, counterfactually, most people came to believe that the factors one takes to be bad-making are in fact irrelevant, or to make other "arcane" conceptual judgments.

This is subject to challenge. The position depends on the viability of a distinction between (a) the object of assessment, in this case Cenci and the "base" properties he bears, and (b) the actual and counterfactual "circumstances" whose adjustment makes possible Cenci's travel through different possible worlds, such that only (a) contains factors relevant to Cenci's assessment. On this account, Cenci and all the factors relevant to his evaluation are together held to be the proper producer of moral properties, and their travel through different possible worlds is determined exclusively by adjustments to morally irrelevant factors. One can see how this view of the matter makes SCS a plausible form of moral supervenience claims. Only SCS requires a consistent evaluative response to Cenci and $\{N_1...N_n\}$ regardless of any conceivable adjustments to the circumstances in which Cenci goes about his business. Any weaker concept allows that such adjustments might (conceivably) merit a different response. However, it might be objected that this account presents a fundamentally flawed picture of our moral thinking. We do not, as the picture assumes, regard both Cenci and $\{N_1...N_n\}$ as together constituting a single proper object of evaluation. Rather, we assess Cenci, in light of factors that are pertinent to his assessment, but are nonetheless not themselves objects of that assessment. For example, Cenci's recent behaviour toward others is no doubt relevant to the assessment of his character, but how we regard his behaviour can depend on a wide variety of factors, such as Cenci's motives and intentions, his past experiences and upbringing, the circumstances in which his morally significant actions occurred, etc. While all of these may be relevant to our assessment of Cenci, they are not themselves the objects of that assessment. Rather, only Cenci is. This suggests a different picture: Cenci travels alone through possible worlds, and his transit from world to world is determined by shifts in circumstances, some of which are relevant to, and hence could warrant revising, how we morally regard him. Since most of the morally relevant shifts in circumstance that normally come to mind in cases of this sort are

conceivable, the picture precludes SCS, for SCS requires a consistent evaluative response regardless of any conceivable shifts in the circumstances surrounding Cenci.

My critique of AMR depends on the first picture of moral supervenience being the correct one.¹¹⁷ And in fact that picture does provide a more adequate account of how we ascribe moral properties. The second picture mistakenly infers from the fact that we regard Cenci as the "bearer" of moral properties, and the fact that we regard 'Cenci' as the subject to which our moral property-ascribing predicates may be attached, that making a supervenience judgment about Cenci involves identifying Cenci (alone) as the subvening producer of moral properties. But one normally includes as part of the subvenient "base" for a moral property all the factors that one takes to be "relevant" to such moral judgments. Moral properties are not normally held to supervene upon individuals, they are held to supervene upon complex constellations of natural, social, and psychological properties (Brink 1989: 161). Such judgments affirm that all other occurrences of a given constellation will realize the same moral property or properties (not "all other occurrences of Cenci"). We can and do regard such complex constellations of subvening properties as the producers of moral value, while regarding individuals as the proper bearers of that value. The objection to SCS based on a rejection of this view of supervenience fails.

This concludes my argument for, and defence of, an SCS construal of moral property ascriptions. If my case is sound, then there is no exact modal parity between moral and nonmoral macrophysical property ascriptions. Of course, the absence of parity

¹¹⁷ It is worth noting that the second view rules out both SCS and SPS as appropriate forms of moral supervenience. SPS requires a consistent evaluative response regardless of any physically possible changes in surrounding circumstances. This too is ruled out by an approach that regards Cenci as traveling alone through worlds, where his transit from world to world is potentially determined by adjustments to morally relevant circumstances. If some of the morally relevant adjustments are physically possible, it follows that some physically possible changes in circumstance could warrant a revision of moral judgment, which precludes SPS. Consequently, even if the second view is accepted, the desired assimilation fails.

does not itself entail that moral supervenience claims must be interpreted antirealistically, but it does not rule out modally-inspired doubts about a realistic construal of moral supervenience in the way that exact parity clearly would. For instance, SCS claims may turn out not to lend themselves to a realistic interpretation, or if they do, they may do so only given certain further conditions that do not obtain in the moral case. I will examine these issues in Section 6.8. However, before doing so I will address one further objection to the argument as it stands, viz., that there may be no modal asymmetry between moral and psychofunctional supervenience claims, and thus that, notwithstanding the asymmetry mentioned here, there is at least one ostensibly-realist class of claims with which moral ascriptions do have exact modal parity.

6.7: The Psychological Functionalist Option

It might be objected that even given the SPS/SCS asymmetry described here, the AMR theorist has another option. She can try to assimilate moral properties to mental states as they are construed by functionalist theories in the philosophy of psychology.¹¹⁸ There are, of course, many versions of functionalism; the leading idea is that mental states are construable in terms of their causal connections (to environmental inputs, behaviour, and other mental states similarly construed).¹¹⁹ Two features of functionalism make it seem a viable model for the AMR theorist. First, it is arguably realist about mental states (mental states need to be real in order to assume their causal roles).¹²⁰ Second, although functionalists are divided about whether folk psychological terminology is capable of fully

¹¹⁸ This strategy is, in fact, suggested by Brink explicitly in 1984: 121-122; implicitly by his use of examples in 1989: 158, 161.

¹¹⁹ See Block 1980: 171-184.

¹²⁰ This distinguishes functionalism from behaviourism. See, e.g., Armstrong 1970: 72.

describing the complex causal roles that form mental states, functionalism in general appears to affirm a very strong necessary connection between mental state types and those roles. Thus, for analytic functionalists, it is in virtue of the meaning of ordinary mental property-ascribing terms that they specify certain functional roles, while for psychofunctionalists, the nature of the functional states is fully describable only via the terms of a mature scientific psychology.¹²¹ But whether they are "folk" or "mature," it seems that on such accounts the mental ascriptions are conceptually necessitated by their functional roles.

