Answering Harman's Relativism

by

Peter Justin Kalef
B.A., University of Victoria, 1998
Postgraduate Diploma, University of Leeds, 1999
M.A., University of Leeds, 2000

A Dissertation Submitted in Partial Fulfillment of the Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Philosophy

Answering Harman's Relativism

By

Peter Justin Kalef
B.A., University of Victoria, 1998
Postgraduate Diploma, University of Leeds, 1999
M.A., University of Leeds, 2000

Supervisory Committee

Dr. Scott Woodcock, Supervisor
(Department of Philosophy)

Dr. Patrick Rysiew, Departmental Member
(Department of Philosophy)

Dr. Conrad Brunk, Departmental Member
(Department of Philosophy)

Dr. William Wadge, Outside Member
(Department of Computer Science)

Supervisory Committee

Dr. Scott Woodcock, Supervisor
(Department of Philosophy)

Dr. Patrick Rysiew, Departmental Member
(Department of Philosophy)

Dr. Conrad Brunk, Departmental Member
(Department of Philosophy)

Dr. William Wadge, Outside Member
(Department of Computer Science)

ABSTRACT

In this dissertation, I propose a new method of dealing with moral relativism. This method, which I call the 'parallel arguments' approach, has the unusual feature that -- if successful -- it neutralizes the force of relativistic arguments in metaethics without making commitments to any particular (or even general) antirelativist position. While other writers have employed this approach in limited ways, I believe that this is the first self-conscious and systematic use of this approach in the particular area I deal with.

The bulk of the dissertation is devoted to a demonstration of the 'parallel arguments' approach against the arguments for moral relativism advocated by one of its most famous contemporary proponents, Gilbert Harman.

The aim of the introduction is to motivate the overall project by showing why a new approach seems to be desirable in dealing with moral relativism. In the introduction, I clarify and justify the 'parallel arguments' approach that will be employed in the chapters to come.

The first chapter sets out the target of the dissertation: the relativistic arguments of Gilbert Harman. In order for the 'parallel arguments' approach to be applied to these arguments, they must first be clarified and cleansed of simpler oversights. For that reason, while the first chapter contains an attempt at an exegesis of all Harman's arguments for moral relativism, that exegesis is accompanied throughout with a critical philosophical gloss. In that gloss, I present and discuss a number of textual and argumentative difficulties in Harman's writings that seem to have been missed by previous critics.

The second chapter is the beginning of the application of the 'parallel arguments' approach to Harman's case for relativism. The chapter is devoted to those relativistic arguments for which Harman is most famous: namely, arguments for moral relativism that stem from an analogy or disanalogy between morality and science. I deal with the first of these arguments quickly, and spend the bulk of the second chapter discussing Harman's most famous relativistic argument. This argument of Harman's is based on a disanalogy between the discovery that there is a proton in a cloud chamber and the discovery that the burning of a cat is immoral. After clarifying more clearly what is at issue in this argument, I present and discuss two distinct 'parallel arguments' responses to it.

The third chapter deals with the other two arguments Harman presents for moral relativism: the argument from moral disagreement and what I call the 'argument

from moral reasons'. I clarify both arguments and, again, present a 'parallel arguments' response to each.

In the conclusion, I return to an issue that was raised in the introduction: might there not be ethicists of a particular philosophical temperament such that they could rightly reject the 'parallel arguments' approach as ineffective? I argue in response that, while this is possible, it does not seem to be a problem for my project.

Table of Contents

List of Figures

Introduction

The Problem:

Those with no formal training in philosophy, even when unfamiliar with the term 'moral relativism', tend to have strong feelings as to the truth or falsity of that view. One often overhears non-philosophers espousing the position that what is true for one person on some moral issue may not be true for someone else (whatever that is meant to say); and somewhat less often, it seems, one hears other non-philosophers greeting such espousals with doubt or derision. Perhaps broad-minded lay individuals tend to find relativism attractive when they do because they feel that, since there are many details of our public and personal lives that should be left up to us, and since spurious claims of objective rightness and wrongness are frequently used as a means of harnessing public support against those whose ways seem foreign, we should be wary of accepting claims of moral objectivity or universality. This, coupled with the discovery of the human tendency to rationalize our way from personal or cultural likes and dislikes to bold claims of non-relativistic goodness and badness, raises doubts about both the viability and desirability of the search for objective morality. Better, it seems to these lay individuals, to abandon the view that morality is anything other than a highly contingent social construction. We should instead look upon each other's moral beliefs and practices with an enlightened respect, understanding that what others hold as good or true in morality is neither any worse or less true, objectively, than anything *we* have.

However, beginning students of philosophy find themselves – or, at least, are meant to find themselves – disabused of this rosy picture of relativism. One of the first things one is meant to discover in studying ethics is that respecting *all* ways of life implies more than just smiling benevolently on unfamiliar, but harmless, sexual practices and funeral rituals: it also implies respecting those social forces that lead to the death camp and the gulag. In fact, it seems a true relativist would even have to respect the very forces of conservatism that lead many to run off and become relativists in the first place.[1] This suggests something self-stultifying about relativism: those who embrace it seem thereby compelled to endorse views that imply that one should not think of relativism as objectively or universally the right view to adopt.

Hence, the most popular textbooks in introductory ethics courses – for instance, James Rachels' *The Elements of Moral Philosophy*, Russ Shafer-Landau's *Whatever Happened to Good and Evil?*, Mary Midgley's *Can't We Make Moral Judgments?* and Julia Driver's *Ethics: The Fundamentals* – tend to dismiss ethical relativism out of hand after offering some apparently decisive objections. Nor are those who learn ethics from general anthologies such as Feinberg and Shafer-Landau's popular *Reason and Responsibility* left with a rosier view of relativism: Feinberg and Shafer-Landau include two pieces critical of relativism and subjectivism, but nothing to *support* either position! Such beginning students also tend to believe, perhaps correctly, that the normative theories they subsequently encounter under the names of 'utilitarianism', 'deontology'

---

[1] But this does not mean, of course, that there is any important connection between social or political conservatism and anti-relativism: those who accept the objective truth of *any* morality, conservative or progressive, must be opposed to relativism.

and 'virtue ethics' would, together with any verdicts reached in their applied ethics courses, be worthless or misguided if relativism were true. All this makes it seem both obvious and fortunate that ethical relativism is a failure.