However, functionalist accounts generally regard psychological supervenience laws as governing the relation between physical states and (functionally construed) mental properties, not functional states and mental properties. Indeed, perhaps the main attraction of the functionalist approach is its apparent ability to explain mental/physical supervenience, i.e., to explain how mental states could be necessitated by a given physical state but not (given multiple realizability) be identical or reducible to any physical state. Functionalists explain the apparent supervenience of the mental on the physical as resulting from the functional nature of mental states.¹²² Since it is easy to understand how the same functional (e.g., computational) state of a thing could be realized by different physical states, the assumption that mental states are functional states seems to explain how mental states could be realized by multifarious physical or neurological events. But this sort of account clearly locates the supervenience relation as holding between (functionally

¹²¹ The distinction originally appears in Block 1978; see also Shoemaker 1984: 272; Horgan and Timmons 1992: 234-240.

¹²² Functional theories do differ widely in how they construe the functional nature of mental states. A central difference is that between so-called "functional state identity theorists" (e.g., Putnam 1967) who regard mental states as identical to certain causal complexes or roles, and so-called "functional specifiers" (e.g., Armstrong 1970) who would identify a mental state as being whatever physical thing assumes a given causal role. We need not worry about these differences here.

construed) mental states and physical states: the idea is that mental states can be viewed as supervening multiply-realizable items because they are supervening multiply-realizable functional states. This is important because, as far as I can tell, no functionalist account makes it conceptually or analytically necessary that any particular physical state of a thing assume a specific causal role. On such accounts, it is perfectly conceivable, and analytically possible, that a brain state that in fact realizes some functional role, could nonetheless not realize it.¹²³ Indeed, the contrary view would appear implausible. Whether a particular physical brain state in fact assumes a particular functional role would appear to be a scientific question, and to depend on what physical laws actually obtain. A nonmental example may make this point clearer. It may seem inconceivable that a thing could be "what mixes gas and air and sends the mixture to something else, which in turn" and not be a carburetor.¹²⁴ But (notwithstanding the claims car vendors might make) it does not seem conceptually or analytically necessary that any specific physical item or type actually carburet! Moreover, suppose that, inspired by advances in automobile technology, the term "carburetor" is rejected by certain mechanics in favour of a term germane to a more "mature," empirically adequate, engine theory - say, "gas/air micromodulator." It might be inconceivable to these sophisticated mechanical functionalists that something could assume the function identified by that term and not be a "gas/air micromodulator," but one would

¹²³ Indeed, it is sometimes asserted that the same physical state type can realize (in the sense of physical possibility) different functional states at different times, or in different creatures (see, e.g., Block and Fodor 1972: 163). This suggests that an even weaker law than SPS governs the relation between physical and functional states (WPS perhaps?), which would make the asymmetry with the moral case deeper here than it is in the case of nonmental macrophysical properties. But Shoemaker (1984: 265-266) has noted that whether alternative realizations are possible depends on how broadly the physical base is construed.

¹²⁴ The carburetor example is from Block 1980: 180.

nonetheless expect such persons to find it conceivable that any physical machine or machine type could fail to "micromodulate."¹²⁵

Second, even if some versions of functionalism do regard mental states as "supervening" on functional states, I doubt any functionalist could tenably hold there to be exact modal parity with the moral case. Exact modal parity with the moral case would imply that there is a class of functional states any member of which would conceptually necessitate a certain mental state (e.g., being in pain): i.e., that just as it would seem inconceivable to inquirers inclined to condemn Cenci in light of his character and actions that he could have done what he did, and so forth, but not have been a bad person, it would seem inconceivable to inquirers inclined to say that pain is just a certain sort of functional state that an organism (or, for that matter, a machine) could have been in that functional state without being in pain. Now, this claim may seem tenable given some varieties of functionalism. For instance, analytic functionalism (see above) holds that our ordinary mental property-ascribing terms function as semantic constraints on the sort of functional complexes that could comprise a mental state.¹²⁶ It insists that any adequate functional account of a given mental state must be compatible with our ordinary conceptions of the state and with our ordinary uses of mental property-ascribing terms. It may be thought that, on such an account, it would seem inconceivable to users of certain mental terms that something could assume a functional role and not be the proper bearer of a given mental property. The problem, however, is that even if our uses of ordinary mental property-ascriptive terms do function as constraints on the specific functional accounts of mental states we can accept, it does not follow that any particular functional state must appear to

¹²⁵ Of course, it may turn out to be physically necessary that certain items carburet and micromodulate, though it would not be metaphysically necessary, for the same reasons raised above in relation to water property ascriptions.

¹²⁶ Horgan and Timmons 1992: 237.

the users of a given ascriptive mental term to conceptually necessitate the mental state affirmed by the term. The reason is this: functional/mental states are themselves commonly thought to be comprised partly by relations among causally-efficacious internal physical (neurological) states. But since we are at present far from having an adequate understanding of the internal causal complexes that, on functional accounts, partly comprise our mental states, it could hardly seem analytically true to anyone that a given causal complex realize a particular mental state. Nor, for the same reason, could one expect it to be inconceivable to ordinary inquirers that the particular functional state or class of states held by our best (but still inadequate) psychological and neurological theories to realize a particular mental state could nonetheless not realize it. Moreover, even if inquirers did come to acquire complete knowledge of the functional state underlying, say, pain, I doubt they would come to see it as inconceivable that that functional state exist without realizing pain. For knowledge of the internal causal processes that might underlie being in pain would presumably be gained empirically, and would be largely comprised by a theory of neurological processes and their causal interactions. And one would not expect the holders of such a theory to regard it as inconceivable that their theory is false.

All of this suggests that an SPS/SCS asymmetry with respect to the moral case emerges here as well, and that there is no (plausible) functionalist account of mental state ascriptions to which ordinary moral property ascriptions may be modally assimilated. Of course, I cannot hope to have settled the issue here with these brief and general remarks. There may still be a tenable version of functionalism that is modally on par with the moral case that I have overlooked. But hopefully I have said enough to place doubt on the prospects for an exact modal analogy between mental and moral property ascriptions.

6.8: Why is the Absence of Modal Parity Important?

The arguments of 6.5 and 6.6, if correct, indicate that there is no exact modal parity between moral and nonmoral macrophysical supervenience. In 6.1, I suggested that a lack of parity in this regard might preclude the assimilation that the AMR theorist desires. Yet, even if my case for the aforementioned asymmetry is sound, one might wonder whether its existence is really incompatible with AMR. In fact, there are a number of grounds to suppose that they are compatible after all.

First, my SCS account would not obviously preclude any form of moral realism. It is not obviously incompatible with, for instance, a moral realism that affirmed the existence of certain objective "normative" truths, as Nagel has proposed. Nor do moral SCS claims seem to be essentially "relativistic," or irrealist ones. Relativism about a class of claims requires, at least, that there be no standards by which to assess or justify those claims that are independent of each cognizer's beliefs. And there may turn out to be such standards even if the class of claims involved are SCS property ascriptions. So if the SCS account of moral claims poses a problem for AMR, this could not be because SCS claims are eo ipso relativistic.

Second, since Brink himself does not distinguish the modal strengths of possible supervenience principles in any precise way (beyond the distinction between "weak" and "strong" concepts), he does not state explicitly that the modal concepts need to be exactly similar. And there are exegetical grounds to suppose that Brink would not regard his view as threatened by the absence of exact parity. For instance, he regards his own attempt to establish that morality and science are "objective in much the same way" as compatible with the claim that "there are important disanalogies and discontinuities between them" (1989: 6). Evidently then, the general strategy of his defence of moral realism is not to show that morality and science are alike in every regard, but just to show that there is no

single decisive argument against moral realism, and that the cumulative weight of all the evidence ultimately favours that theory. (Brink, in fact, introduces his discussion of supervenience and constitution as part of this strategy, viz., to show that the existence of an "is/ought gap" does not itself undermine the case for moral realism, and I do not pretend to have shown that his attempt is unsuccessful.) Since I am construing Brink's theory as a representative account of AMR, there thus seem to be strong exegetical grounds to doubt that the SCS account is incompatible with AMR.

Third, there may seem to be a way in which the SCS account of moral property ascriptions might be accommodated by AMR. Recall that on the AMR view, moral judgments are justified abductively, on the basis of inferences to the best explanation of certain facts. Briefly, the idea is that one would have reason to believe that a particular judgment (or moral theory) is true if the best explanation of a certain fact or set of facts includes the assumption that the moral judgment or theory in question is true. Among the kinds of facts that (on the view in question) can sometimes be explained on moral assumptions are, variously, other moral facts, moral "observations," moral "conversions," human behaviour, and historical events.¹²⁷ But suppose moral judgments, such as the judgment about Cenci, are indeed, as I have argued, SCS claims. If my account is correct, then the person who makes this judgment will possess a specific cluster of conceptual abilities and inabilities; in particular, they will find it inconceivable that Cenci could have done what he had done, and so forth, but not have been a bad person. But it seems that if conceptual abilities such as these could themselves be explained as resulting from the influence of real moral properties, then the explanations might help justify the ascriptions that cognizers are disposed to make on the basis of their abilities, and the explanations could thereby represent an objective standard for the assessment of the ascriptions. For

¹²⁷ See Brink 1989: 182-297; Sturgeon 1986; 1984; Werner 1985; 1983.

instance, it may be argued that cognizers find it inconceivable that Cenci should do what he did and not be bad because they apprehend (somehow) the fact that he and his actions are bad. Moreover, if, as AMR theorists argue, the cognitive mechanisms that facilitate the apprehension of such moral facts are the ordinary physiological and psychological ones postulated by currently accepted physiological and psychological theories, then the modalities could apparently be brought into line. The conceptual necessity of moral property ascriptions could be regarded as a constraint under which cognizers labour that is imposed on them by their own moral perspicacity: just as observant individuals could not help noticing what confronts them, perspicacious moral appraisers could not help regarding Cenci as bad. It would be physically necessary that cognizers find certain judgments inconceivable. The "objective standard" for assessing moral property ascriptions would thus be comprised by the criteria by which we assess ordinary empirical claims.

I see two problems with this approach. First, even if it were possible to explain the moral/conceptual abilities of cognizers in the proposed manner, there is some reason to doubt that this could allow us to justify moral SCS judgments. If moral judgments are indeed SCS judgments, then being justified in making them would appear to be a matter of having a reason to regard certain scenarios as inconceivable. So, for instance, to have a reason to make the judgment that Cenci is bad in light of his character and actions would seem to require having a reason to regard it as inconceivable that Cenci could have had exactly the character he had, and so forth, but not have been a bad person. But it is unclear how the AMR theorist's favoured account of the epistemic justification of moral claims could succeed in justifying such a stance. On the AMR view, moral claims are justified by their coherence with moral and nonmoral theories and particular judgments; in particular, AMR theorists insist that moral claims can be and are coherent with our best empirical (scientific and ordinary observational) judgments and theories. (For instance, Sturgeon

thinks that moral principles can be tested empirically and that they can pass such tests.) But it is unclear how coherence with these a posteriori theories - theories that are justified largely by empirical evidence - could warrant the sort of modal stance represented by SCS claims. After all, the ordinary empirically-discernible facts that comprise such evidence themselves presumably obtain with, at most, physical necessity (e.g., by obtaining as a matter of "scientific" law), and as such it seems they could furnish no metaphysical, conceptual, analytic, or logical grounds to suppose that a particular set of base properties will necessitate a particular moral property. This is how it works in science: the nomological claim that H₂O, say, always boils under certain conditions (e.g., at 100° C. and so forth) is supported by empirical evidence, and can be said to obtain as a matter of physical (scientific) necessity, but this in itself establishes no stronger nomological claim about the connection between being H₂O and boiling at 100° C.

Second, as we noticed in Chapter One, whether assumptions about moral properties can provide the best explanations of our moral judgments (or indeed any of the other aforementioned items) is itself a contentious issue in the literature. Harman, as we have seen, has argued that the best explanations of an important class of moral judgments (i.e., those made spontaneously in response to certain situations) need make reference only to nonmoral features of the world, and to what he called the "psychological set" of the cognizer, without presupposing that the moral elements of that set are true or are capable of truth (1977). In Chapters Four and Five I have defended Harman's view against a number of recent criticisms. I have argued that given a plausible construal of what moral facts are, moral assumptions must be incapable of explaining (among other things) why cognizers make the moral judgments they make. What is important to note here is that if moral assumptions fail to be explanatory for the reasons I have described in those Chapters, then

they could not plausibly be held to be capable of explaining cognizers' moral/conceptual abilities.