Nor is there anything particularly new (to philosophers, at least) about these basic objections. It has been well over two millennia since Plato set down in the *Theaetetus* a similar argument against Protagoras' relativism about truth: to say that relativism about truth is true is to imply that those who deny this are also saying something as true as true can be, so relativism about truth entails its own refutation.[2] Plato's attack on relativism in the *Theaetetus* also includes some (perhaps all!) of the more familiar anti-relativistic staples seen today: the charge that relativism cannot intelligibly explain disputes over what is right;[3] the objection that, if relativism is true, all hope of learning must be given up, since everyone already knows the truth of every matter on which he or she has an opinion;[4] etc. One could fairly take from all this the impression that relativism in both its moral and alethic forms does little more than repeat errors identified as such in antiquity, and that the main interest philosophers have in relativism nowadays is to warn outsiders and newcomers of its hazards.

How surprising it is, therefore, when one hears very different verdicts from this from philosophers. Michael Smith, in *The Moral Problem*, summarizes the contemporary scene as follows: "[w]e are told that morality is objective, that there is a single 'true'

---

[2] Plato, 170a-171c. Some philosophers (see, for example, Chapter 3 of Hales 2006 for a recent example) have expressed doubts as to the effectiveness of Plato's treatment of this issue. However, Ferraiolo 2002 offers an interesting response.

[3] Plato, 170c

morality… And we are told that morality is not objective, that there is not a single true morality… Nor should it be thought that this account of the deep disagreements that exist is misleading; that though there are disagreements there are certain dominant views. The situation is quite otherwise. There are no dominant views"[5]. Gilbert Harman, whose relativistic works Smith cites in support of his view that such disagreement exists at the highest level, has recently claimed that "people appear to be quite divided, with about half being moral relativists and half moral absolutists."[6] Harman recently co-wrote a book-length debate on the subject of relativism with Judith Jarvis Thomson (Harman and Thomson 1996): its back cover goes so far as to describe the objectivism/relativism issue as "the key question of contemporary ethical theory".[7] *The Stanford Encyclopedia of Philosophy*'s entry on moral relativism, while somewhat less sanguine than Smith about the popularity of the view, nonetheless asserts that "moral relativism is a standard topic in metaethics, and there are contemporary philosophers who defend forms of it: The most prominent are Gilbert Harman and David B. Wong."[8] Many other contemporary metaethicists have paid relativism at least the backhanded compliment of a lengthy

---

[4] Plato, 161d-162a

[5] Smith 1994, p.4

[6] Harman 2000, p.xiii

[7] It must be admitted that this work (Harman and Thomson 1996) and, indeed, its back cover blurb, have come in for some criticism. Gordon Graham, for example, begins his review of the book in *Philosophy* (Graham 1996) by saying that "[t]he text on the back cover of this book contains two striking falsehoods. According to Michael Smith, it is 'an introduction to meta-ethics of the very best kind', and according to the publisher it will be 'invaluable to all students of ethics, moral philosophy and meta-ethics'. By no stretch of anyone's imagination, and allowing for all reasonable differences of opinion, could either remark be justly made of this book" (p.622). However, I think it noteworthy that even as trenchant a critic as Graham feels compelled to admit that neither Harman nor relativism can be dismissed out of hand or even blamed for what Graham takes to be the failings of the Harman/Thomson debate: he concludes his review with the remark that "[t]he argument between relativism and objectivism has reached something of an impasse, and if two fine minds cannot advance it, this may be because available strategies are, for the moment, exhausted". (*Ibid.*, p.624).

[8] Gowans, Chris, "Moral Relativism", *The Stanford Encyclopedia of Philosophy (Winter 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2008/entries/moral-relativism/>.

objection.[9] Harman's objections against nonrelativistic forms of moral realism[10] from his seminal work on moral relativism, *The Nature of Morality*, have become a central fixture in contemporary metaethical debate in their own right. An illustration of this is the fact that three of the twelve papers in Sayre-McCord's respected anthology, *Essays on Moral Realism*, are devoted to Harman's argument from the causal inefficacy of moral properties.

So what is it about relativism that makes more seasoned philosophers give it more credit than beginning philosophers are meant to? Perhaps the answer has less to do with the virtues of relativism than with the ostensible vices of everything else. Whereas there is doubtless much to be said against relativism, more prominent philosophers tend to feel that there is some difficulty in proposing a viable alternative: the best formulations of moral realism (as well as the most promising non-realist views on hand) have been accused, at least somewhat forcefully, of entailing unfortunate metaphysical or psychological commitments. That this may be the chief reason for the popularity of relativism among serious philosophers is borne out to some extent by the structure of *The Nature of Morality*, Harman's chief work on the subject: Harman's objections against the alternatives to relativism occupy a far more prominent position in his book than his arguments for it. While it is true that many contemporary metaethicists make persuasive attempts to get shock-proof anti-relativistic models up and running, the differences

---

[9] E.g. Brink 1989, pp. 52ff

[10] Moral realism, as will be seen, is a term with many definitions. What they have in common, if anything, is the view that the properties of moral goodness and badness are, in J. L. Mackie's famous phrase, "part of the fabric of the world" (Mackie pp.15, 24). Of course, this is only a metaphor; but as I will argue in the conclusion, expressing what we are talking about in theoretical metaethics often seems to depend on the use of metaphors whose meanings are not always as clear as one might wish.

among the metaphysical and psychological commitments of these models tend to make them targets for philosophical attacks. The ensuing dust storm seems, in the eyes of some, to cast a shadow over the whole anti-relativist enterprise. Meanwhile, other metaethicists have tried to articulate forms of relativism in a way that allows for everything the moral realists want without taking on the burdensome commitments that seem to come with all versions of nonrelativistic realism. Naturally, if these attempts can be made plausibly, and it seems to many relativists that they can, then it becomes even less clear whether Plato has the last word against Protagoras in the moral arena.