6.9: Conclusion

Brink introduces the notions of supervenience, constitution, and multiple realizability in the course of defending naturalism against the so-called "is/ought" objection. That objection infers from the assumption that there are no entailment relations between ascriptions of moral properties and ascriptions of natural properties that moral properties are not natural properties. The "constitution" account of moral property ascriptions seems to lend the account considerable benefits in this regard. The assumption that moral property ascriptions are ascriptions of properties that are constituted by (though nonidentical to) their naturalistic bases seems to explain, in a way that is both compatible with naturalism and that facilitates the desired parity between morality and science, how moral property ascriptions could be ascriptions of natural properties even though there are no entailment relations between moral claims and scientific claims. However, if the argument of this Chapter is correct, even if moral and nonmoral property ascriptions possess these common features, they do not have exactly similar modal implications. The nonmoral property ascriptions germane to the ostensibly realist disciplines to which AMR theorists seek to assimilate morality are SPS claims, but moral property ascriptions are SCS claims. While the existence of rough modal asymmetry is not, I think, in itself incompatible with AMR (Brink, after all, never insists that there must be exact modal parity between morality and science), the SCS account of moral property ascriptions poses a problem for AMR because it seems incompatible with AMR's own account of the epistemic justification of moral property ascriptions. In order for the SCS account to be compatible with AMR, it must be possible to explain cognizers' conceptual abilities on moral realist assumptions. For on the AMR account, moral judgments are justified largely

by their explanatory power. But if, as I have argued in Chapters Four and Five, moral facts (properly construed) cannot figure as causes and so cannot enter into tenable causal explanations, then assumptions about those facts could not plausibly explain anyone's moral/conceptual abilities, and so could not explain anyone's moral SCS judgments. This suggests that moral property ascriptions could not possibly be justified along AMR lines. And the failure of the AMR account of the epistemic justification of moral property ascriptions is arguably a sufficient reason to reject AMR itself.

CHAPTER SEVEN

CONCLUSION

Brink's Moral Realism and the Foundations of Ethics is arguably the most sustained and systematic articulation and defence of the claim that morality differs in no significant respect from science in terms of objectivity. It contains powerful and sophisticated arguments for scientific/moral parity - arguments that have themselves inspired considerable interest and discussion in the literature. Indeed, not only are Brink's arguments sophisticated, the thesis itself is compelling. It seems that, if AMR is true, then we would have no more reason to dismiss the claims of our best moral theories than we do of our best scientific theories: morality and science would stand or fall together. And AMR theorists appear able to defend naturalistic moral realism against traditional semantical objections to ethical naturalism. However, if the arguments presented here are sound, Brink and other AMR theorists have not provided sufficient reason to accept their theory.

I take the heart of the debate about naturalistic moral realism to be a debate not about the existence of certain complexes of natural and social properties, but about whether the things we call good, bad, and so forth, really do bear any moral properties (1.3). Put most simply, the AMR theorist says that some things really do bear moral properties in virtue of possessing certain natural and social properties; the moral antirealist denies that this is so. The AMR theorist insists that knowledge of moral facts is gained in the same way knowledge of ordinary empirically-discerned facts is obtained. The main threat to AMR is posed by the fact that moral facts, such as the fact that an action is wrong, seem not to be the kind of thing that could plausibly be regarded as the cause of anything. I take

this to be readily apparent. To affirm a moral fact is to assign a moral property to a naturalistic base; it is not just to affirm the existence of things that are thought to bear moral properties, but to affirm that those things do bear moral properties. But even though there is no question that the things we call good, bad, and so forth, can be causes, the fact that things bearing certain natural properties bear a certain moral property in virtue of possessing those natural properties is not itself something that can plausibly figure as a cause (1.4). Such facts are not plausibly regarded as causes just because, as Harman has pointed out, we have no idea how facts of that sort could figure as causes. This is, of course, just a claim about how things seem, but I have offered reasons to suppose that the account is not threatened by those AMR arguments that might be thought to tell against it, viz., Boyd's application of Putnam's "causal" theory of meaning to moral terms (5.5), the Counterfactual and Generality Arguments given by Brink and Sturgeon (4.3, 5.2), and Werner's Argument from *Moral Conversions* (5.3).

If the arguments of this thesis are sound, the worry about efficacy threatens the account in three ways. First, it casts doubt on Brink's dialectical thesis because it is a reason to reject arguments for that thesis that cite facts about our moral phenomenology (2.2, 2.3). I have suggested that such arguments depend (at least implicitly) on Swinburne's Principle of Credulity (or on a close approximation of that principle) (2.3). According to the PC, if it seems to a subject that x is present, then probably x is present, so a cognizer to whom x seems to be present is warranted in presuming that x is present. Thus, it is argued, if it seems to a subject that he is confronted by an inquiry-independent moral fact (or by a realm of such facts), then it is probably so. But as Swinburne himself has noted, one would not in general be warranted in presuming that x is present if one has reason to suppose that x could not have caused the appearance of its presence; this is one of the "limits" to the PC that Swinburne mentions (1979: 261). Of course, if one has reason

to believe that x is not the sort of item that could be the cause of anything, then one could not regard it as the cause of anyone's experience of it (2.3). So, given Harman's insight that moral facts are not reasonably regarded as efficacious, and Swinburne's own limit to the PC, it seems we are not warranted in presuming the truth of AMR on the grounds that the world seems to contain the items whose existence AMR affirms. Moreover, critics of Swinburne have pointed to another plausible limit to the PC: it is reasonable to suspend judgment when controversy among cognitively and perceptually unimpaired persons exists about the existence and efficacy of the fact affirmed by the judgment (2.3). This suggests that we need not even accept Harman's view that moral facts are inefficacious in order to reject the PC-based argument for the realist's dialectical stance. We need only note that the existence and efficacy of moral facts is itself subject to controversy among reasonable and perceptually unimpaired persons. Now, it might be contended that all of this tells only against an argument that Brink does not ever explicitly make, and that his more considered view is that the dialectical stance is warranted by the PC* Argument and not the PC Argument. But if the arguments of Section 2.4 are sound, then the PC* Argument should be rejected as well. The PC* Argument fails because the PC* is not a viable epistemic principle (the fact that an inquiry in an area is generally conducted in order to discover inquiry-independent facts does not itself make realism about that area very probably true), and because a key assumption on which it depends, viz., the assumption that moral inquiry is generally conducted with the goal of discovering real moral facts (of the sort affirmed by AMR), is arguably false.

Second, the assumption that moral facts cannot figure as causes is a (sufficient) reason to reject Brink's account of moral belief justification. That account, as we have seen, is an application to the moral case of a more general attempt to combine realism about empirical claims with a coherentist account of their justification (1.3). Of course, realist

coherentists about empirical beliefs must explain why the coherence of such beliefs can be an indicator of truth (realistically construed). Brink suggests that in general the coherence of an empirical judgment *p* with other beliefs that a person holds can be an indicator of the truth (so construed) of *p* because *p* can be coherent with second-order commonsense and scientific convictions about the reliability of beliefs formed in the way *p* was formed; beliefs that form part of a more general commonsense and scientific account of how creatures such as ourselves obtain knowledge of the world. According to Brink, moral beliefs can be justified in the same way, by the same sort of second-order account. In particular, on Brink's account, considered moral judgments are initially credible because they are coherent with second-order beliefs about the reliability of beliefs that are formed in the way CMJs are formed. So other moral and nonmoral beliefs are lent support by their coherence with CMJs. Included among the second-order beliefs with which CMJs cohere are assumptions about the explanatory capacities of moral facts and about the empirical testability of moral assumptions. Simply stated, it is contended that CMJs and the beliefs with which they cohere are credible because they form part of a comprehensive explanatory account of our observations (just like ordinary observational beliefs ostensibly do). The assumption that moral facts are inefficacious undermines this account because if it is true, then moral fact-affirming beliefs cannot service explanations of our observations and CMJs cannot themselves be explained on the assumption that they have been adopted, as are our observational judgments, as a result of our exposure to the facts that they affirm.

Of course, Brink and Sturgeon do present arguments for the explanatory relevance and indispensability of moral fact-affirming assumptions. But, as we have seen, those arguments are themselves threatened by the worry about efficacy and do not succeed in ameliorating it. The Counterfactual Argument depends on a view of relevance assessment that is generally implausible. It asks us to assess the explanatory relevance of assumptions

about moral facts by considering counterfactual scenarios that are themselves irrelevant to any reasonable assessment of the relevance of those facts. It would ask us, for instance, to assess the relevance of the assumption that cat burning is wrong to a reasonable explanation of the fact that a person has responded to an instance of cat burning by assigning that moral property to the action by considering whether the person would have judged as he did if he had confronted an action naturalistically different enough from cat burning not to be wrong. But considering such a scenario could tell us only about the relevance of the assumption that a person has confronted an instance of cat burning, not about the relevance of the assumption that burning cats is wrong. The only way to assess the relevance of the latter assumption is to consider what would have happened (what judgments the person would have made) if the theory that assigns wrongness to cat burning had been false, just as the only way to assess the relevance of the assumption that water is H_2O to a reasonable explanation of our acceptance of that theory is to consider what would have happened if that theory were false. But in order for the assumptions in question to pass such counterfactual tests, we would need an account of why things would have been different in the counterfactual scenario than they were in the actual one. We have such an account in the water case: if what we call water had had a different microcomposition, then scientists could presumably have discovered this using the same techniques they employed to detect its actual composition. But we have no idea how moral inquirers could discover a counterfactual difference in the moral status of cat burning, for we have no idea how moral facts of the sort affirmed by the judgment that cat burning is wrong could figure as the cause of anything. In short, the apparent inefficaciousness of moral facts is a reason to doubt that assumptions about them can pass any plausible counterfactual test of explanatory relevance (4.3).

Nor, for similar reasons, can the apparent "generality" of moral facts make assumptions about them indispensable to certain causal explanations. We could regard such assumptions as indispensable only if we had some understanding of how structural moral facts could be causes. But the Generality Argument leaves this problem unaddressed. According to the Generality Argument, what makes moral explanations necessary is the fact that in many cases there are no basic properties common to all of the items that appear to precede what we want to explain. The only thing that those items have in common is the fact that they seem to us to realize some moral property. But this does not address the problem of moral efficacy, for it provides no insight into how the fact that a thing bears a constituted moral property could be a cause of anything. Indeed, the thesis that assumptions about moral structural facts are indispensable to certain causal explanations is untenable except on the assumption that facts of the kind postulated can be causes. This is, I think, a very general constraint on the making of viable structural explanations. If compositional facts L_1 , L_2 , and L_3 , precede E (i.e., a fact that we want to explain) but seem to possess no common basic properties, we cannot just postulate any structural fact they may appear to possess as the cause of E ; we can only postulate structural facts that can plausibly function as causes. This is in part why we do and should reject explanations that cite supernatural structural facts: we have no idea how the fact that a naturalistically disparate class of items could assume a causal role in virtue of belonging to a class of items that realizes some supernatural fact (such as being tainted by a curse, blessed by a spiritual healer, and so forth). All of this suggests that moral (fact-affirming) beliefs cannot possess the sort of coherence that on the AMR account is required for justification; i.e., they cannot be coherent with plausible second-order beliefs about their etiology.