To put it more bluntly, and more depressingly: it seems that Gordon Graham could well be right in worrying that "available strategies are, for the moment, exhausted."[11]

Toward a solution: the split personality of philosophy

But *must* the prospect of such a stalemate be depressing? To some, perhaps, it is not. There are many philosophers who feel (or claim to feel) no concern over whether one can solve any of the problems of philosophy. On one variant of this view, according to a pithy saying, philosophy does not seek to answer questions; only to question answers. Another variant, famously defended by Richard Rorty, has it that we should abandon the aim of solving problems in philosophy: instead, we should seek to 'edify' ourselves by discovering "new, better, more interesting, more fruitful ways of speaking." This process, apparently, involves either "the hermeneutic activity of making connections between our own culture and some exotic culture or historical period, or between our own discipline and another discipline which seems to pursue incommensurable aims in an

incommensurable vocabulary", or else "the 'poetic' activity of thinking up… new words, or new disciplines, followed by, so to speak, the inverse of hermeneutics: the attempt to reinterpret our familiar surroundings in the unfamiliar terms of our new inventions."[12] I, however, cannot help but find these views of philosophy trite and unfulfilling: to me, it is precisely the fact that the relativism issue is one that we must resolve for reasons external to mere philosophical concerns that gives it its special importance, and that largely justifies the amount of time philosophers spend grappling with it. If it does turn out to be unresolvable, then I find it difficult to see much point in reasoning or arguing about it any further.

Nonetheless, I must accept the possibility that this problem lies with me. Perhaps my conception of philosophy and its purposes is importantly mistaken. Plato – in at least some of his philosophical moods, it seems – would have responded to my worries by suggesting that I *am* misguided in this regard. Immediately after the section of the *Theaetetus* I cited previously, he has Socrates draw a contrast between the philosopher and the practical man.[13] The former is "the one who has been brought up in true freedom and leisure, the man you call a philosopher; a man to whom it is no disgrace to appear simple and good-for-nothing when he is confronted with menial tasks, when, for instance, he doesn't know how to make a bed, or how to sweeten a sauce or a flattering speech."[14] Since such a person is not particularly concerned with the practical results of any philosophical inquiry, reaching a decision is nothing to hurry: "It does not matter to such

---

[11] Graham, 1996
[12] Rorty 1981, p.360.
[13] Plato, 172c-176a. The gender-specific noun is, of course, Plato's.
[14] *Ibid.*, 175e.

men whether they talk for a day or a year, if only they may hit upon that which is."[15] Whereas "a witty and amusing Thracian servant girl" made fun of Thales for falling into a well while stargazing, on the grounds that he was "wild to know about what was up in the sky but failed to see what was in front of him and under his feet", in fact Thales is just the kind of person we should emulate: "It really is true that the philosopher fails to see his next-door neighbour; he not only doesn't notice what he is doing; he scarcely knows whether he is a man or some other kind of creature".[16]

The more practical sort of thinker is, to Plato, not a philosopher at all: he is a "man of the law-courts", and "is always in a hurry when he is talking; he has to speak with one eye on the clock… And the struggle is never a matter of indifference; it always directly concerns the speaker, and sometimes life itself is at stake. Such conditions make him keen and highly-strung, skilled in flattering the master and working his way into favour; but cause his soul to be small and warped."[17]

When one considers the sort of question Plato deems typical of serious philosophical inquiry, it becomes apparent that the practical benefits of seeing that Protagorean relativism is mistaken are hardly essential to its philosophical importance:

> …within the space of a year, I (a full-grown man) without having been either increased or diminished, am now bigger than you (who are only a boy) and, later on, smaller – though I have lost nothing and you have grown. For this means that I am, at a later stage, what I was not before, and that, too, without having become – for without becoming it is not possible to have become, and without suffering any loss in size I could never become less.… [T]his is an experience which

[15] *Ibid.*, 172d.
[16] *Ibid.*, 174 a-b.

> is characteristic of a philosopher, this wondering: this is where philosophy begins and nowhere else.[18]

Not all philosophers would agree with this characterization of the discipline. The opposite extreme to this can be found in what William James seems to have meant by the 'pragmatic' approach to philosophy. As Bruce Kuclick puts it, "[James] would seek the meaning of 'true' by examining how the idea functioned in our lives. A belief was true, he said, if in the long run it worked for all of us, and guided us expeditiously through our semihospitable world."[19]

Still other philosophers combine these two approaches. An important theme in Hume's writings is that philosophers tend to have *both* conflicting attitudes in relation to the philosophical questions they ponder, and that both these attitudes, or 'humours', must be taken into consideration when we attempt to solve philosophical problems.[20]

In my view, the most important insight by far in all of Harman's ethical writings (though seemingly neglected by commentators) is the recognition that it is enlightening to discuss the moral relativism/antirelativism issue along similar lines. This is an important theme in his 'Is There a Single True Morality?'. In the following passage, he considers why relativists and anti-relativists seem to have such a difficult time persuading, or even understanding, one another:

---

[17] *Ibid.*, 172e-173a.
[18] *Ibid.*, 155b-d.
[19] Kuklick 1981, p.xiv

> …I want to suggest that part of the explanation for this mutual incomprehension is that there are two different ways to do moral philosophy. If one approach is taken, moral relativism, noncognitivism, or scepticism may seem obviously correct and moral absolutism may seem foolish. If the other approach is taken, absolutism may seem clearly right and scepticism, relativism and noncognitivism may seem foolish.
>
> The difference in approaches is, to put it crudely, a difference in attitude toward science. One side says we must concentrate on finding the place of value and obligation in the world of facts as revealed by science. The other side says we must ignore that problem and concentrate on ethics proper. Both sides agree that we must begin at the beginning with our initial beliefs, moral and nonmoral, and consider possible modifications that will make these beliefs more coherent with each other and with plausible generalizations and other explanatory principles… The process will leave some issues unresolved; in particular, we may find ourselves with no account of the place that value and obligation have in the world of facts. This will not dismay someone who is willing to leave that question unanswered. But it will be disturbing to someone who, on the way to reflective equilibrium, has come to think that the basic issue in moral philosophy is precisely how value and obligation fit into the scientific conception of the world.[21]

Harman does not develop this idea much further, and I have serious doubts about some aspects of this view: I see no good reason to believe, for instance, that the division between those concerned with the place of values in the world and those concerned with solving practical problems in ethics will have much to do with the relativism/nonrelativism fault line. However, as I will now argue, this discussion of Harman's contains an important insight that can be developed into a way out of this metaethical stalemate.