Finally, the assumption that moral facts cannot be causes undermines Brink's response to objections to ethical naturalism that are motivated by the semantic test of properties. That response depends largely on a particular construal of moral facts. According to Brink, moral facts are empirically-discerned facts about supervening, constituted (hence multiply realizable) upper level properties. This assumption seems to explain how a particular basic property or set of properties can realize a moral property even if ascriptions of the moral property are not reducible to ascriptions of underlying basic properties. However, appealing to these modal concepts does not assist the case for AMR because the apparent modal features of moral property ascriptions can be explained without making any realist assumptions. In Chapter Six, I argued that moral property ascriptions differ modally from scientific property ascriptions. Moral property ascriptions were found to be SCS claims (6.6); nonmoral macrophysical and psychofunctional property ascriptions were found to be SPS claims (6.5, 6.7). The fact that moral property ascriptions are SCS claims is, I suggested, tentative evidence that they do not ascribe real moral properties to the world. Our scientific property ascriptions obtain with physical necessity because they identify basic properties which, as far as our best scientific theories are concerned, necessitate and explain the properties they produce. Such ascriptions may be explained on the assumption that the facts they affirm are real (i.e., inquiry-independent) and are capable of causing cognizers to affirm them. By contrast, it is natural to explain SCS claims by citing facts about a cognizer's conceptual abilities rather than facts about the world. In 6.8, I suggested that this tentative evidence against AMR could be overturned only if it could be shown that certain conceptual abilities, (i.e., those involved in moral judgment), could themselves be explained by the moral facts affirmed by moral SCS judgments. That is, we could regard SCS property ascriptions as ascriptions of real properties if we could explain them by citing the properties they affirm, but we can so

explain them only if we can explain the conceptual abilities that motivate moral SCS claims by citing the properties they affirm. However, if the arguments of Chapters Four and Five are sound, then assumptions about moral facts are irrelevant to, and unnecessary for, such explanations. Consider Sturgeon's judgment that Hitler is morally depraved. If my arguments are correct, then there is no (moral) fact about the world that explains a person's professed inability to conceive of a world in which Hitler acts as he did but was not morally depraved. For even if the judgment is true, there is no apparent way in which the fact it affirms could causally affect Sturgeon, or anyone else, so as to make them unable to conceive that counterfactual scenario. Moreover, there is no need to cite such facts in order to explain such conceptual inabilities, for we can adequately explain them by citing nonmoral facts about the world and about our moral psychology. Given this, there is no reason to suppose that the tentative evidence against AMR furnished by the modal difference discovered in Chapter Six has been refuted by recent arguments for moral explanations.

BIBLIOGRAPHY

- Armstrong, D. M. 1970. "The Nature of Mind." In C. V. Borst (ed.), The Mind/Brain Identity Theory. London: Macmillan. Reprinted in Block 1980: 191-199.
- Audi, R. 1990. "Moral Epistemology and the Supervenience of Ethical Concepts." The Southern Journal of Philosophy, Supplement 29: 1-24.
- Ayer, A. J. 1952. Language, Truth, and Logic. New York: Dover.
- Baier, K. 1958. The Moral Point of View. Ithaca: Cornell University Press.
- Bambrough, R. 1979. Moral Skepticism and Moral Knowledge. Atlantic Highlands, New Jersey: Humanities Press.
- Bergmann, M., J. Moore, and J. Nelson. 1980. The Logic Book. New York: Random House.
- Blackburn, S. 1971. "Moral Realism." In J. Casey (ed.), Morality and Moral Reasoning. London: Methuen: 101-124.
- _____. 1984. Spreading the Word. Oxford: Clarendon Press.
- _____. 1985. "Supervenience Revisited." Rpr. in 1988 Geoffrey Sayre-McCord, Essays on Moral Realism. Ithaca and London: Cornell University Press: 59-75.
- Block, N. 1978. "Troubles with Functionalism." In C. W. Savage (ed.), Perception and Cognition. Issues in the Foundations of Psychology, Minnesota Studies in the Philosophy of Science, Vol. 9, Minneapolis: University of Minnesota Press: 261-325.
- _____. (Ed.). 1980. Readings in Philosophy of Psychology, Vol. 1. Cambridge, Mass.: Harvard University Press.
- _____. 1990. "Can the Mind Change the World?" In George Boolos (ed.), Meaning and Method: Essays in Honour of Hilary Putnam. Cambridge: Cambridge University Press: 137-170.
- Block, N. and J. Fodor. 1972. "What Psychological States are Not." The Philosophical Review 81: 159-181.
- Bohr, N. 1934. Atomic Theory and the Description of Nature. Cambridge: Cambridge University Press.

- BonJour, L. 1986. The Structure of Empirical Knowledge. Cambridge, Mass.: Harvard University Press.
- Boyd, R. 1981. "Scientific Realism and Naturalistic Epistemology." In Peter D. Asquith and Ronald N. Giere (eds.), Philosophy of Science Association 1980, Vol. 2: 613-662.
- _____. 1988. "How to be a Moral Realist." In G. Sayre-McCord (ed.), Essays on Moral Realism. Ithaca and London: Cornell University Press: 181-228.
- Brandt, R. 1979. A Theory of the Good and the Right. New York: Oxford University Press.
- Brandt, R. and J. Kim. 1963. "Wants as Explanations of Actions." The Journal of Philosophy lx, no. 15.
- Brink, D. 1984. "Moral Realism and the Skeptical Arguments from Disagreement and Queerness." Australasian Journal of Philosophy 62: 111-125.
- _____. 1989. Moral Realism and the Foundations of Ethics. Cambridge: Cambridge University Press.
- Broad, C. D. 1939. "Arguments for the Existence of God. II." The Journal of Theological Studies 40: 157-167.
- Butchvarov, P. 1989. Skepticism in Ethics. Bloomington and Indianapolis: Indiana University Press.
- Clarke, R. W. 1984. "The Evidential Value of Religious Experiences." International Journal for the Philosophy of Religion 16: 189-202.
- Cohen, G. A. 1982. Karl Marx's Theory of History. Oxford: Oxford University Press.
- Copp, D. and D. Zimmerman. (Eds.), 1984. Morality, Reason and Truth. Totowa, NJ: Rowman & Allenheld.
- Dancy, J. 1986. "Two Conceptions of Moral Realism." Proceedings of the Aristotelian Society, Supplement 60: 167-187.
- Daniels, N. 1979. "Wide Reflective Equilibrium and Theory Acceptance in Ethics." Journal of Philosophy 76: 256-282.
- _____. 1980. "Reflective Equilibrium and Archimedean Points." Canadian Journal of Philosophy 10: 83-103.
- Devitt, M. 1984. Realism and Truth. Princeton: Princeton University Press.

- Dreier, J. 1990. "Internalism and Speaker Relativism." Ethics 101: 6-26.
- _____. 1992. "The Supervenience Argument Against Moral Realism." The Southern Journal of Philosophy 30 (No. 3): 13-38.
- Duhem, P. 1962. The Aim and Structure of Physical Science (2nd ed. 1914). Trans. by P. P. Wiesner, New York: Athenium.
- Einstein, A., B. Podolsky, and N. Rosen. 1935. "Can the Quantum-Mechanical Description of Reality be Considered Complete?" The Physical Review 47: 777-780.
- Fine, A. 1984. "The Natural Ontological Attitude." In J. Leplin (ed.), Scientific Realism. Berkeley and Los Angeles: University of California Press: 83-107.
- Forgie, J. W. 1986. "The Principle of Credulity and the Evidential Value of Religious Experience." Philosophy of Religion 19: 145-159.
- Gewirth, A. 1978. Reason and Morality. Chicago: University of Chicago Press.
- Gilbert, A. 1990. Democratic Individuality. Cambridge: Cambridge University Press.
- Gillespie, N. (ed.). 1986. Spindel Conference on Moral Realism. The Southern Journal of Philosophy, Supplement 24.
- Goodman, N. 1954. Fact, Fiction, and Forecast. London: University of London Athlone Press.
- Hare, R. M. 1963. Freedom and Reason. New York: Oxford University Press.
- _____. 1981. Moral Thinking. Oxford: Clarendon Press.
- _____. 1984. "Supervenience." Proceedings of the Aristotelian Society, Supplement 58: 1-16.
- Harman, G. 1965. "The Inference to the Best Explanation." Philosophical Review 74: 88-95.
- _____. 1977. The Nature of Morality. New York: Oxford University Press.
- _____. 1984. "Is There a Single True Morality?" In David Copp and David Zimmerman (eds.), Morality, Reason and Truth. Totowa, NJ: Rowman & Allenheld: 27-48.
- _____. 1986. "Moral Explanations of Natural Facts - Can Moral Claims be Tested Against Moral Reality?" The Southern Journal of Philosophy, Supplement 24: 57-68.

- Heisenberg, W. 1958. Physics and Philosophy. New York: Harper and Row.
- Hempel, C. 1942. "The Function of General Laws in History." The Journal of Philosophy 39: 35-48.
- Hempel, C. and P. Oppenheim. 1948. "Studies in the Logic of Explanation." Rpr. 1988 Joseph C. Pitt (ed.), Theories of Explanation. New York, Oxford: Oxford University Press: 9-46.
- Hodges, W. 1977. Logic. Harmondsworth, Middlesex, England: Penguin Books Ltd.
- Honderich, T. (ed.). 1985. Morality and Objectivity. London: Routledge and Kegan Paul.
- Horgan, T. and M. Timmons. 1990-1991. "Moral Realism Meets Moral Twin Earth." Journal of Philosophical Research 16: 247-265.
- _____. 1992. "Troubles on Moral Twin Earth: Moral Queerness Revived." Synthese 92: 221-260.
- Hudson, W. D. (ed.). 1969. The Is/Ought Question. London: Macmillan.
- Hume, D. 1739. A Treatise of Human Nature. Ed. P. H. Nidditch. New York: Oxford University Press, 1978.
- Huxley, A. 1946. Brave New World. 2nd ed. New York: Harper and Row.
- Jobe, E. K. 1990. "Sturgeon's Defence of Moral Realism." Dialogue XXIX: 267-275.
- Kaplan, D. 1989. "Demonstratives." In Joseph Almog, John Perry, Howard Wettstein, (eds.), Themes from Kaplan. New York: Oxford University Press: 481-564.
- Kim, J. 1972. "Phenomenal Properties, Psychophysical Laws, and the Identity Theory." Monist 56: 177-192.
- _____. 1979. "Causality, Identity and Supervenience in the Mind-Body Problem." In Peter A. French, Theodore E. Uehling, Jr., and Howard K Wettstein, (eds.), Midwest Studies in Philosophy 4, Studies in Metaphysics. Minneapolis: University of Minnesota Press: 31-49.
- _____. 1983. "Supervenience and Supervenient Causation." The Southern Journal of Philosophy, Supplement 22: 45-56.
- _____. 1984. "Epiphenomenalism and Supervenient Causation." Midwest Studies in Philosophy 9: 257-270.

- _____. 1992. "Multiple Realization and the Metaphysics of Reduction." Philosophy and Phenomenological Research LII, (No. 1): 1-19.
- Kitcher, P. 1981. "Explanatory Unification." Philosophy of Science 48: 507-531.
- _____. 1982. Abusing Science. Cambridge, Mass.: MIT Press.
- Klagge, J. 1984. "An Alleged Difficulty Concerning Moral Properties." Mind 93: 370-380.
- _____. 1987. "Supervenience: Perspectives vs. Possible Worlds." Philosophical Quarterly 37: 312-315.
- _____. 1988. "Supervenience: Ontological and Ascriptive." Australasian Journal of Philosophy 66: 461-470.
- Kripke, S. 1980. Naming and Necessity. Cambridge: Harvard University Press.
- Kupperman, J. 1987. "Moral Realism and Metaphysical Anti-Realism." Metaphilosophy 18: 95-107.
- Lear, J. 1983. "Ethics, Mathematics and Relativism." Mind 92: 38-60.
- _____. 1984. "Moral Objectivity." In S. C. Brown (ed.), Objectivity and Cultural Divergence. Cambridge: Cambridge University Press: 135-170.
- Lemos, N. 1994. Intrinsic Value. Cambridge (England) and New York: Cambridge University Press.
- Lepore, E. and B. Loewer. 1989. "More on Making Mind Matter." Philosophical Topics 17: 175-192.
- Lipton, P. 1993. Inference to the Best Explanation. London and New York: Routledge.
- Little, D. 1984. "Reflective Equilibrium and Justification." The Southern Journal of Philosophy 32, (No. 3): 373-387.
- Lovibond, S. 1983. Realism and Imagination in Ethics. Oxford: Basil Blackwell.
- Mach, E. 1960. The Science of Mechanics. La Salle, Ill.: Open Court, trans. T.J. McCormick.
- Mackie, J. 1946. "A Refutation of Morals." Australasian Journal of Philosophy 24: 77-90.
- _____. 1977. Ethics: Inventing Right and Wrong. Harmondsworth: Penguin Books Ltd.