An Outline of the Solution

The central idea I pursue here is that there may well be difficulties in presenting a response to relativism in a way that will satisfy philosophers in both humours; but that, if

---

[20] I take this to be the central point made in Book I, Part IV, Section 7 of the *Treatise*.

one separates out the two types of philosophical aims, it is possible to provide to each a more persuasive answer to the challenge of relativism than could be given in the attempt to satisfy both aims at once.

To see how this works, let us consider each of these two philosophical attitudes in its most sophisticated form. I begin with the philosopher who, in Hume's phrase, "considers man chiefly as born for action."[22] This philosopher need not, *pace* Harman, be ready to *ignore* the arguments for relativism in favour of "concentrating on ethics proper", even if this form of extreme pragmatism seems to have its champions in the likes of James.[23] If it really *is* the case, or if we really have good reason to *believe*, that the people we think ought not do something are in fact permitted or even obliged to do it, then that is something a more sophisticated practical philosopher will care about. If claims about morality seem to be relative to culture or personal inclination in a way in which (say) scientific claims, claims about reasoning, etc. are not, then no sincere philosopher is entitled to ignore that fact. However, if the best arguments for the relativity of morality turn out to be no better than the best arguments for the relativity of science, logic, etc., then the practical philosopher can legitimately remain undisturbed by relativistic arguments. The question of whether moral claims are best understood as statements of natural facts, statements of nonnatural facts, expressions of norms, etc. seem unlikely (I will argue) to have any effect upon practical matters; so they can safely be ignored. A practical person cannot consistently hold his or her ability to reason at the most basic

---

[21] Harman 2000e, p.79
[22] Hume 1993, p.5
[23] Objections against James' extreme view can be found in A. J. Burger's essay in Clifford, James and Burger 2001 and in Blackburn 2005, pp.4-13.

level in doubt, since this will make living one's life impossible. Hence, from the viewpoint of practical philosophy, trust in the claims of reason are as secure as secure gets. Hence, any doubts about morality that, if we were consistent in accepting them, would commit us to doubts about reasoning itself can legitimately be dismissed. A full-blooded relativism about science, according to which (for instance) the claim that the earth is flat is no more true objectively than the claim that it is not, will also be unacceptable to the sort of philosopher I am discussing here. If these doubts can be dismissed by means of linguistic or metaphysical arguments, well and good; but this is not necessary for the practical philosopher. The practical philosopher need not take an interest in metaphysics per se: its value is merely instrumental to him or her. And this is a good thing for our purposes, since it is precisely in attempting to articulate and defend metaphysical hypotheses that anti-relativists seem to run into problems.

From this, a natural way of responding to relativism while avoiding metaphysical and other burdens becomes clear: I call it the **parallel arguments approach**. The approach is as follows: *For every argument offered for moral relativism, avoid the temptation to defend a particular nonrelativistic metaphysical, psychological, etc. position: instead, show that, in each case, a parallel argument could be constructed in favour of scientific relativism or relativism about basic reasoning*. Since one must accept the conclusions of both or neither of a pair of parallel arguments, and since the practical philosopher will not accept any argument in favour of any interesting relativism regarding science or basic reasoning, this strategy should be all that is needed.

The term 'interesting' in the previous sentence deserves special comment, for it points to one sense in which it *might* not matter at all to a practical philosopher whether relativism is true. To see this, consider the range of positions one might call moral relativism. Let us begin by considering this position:

> Trivial Moral Relativism**:** *some* moral truths are culturally and personally relative.

This statement of the view is trivial because it says something everyone seems to accept. Which relevantly-informed person denies, for instance, that showing up five minutes late to a dinner party is permissible in the culture of the British Columbian gulf islands, impermissible in mainstream German culture, and that there is no objective right or wrong in the matter?

A better statement of relativism (since it actually says something that others might reasonably doubt) is as follows:

> Interesting Moral Relativism: *all* moral truths are culturally and personally relative, and anything one chooses to do is perfectly morally right.

This is certainly the sort of relativism the practical philosopher will care about. However, it will be very difficult (if not impossible) to defend it against the typical anti-relativist objections. The more sophisticated relativist will, therefore, see the wisdom in moving the target a little. We might expect this to be the result:

> Fairly Interesting Moral Relativism: Though there is a morality that correctly applies to nearly everyone, there are a few actual people (namely, sociopaths) and a number of nonexistent people (whom we can imagine in fictional scenarios) to whom this morality does not apply.

The relativist's desire to make a safer and only vaguely controversial claim might be felt more strongly, however, in which case we might expect this:

> Boring Moral Relativism: There is an objective morality that applies to every real or conceivable person in this world, but one can imagine possible worlds in which this morality does not apply.

Or even this:

> Extremely Boring Moral Relativism: There is an objective morality that applies to every real or conceivable person in this or any other possible world. However, conceivability does not entail possibility, and it turns out that there are some people who could hypothetically be conceived, though they cannot belong to any possible world, and this objective morality may not apply to them.

The point to note here is that, to a practical philosopher, it makes no difference whether moral objectivism, Boring Moral Relativism, or Extremely Boring Moral Relativism is true, so long as (as seems to be the case) none of these views has any distinguishable practical consequences. In these cases, showing that nothing practically interesting can follow from a particular form of relativism can be a perfectly adequate response to it.

Hence, the task of answering moral relativism need not, to a practical philosopher, involve refuting it.

I turn now to the other variety of philosopher. Philosophers of this type "consider man in the light of a reasonable rather than an active being", who "rest not satisfied till they arrive at those original principles, by which, in every science, all human curiosity must be bounded."[24] To thinkers such as these, it seems the practical philosopher makes the error of dropping out of the game just when things get interesting. These more *profound* philosophers (for want of a better word) may rightly feel undeterred by the threat of a self-stultifying claim on the horizon: we may not be able to *think* consistently that our beliefs are neither more nor less justified than any others, but that might still be true of our beliefs even if we cannot think it. More than this is needed to provide such philosophers with an answer to relativism.

However, it may also be possible to provide satisfactory but metaphysically noncommittal answers to the *profound* philosophers. The trick here will be to show that the attempt to articulate and defend precise positions on these issues is somehow problematical or misguided. Such moves as the one I am proposing form a long and respected tradition in philosophy dating from at least as far back as Locke's *Essay*. In essence, these moves involve the presentation of an *a priori* argument to the effect that all metaphysical claims are either a) nonsensical; b) meaningful, but discuss things of which we can have no knowledge; or c) relativistic. While many philosophers today seem to find arguments of at least one of these varieties to be plausible, some do not. However, I

propose to offer to the profound philosophers what are, I think, three important improvements on the standard line. First, my scope is more narrow: I seek only to cast doubt on the prospects of abstract metaethical arguments and investigations, not of metaphysical ones in general. Second, my reasoning is not based on general *a priori* claims about metaphysics, but rather on particular difficulties that seem essential to abstract metaethical disputes as seen in the recent literature. Finally, my approach has the advantage of attempting something less ambitious than what has been tried in the standard attempts mentioned above: rather than attempt to show that such metaphysical explorations are impossible in theory, I seek only to shift the burden of proof to the other side.

Actually, the decision to treat the problem of answering moral relativism as two different problems brings with it a fourth advantage over the standard attempts. Many readers of the great works of the anti-metaphysical tradition in philosophy continue to feel that the attack on speculative metaphysics is somehow incomplete. Such readers often have a difficult time resisting the urge to return to the scene of the crime, "where many a fog bank and rapidly melting iceberg pretend to be new lands and, ceaselessly deceiving with empty hopes the voyager looking around for new discoveries, entwine him in adventures from which he can never escape and yet also never bring to an end."[25] But with the two philosophical temperaments kept separate, this can be both understood and, I think, avoided: there *is* something unsatisfying in offering a limited critique of metaphysics in place of the establishment of a position there, but what is missing is needed by our

---

[24] Hume 1993, p.5
[25] Kant, B295.

practical side, not our profound side. And where the question at issue has no practical consequences, the lack of philosophical 'closure' is, I think, something we can much more easily learn to live with.

Philosophers in both practical and profound humours agree on one important thing: an argument that commits a fallacy must be rejected. The exposure of these fallacies is therefore not only the simplest, but also the most generally persuasive, method of making philosophical progress for those in either humour. This shows a way to save time: if a relativistic argument contains a fallacy, it will always be better to expose this fallacy *decisively* than to deal with it by employing the parallel arguments approach or by raising doubts about the possibility of a metaphysical solution.

With all this in mind, the way forward becomes evident:

1) First, consider the best relativistic arguments on offer. As the quotes provided earlier in this introduction suggest, Gilbert Harman may well be most esteemed moral relativist today. I set out Harman's arguments in Chapter One.
2) Second, expose any straightforward fallacies found in these relativistic arguments. Most, but not all, of this is done in Chapter One.
3) Third, deploy the parallel arguments approach against what is left of these relativistic arguments. I do this in Chapters Two and Three.

4) Finally, take what is left (which will by now be free of both fallacies and practical import) and show that even the best prospects of making progress on these remaining points seems problematic. I attempt this in the Conclusion.

I hope this promise of a way around the metaethical deadlock is an exciting prospect for the reader. Before turning to Chapter One and my exploration of Harman's arguments, however, I hope the reader will bear with me as I emphasize some crucial features of the parallel arguments approach that may not be obvious.

The Parallel Arguments Approach

The special force of the parallel arguments approach lies in its being noncommittal on many issues. In the Conclusion, I will discuss this point further; but it is important to be clear from the beginning that I am not concerned in this dissertation to argue for moral realism, nor for cognitivism in ethics, nor for reasons or motivational externalism, nor for the view that morality is objective or universal. Nor should the approach be confused with the adoption of a position on the sort of issue discussed by Darwall, Gibbard and Railton in their well-known article 'Toward Fin de siècle Ethics: Some Trends'[26] This article maps out the terrain of meta-ethics as follows:

> [w]e can distinguish two broad trends in contemporary moral theory depending upon how "the problem of placing ethics" is identified and faced, and the implications drawn. The first starts out from the idea that the "problem" is a product not of ethics, but of the wrong-headed notion of seeking to understand the objectivity of moral judgments on the model of the objectivity of empirical science. This approach depends upon finding some substantial contrast or discontinuity between facts (at least, facts of the paradigm sort treated of in natural science) and norms or

[26] Darwall, Gibbard and Railton 1992, 128-130

> values... The second broad trend in response to "the problem of placing ethics" accepts the challenge of showing that moral judgments are factual in the paradigm sense afforded by empirical or theoretical judgments in the natural sciences

The 'Fin de siecle' article terms the first of these trends (i.e. "the view that there is a discontinuity between ethics and science"[27]) the 'discontinuity' approach, and calls the second the 'continuity' approach. However, I must stress here that the continuity approach is *not* what I am discussing when I refer to the parallel arguments approach. The continuity approach, as can be seen from the above quote from Darwall, Gibbard and Railton, is concerned with the sense in which moral statements can be seen as factual: it is a view about moral objectivity. The parallel arguments approach, by contrast, is first of all, a *strategy*, not a view or position; and secondly, it is concerned not with a position per se, but with the strength and implications of *arguments* for that position. Moreover, one's commitment to the parallel arguments approach neither implies nor is implied by commitment to the continuity approach.