- Maffie, J. 1990. "Recent Work on Naturalized Epistemology." American Philosophical Quarterly 27, (No. 4): 281- 292.
- McDowell, J. 1979. "Virtue and Reason." The Monist 62: 331-350.
- _____. 1981. "Noncognitivism and Rule Following." In Steven H. Holtzman and Christopher M. Leich (eds.), Wittgenstein: To Follow a Rule. London: Routledge and Kegan Paul: 141-162.
- _____. 1985. "Values and Secondary Qualities." In Ted Honderich (ed.), Morality and Objectivity. London: Routledge and Kegan Paul: 110-129.
- McFetridge, I. 1985. "Supervenience, Realism, Necessity." Philosophical Quarterly 35: 245-258.
- McGinn, C. 1981. "Modal Reality." In R. Healey (ed.), Reduction, Time and Reality: Studies in the Philosophy of the Natural Sciences. Cambridge: Cambridge University Press: 143-148.
- McNaughton, D. 1988. Moral Vision. Oxford: Basil Blackwell.
- Melden, A. I. 1961. Free Action. London: Routledge and Kegan Paul.
- Miller, R. W. 1983. "Fact and Value in the Social Sciences." Rpr. in 1991, The Philosophy of Science. Cambridge, Mass. and London, UK: Bradford Books: 743-762.
- Moore, B., Jr. 1978. Injustice: The Social Basis of Obedience and Revolt. White Plains, New York: M.E. Sharpe, Inc.
- Moore, G. E. 1903. Principia Ethica. Cambridge: Cambridge University Press.
- Nagel, T. 1980. "The Limits of Objectivity." In S. M. McMurrin (ed.), The Tanner Lectures on Human Values I. Salt Lake City: University of Utah Press: 77-139.
- _____. 1986. The View From Nowhere. Oxford: Oxford University Press.
- Newton-Smith, W. H. 1983. The Rationality of Science. Boston, London and Henley: Routledge and Kegan Paul.
- Plantinga, A. 1974. The Nature of Necessity. Oxford: Oxford University Press.
- Platts, M. 1979. Ways of Meaning. London: Routledge and Kegan Paul.
- _____. 1980. "Moral Reality and the End of Desire." In Mark Platts (ed.), Reference, Truth and Meaning. London: Routledge and Kegan Paul.

- Postow, B. C. 1985. "Werner's Ethical Realism." Ethics 95: 285-291.
- Putnam, H. 1967. "Psychological Predicates." Reprinted in Block 1980: 223-231.
- _____. 1974. "The 'Corroboration' of Theories." Rpr. in 1991, The Philosophy of Science. Cambridge, Mass. and London, UK: Bradford Books: 121-137.
- _____. 1975. "The Meaning of 'Meaning'." In Mind, Language and Reality. Philosophical Papers, Volume 2. Cambridge: Cambridge University Press: 215-271.
- Quine, W. V. O. 1960. Word and Object. Cambridge, Mass.: MIT Press.
- _____. 1969. "Epistemology Naturalized." In Ontological Relativity. New York: Columbia University Press.
- Rabinowicz, W. 1979. Universalizability. Dordrecht, Boston, and London: D. Reidel Publishing Company.
- Railton, P. 1986. "Moral Realism." Philosophical Review 95: 163-207.
- Rawls, J. 1971. A Theory of Justice. Cambridge, Mass.: Harvard University Press.
- Rowe, W. L. 1982. "Religious Experience and the Principle of Credulity." International Journal for the Philosophy of Religion 13: 85-92.
- Sayre-McCord, G. (ed.). 1988. Essays on Moral Realism. Ithaca and London: Cornell University Press.
- Searle, J. 1964. "How to Derive 'Ought' from 'Is'." Philosophical Review 73: 43-58.
- Shelley, P. B. 1970. The Cenci. Ed. by Roland A. Duerkson. Indianapolis: Bobbs-Merrill.
- Shoemaker, S. 1984. Identity, Cause and Mind. Cambridge, Mass.: University of Cambridge Press.
- Sinnott-Armstrong, W. 1990. "Moral Experience and Justification." The Southern Journal of Philosophy, Supplement 29: 89-96.
- Stevenson, C. L. 1944. Ethics and Language. New Haven: Yale University Press.
- Sturgeon, N. 1984. "Moral Explanations." In David Copp and David Zimmerman (eds.), Morality, Reason and Truth. Totowa, NJ: Rowman & Allenheld: 49-78.
- _____. 1986. "Harman on Moral Explanations of Natural Facts." The Southern Journal of Philosophy, Supplement 24: 69-78.

- Swinburne, R. 1979. The Existence of God. Oxford: Oxford University Press.
- Thompson, E. P. 1963. The Making of the English Working Class. New York: Pantheon.
- Timmons, M. 1990. "On the Epistemic Status of Considered Moral Judgments." The Southern Journal of Philosophy, Supplement 29: 97-129.
- _____. 1990a. (Ed.) The Southern Journal of Philosophy, Supplement 29.
- Van Fraassen, B. 1977. "The Pragmatics of Explanation." American Philosophical Quarterly 14: 143-150.
- _____. 1980. The Scientific Image. New York: Oxford University Press.
- Waller, B. N. 1992. "Moral Conversion Without Moral Realism." The Southern Journal of Philosophy 30 (No. 3): 129-137.
- Werner, R. 1983. "Ethical Realism." Ethics 93: 653-679.
- _____. 1985. "Ethical Realism Defended." Ethics 95: 292-296.
- Wiggins, D. 1976. "Truth, Invention, and the Meaning of Life." Proceedings of the British Academy 62: 331-378.
- Wittgenstein, L. 1980. The Blue and Brown Books. Oxford: Basil Blackwell.
- Yasenchuk, K. 1994. "Sturgeon and Brink on Moral Explanations." The Southern Journal of Philosophy 32 (No. 4): 483-502.
- _____. 1995. "Assimilative Moral Realism and Supervenience." Dialogue XXXIV, (No. 1): 75-97.
- Zimmerman, D. 1984. "Moral Realism and Explanatory Necessity." In David Copp and David Zimmerman (eds.), Morality, Reason and Truth. Totowa, NJ: Rowman & Allenheld: 79-103.