It may be clearer to see that this is the case if one considers the following analogous situation: I hear some unphilosophical theist remarking that atheists must be mistaken in their view because God exists whether we believe in Him or not. I respond by pointing out that, by parallel reasoning, it could be argued that those who deny the existence of the Flying Spaghetti Monster must be mistaken. Now, on the basis of this information alone, what is one entitled to assume about my own religious beliefs? Certainly, it would not follow from my having said this that I think that God is no more real than the paradigm case of an unreal thing, the Flying Spaghetti Monster. Nor would it be legitimate to

---

[27] Darwall, Gibbard and Railton (1992), p.130

conclude this even if I made the further claim that *all* the arguments I had heard for the existence of God could be used just as well to justify the existence of the Flying Spaghetti Monster. For these statements could also be made by:

- a believer in God who feels that theists should try to come up with better arguments for his existence than any that have been given so far;
- a believer in God who feels that theists make an error if they try to justify belief in his existence through rational means;
- a believer in God who wishes to get the original speaker to agree that good reasons can be given for many other supernatural beings, and hence that she should be more tolerant; an agnostic;
- a believer in the Flying Spaghetti Monster; someone who really doesn't care whether or not there's a God, but just wants to point out the problem with a particular argument;
- a logical positivist, who thinks that such first-order statements about God are all nonsensical; etc.

Similarly, one might employ the parallel arguments approach and hold that:

- moral judgments are factual in the same way, and for the same reasons, as scientific statements are;
- that they are factual in the same way, but for different reasons;
- that they are importantly factual, but not in the same way;
- that neither they nor scientific judgments are factual;
- that we do not yet know whether ethical judgments are factual;

- that we will never know whether they are factual;

- that they are not even truth-apt; and so on.

Furthermore, it does not follow from the fact that someone believes in what Darwall, Gibbard and Railton term the 'continuity' approach that such a person would employ or even accept the parallel arguments approach. One can hold "that moral judgments are factual in the paradigm sense afforded by empirical or theoretical judgments in the natural sciences" and still have doubts about – or deny outright – the view that for each argument against moral realism, there is an equally powerful argument concerning scientific realism, say. It could be, for example, that one holds that quite different arguments are needed to argue for these two sorts of realism, but that these arguments nonetheless show that we end up with an equally robust form of realism in both cases. Since a commitment to the parallel arguments approach is neither a necessary nor a sufficient condition for a commitment to the continuity approach, they are clearly not interchangeable.

Chapter One: A Critical Exposition of Harman's Case for Relativism

Introduction

In this chapter, I offer a critical exposition of Harman's entire case for relativism. This project has significance of its own: to my knowledge, there has never been a full treatment of Harman's ethical views. Discussions of *The Nature of Morality*, Harman's major work on the subject, tend to consider one or two of his arguments in isolation from the rest, often without any clear explanation of how they might have been intended to work with the others to make up a greater overall case for relativism. That is unfortunate. Harman clearly intends his arguments to be considered together. This is evident in the preface of that work, in which he says "[a]lthough I discuss a variety of views, it will be clear what my own position is. *The book as a whole represents a sustained defense of that position*" (emphasis added).[28]

While Harman's original case for relativism in *The Nature of Morality* is in a sense his most 'current' (in that it seems to have received the most attention in recent years), I will also discuss a few clarifications and amendments Harman has made in the years since 1977.

In writing this chapter, I have tried to find a balance between its two aims: that of introducing the reader to Harman's arguments as clearly and charitably as possible, and

that of exposing what I take to be problematic areas in his reasoning. Ideally, it should be possible to read the chapter conveniently with either end in mind. I have attempted to satisfy both purposes as follows: in the main text, I have limited myself to a purely exegetical role. There, I summarize Harman's arguments, supplementing my synopsis with direct quotes from his writings whenever I feel that this will help the reader in understanding Harman's intentions or make clear the wording of a passage I wish to attack as problematic. The actual philosophical work of the chapter, the critical commentary, is confined to an extensive set of footnotes accompanying the text. I hope that any awkwardness this presents will be outweighed by the clarity it provides.

The Nature of Morality

A Problem in Confirming Ethical Theories

Harman begins his central ethical work by raising some concerns about moral principles and our ability to confirm them. Scientific principles can be tested against the world, Harman feels; but it is not clear to him that moral principles can be[29]. "The difference", Harman claims,

> is that you need to make assumptions about certain physical facts to explain the occurance of the observations that support a scientific theory, but you do not seem to need to make assumptions about any moral facts to explain the occurrence of the so-called moral observations I have been talking about. In the moral case, it would seem that you need only make assumptions about the

---

[28] *The Nature of Morality*, p.ix
[29] *The Nature of Morality*, p.4

> psychology or moral sensibility of the person making the moral observation. In the scientific case, theory is tested against the world.[30]

As Harman sees it, this should make moral antirelativists feel uneasy.

Since this issue (which Harman calls the 'basic issue'[31]) pertains to our ability to *verify* or *support* moral principles, and since the relativism issue in ethics pertains to the question of whether a moral principle could be true *whether or not we believe it to be*, one would expect Harman to provide the reader with some reason to feel that the first, epistemological issue is meant to be connected with the second, metaphysical one. Surprisingly, he does not do this, but spends the rest of Chapter 1 clarifying what he has already said.

Toward the end of this chapter, he acknowledges that not all legitimate scientific hypotheses can be tested by direct observation after all (he gives as an example certain hypotheses about black holes), which seems to present some difficulties for his major point in the chapter. However, these difficulties are never discussed, nor is the matter mentioned again. Lastly, Harman considers the possibility that moral principles should be seen as more closely comparable with mathematical than with scientific principles. One does not, after all, observe mathematical facts directly; and Harman agrees that this does not imply a problem with mathematics. However, he rejects this possibility on the grounds that we do confirm mathematical principles through observation every time we

---

[30] *Ibid.*, p.6. Discussion of this point has become a significant fixture of contemporary metaethical debate. I will consider it in some depth toward the end of Chapter 2.
[31] *Ibid.*, p.3

use mathematical principles to support a scientific theory, whereas (he feels) this does not happen in the case of ethics.[32]

Nihilism

After setting out this problem concerning observation in ethics, Harman turns to what he calls 'extreme nihilism', which is the view that "morality is simply an illusion: nothing is ever right or wrong, just or unjust, good or bad. In this version, we should abandon morality, just as an atheist abandons religion after he has decided that religious facts cannot help explain observations"[33]. Harman immediately rejects this view on the grounds that it is "hard to accept"[34], since it would follow from such a view that there is nothing wrong, say, in murdering one's own parents. Though this treatment of extreme nihilism is very brief, it is noteworthy in that it shows Harman to be committed to the principle that one can rightly reject a position merely on the grounds that it is difficult to accept it[35]. I will refer back to Harman's reliance on this principle in what follows.

Noncognitivism

---

[32] *Ibid.*, p.10. The most important critical responses to these issues, and my own attempt at a contribution to them, will be presented in Chapter Two.

[33] *Ibid.*, p.11. Harman is less careful than one might wish him to be here: someone who "has decided that religious facts cannot help explain observations" would presumably hold that we can have no reason to believe that God does or does not exist (and hence would become an agnostic, not an atheist); or else hold, in the manner of A. J. Ayer, that first-order claims about God's existence are nonsensical. Such a person might even become a mystical sort of theist. It is not at all clear, however, how the explanatory impotence of religious facts would provide good grounds for becoming an *atheist*.

[34] *Ibid.*

[35] Harman does, later on, present the caveat that "[a]lthough such a conflict [i.e. one between a philosophical position and 'ordinary ways of talking and thinking'] does not refute a theory, we must ask whether we can accommodate the point about ethics and observation without having to give up our ordinary views and endorsing some form of nihilism" (p. 13). Still, this seems to leave him committed to the view that his relativism should be rejected if it is *at least* as absurd as extreme nihilism.

Harman believes that there is another sort of nihilism, which he calls 'moderate nihilism', to consider. Moderate nihilism is the view that "the purpose of our moral judgments is not to describe the world but to express our moral feelings or to serve as imperatives we address to ourselves and others"[36]. Readers with a background in ethical theory will recognize the view described here as *noncognitivism*, one of the dominant metaethical positions of the 20th century. (One must wonder, though, whether the great proponents of noncognitivism would generally agree to be characterized as 'nihilists'!) In any event, Harman does not give much consideration to noncognitivism/moderate nihilism either. He dismisses it swiftly from consideration (at this point, at least) on the grounds that it is at variance with our common ways of speaking: "We ordinarily do speak of moral judgments as true or false; and we talk as if we knew certain moral truths but not others"[37]. Again, what is of greatest interest here is the standard Harman must implicitly have in place if he is to be justified in dismissing noncognitivism in this way: whereas his rejection of extreme nihilism rests on a commitment to common sense beliefs, his rejection of noncognitivism rests on a *prima facie* commitment to the correctness of ordinary language.

Naturalism

If we are to accept that moral judgments are statements that express genuine facts about the world (which is what Harman feels his provisional rejection of both forms of 'nihilism' entails), then what sorts of genuine facts could these be? Perhaps some form of reductionism is true: moral facts might be reducible to some other sort of facts

---

[36] *Ibid.*, p.12
[37] *Ibid.*, p.13

(preferably to natural facts – the sorts of facts that are, or could be, dealt with by the natural or social sciences. 'Reductive naturalism' is the view that such reduction is possible). The problem Harman claims to have found in his first chapter could be solved, he feels, if "moral facts can be reduced to other sorts of facts and… assumptions about these facts do help explain observations. In that case, there could be evidence for assumptions about moral facts".[38] He attempts to illustrate this by comparison with the case of colours. We "might be able to account for color perception", he says, "without making the supposition that objects actually have colors"[39]. This could be achieved if we could account for colour perception in terms of the tendency of certain natural objects to reflect light in certain ways, and of the effects of this light upon the observer. "This would not prove that there are no facts about colors", Harman explains; "it would only show that facts about colors are not additional facts, over and above physical and psychological facts"[40].

Functionalism

If we are to take the naturalist/reductionist route, then it seems reasonable to attempt a sort of functionalism regarding moral properties. This would imply that moral properties (like goodness) are reducible to the property of serving a particular function. To take an example from Harman, to say that a knife is good is to say that it performs its function (cutting) well. Harman is comfortable with this analysis of 'goodness' in the case of the

---

[38] *Ibid.*
[39] *Ibid.*
[40] *Ibid.*, p.14. It is difficult, given the two quotes just presented, to be sure what status colours (and by extention, moral properties) are meant to have on this view. The point of the first of these quotes, it initially seems, is that objects would not have colours on the hypothesis Harman is considering; but in the second, he says that there could be such colours. Rather than engage in speculative interpretation, I will

knife, and even with the claim that a knife ought to have certain properties and function in certain ways.[41] Whether the object in question is an artefact like a knife or a watch or a natural organ like a heart, "it is a factual question what the relevant function of [the object] is", because in the case of artefacts the function is "determined by their makers and users" and in the case of natural organs it is "determined by their roles in sustaining those systems"[42]. It follows from this that, unless Harman can provide a reason in the form of a relevant difference for why the case of morality is relevantly different from these sorts of cases, the alleged problem he has found with ethics has a simple, nonrelativistic and non-nihilistic[43] solution.

In his attempt to provide such a relevant difference, Harman considers a sequence of cases spanning the gap between facts about watches and hearts and facts about morality, in which the difficulties of functionalism in morality are meant to become clear. He begins by discussing the terms 'a good meal', 'a good swim', and 'a good time'. These terms "answer to the relevant interests" of the individuals who have a stake in these matters – someone might consider a 'good' meal to mean a tasty one, while someone else might consider a 'good' meal to mean a nutritious one – so "where different interests are relevant, we get ambiguity".[44]

---

leave this puzzle to the reader and will try to account for the rest of Harman's discussion of naturalism as best I can without running afoul of his somewhat obscure intentions here.

[41] *Ibid.*

[42] *Ibid.*, p.15

[43] From here on, I will use the term 'non-nihilistic' to mean what Harman would mean by 'opposed to or incompatible with extreme nihilism', in his terminology.

[44] Harman seems to feel that this ambiguity provides some support for relativism. However, it appears to provide no such support. It is true that ambiguity in some sense appears to follow *from* relativism about values in that, if there were no universal or objective values, then statements of the form 'X is of value' would be ambiguous insofar as they would be incomplete: the ambiguity would disappear only when they were made more specific by stating whose values are at issue (i.e. 'X is of value **to Y**'). However, the

Harman follows this by discussing the term 'a good person'. Since "a person is evaluated in terms of functions, roles, and various interests in a way that is hard to specify"[45], it can be difficult to determine whether someone is good overall.[46] Still, Harman admits that in this case, "the words 'ought' and 'wrong' are relevant as before", at least when we are clear on the context.[47]

Finally, he discusses cases involving "kinds of things… not associated with functions, purposes, or sets of interests":[48] a rock, for example, could be good if used as a

---

converse does not hold. A term can be genuinely ambiguous without relativism of any kind being true. For instance, I may talk about the day I met someone, meaning by this only that I met that person during a given 24-hour cycle. My audience might incorrectly take me to be saying that the meeting took place during sunlight hours. There is clearly a ambiguity here arising from two uses of the term 'day', both of which are correct and either of which my hearer and I might accept as legitimate; but this would be so whether or not relativism concerning either facts or values were true. It is not clear why the case Harman discusses is meant to be relevantly different. Nor should the presence of ambiguity be seen as a necessary or sufficient condition for (extreme) nihilism. There is ambiguity in the 'day' case described above, but that case does not involve nihilism. Furthermore, one could be an extreme nihilist about moral values without being committed to using terms ambiguously as a result.

[45] *Ibid.*, p.15

[46] Harman seems to think that this provides support for his point about ambiguity and relativism. However, the problem of determining how to weigh competing and perhaps incommensurable factors in sorting out whether some person or action is good or bad is not only distinct from the problem of resolving the *ambiguity* in the term 'good' in some particular case; it is also an entirely separate matter from both relativism and (extreme) nihilism in ethics and gives no support to either view. Consider, for instance, a case where I have to decide whether to kill a stray dog or allow some important document to be eaten by it. Determining whether such an action would be good, or whether I would be a good person if I were the sort who would typically kill the dog in that sort of situation, may be difficult due to the very different sorts of competing moral obligations I have toward the dog and toward those interests and principles that may somehow be served by the document's preservation. Nonetheless, it clearly does not follow from this fact that it cannot be objectively right or wrong overall for me to choose one way or another, or that moral values do not exist in this situation. It is also clear that the difficulties offered by this dilemma do not hinge on the ambiguity of the term 'good', or of any other term. If Harman is not running this issue together with the one about ambiguity, however, then it is very difficult to see why he discusses it at this point. I take it from this that the interpretation I have given is correct.

[47] *Ibid.*

[48] *Ibid.*, p.15

paperweight but bad (because too light) as a doorstop, though neither these nor any other particular functions seem to be especially the purpose of a rock.[49]

Harman does not feel that these last cases present any problem for naturalism: we simply need to employ a more complicated type of terminology, and say that some X is good as an A but not as a B. It also does not seem to be a relevant difficulty (as Harman seems to agree) that we sometimes want to talk about the merits of a rock more generally. For instance, one can imagine a case where I am looking for a rock to use as both a paperweight and a doorstop, and which I may also wish to have an attractive colour. Given this diversity of wants, it may be that I find it difficult to determine, of a given set of rocks, which is best. It may also be that I cannot fully articulate, and may even be unaware of, a precise formula for determining which rock will be best (or even which ones will be good). However, none of these facts, as Harman seems to recognize, leads to relativism or nihilism about value in regard to rocks. In fact, Harman goes on to say, such an analysis even gets around the alleged problems of value and observation he began by discussing: "That my watch is a good one may not explain anything about my observations of it; but that it keeps fairly accurate time does help to explain its continual agreement with the announcements of time on the radio and perhaps the goodness of my watch consists in facts of this sort"[50].

If talk about goodness and 'oughts' are unproblematic in these cases, why should such talk be problematic in ethical ones? Harman's answer is, in the end, not easy to

---

[49] *Ibid.*, p.16
[50] *Ibid.*

understand. He draws our attention to a hypothetical case of a doctor who can save five patients from death by cutting up a healthy patient and transplanting his organs into the other five as needed. According to Harman,

> [t]he problem is that in either case you would be satisfying certain interests and not others. The interests of the five dying patients conflict with the interests of the healthy patient in Room 306. The question is what you ought to do, taking all interests into account. As we saw earlier, our intuitive judgment is that you ought not to sacrifice the one patient in Room 306 to save the five other patients. Is that a factual judgment? If we suppose that it is a fact that you ought not to sacrifice the patient in Room 306, how is that fact related to facts that can help explain observations? It is not at all obvious how we can extend our analysis to cover this sort of case.[51]

However, Harman quickly acknowledges that "the problem is not peculiar to ethics"[52]: a shockproof but unattractively bulky watch, or an effective but misery-inducing teacher, might be good in some ways and/or for some purposes but not for others, and this does not lead us to the special sort of problem to which Harman feels moral goodness is vulnerable. It is worthwhile paying attention to Harman's wording here:

> But even given further specifications of our interests in watches and teachers in this way, there may be no fact of the matter as to which watch or teacher is better – not because these are not factual questions but because of vagueness of standards. Factual questions are still factual even when they cannot be answered because of vagueness (Is a door open or shut if it is slightly ajar?) Furthermore, even in cases where we feel intuitively that one watch or teacher is clearly better, we may not be able to specify very clearly the interests, functions and roles with reference to which

---

[51] *Ibid.* Of course, Harman's analysis here relies on certain assumptions that need not be granted: first, that the intuition that it would be wrong to kill the healthy patient for the sake of the others needs to be followed if we are to act morally – in other words, that it *is*, as Harman says, "a fact that you ought not to sacrifice the patient"; second, that it is no worse to kill a healthy person than it is to allow one or more persons to die; third, that morality (or even more specifically, moral naturalism) is in serious trouble if there are genuine moral dilemmas; etc. Since utilitarians would reject the first assumption, and deontologists and virtue ethicists reject the second, and there appears to be no good reason for anyone to accept the third, the force of this passage is questionable for reasons quite separate from those Harman goes on to discuss.

[52] *Ibid.*

> one is better, as a watch or teacher, than the other. Still, it may well be a fact that one is better – a fact constructed in a way that we can only vaguely specify from facts of a sort that can help explain observations.[53]

Harman admits that if it were possible to break down, in a similar manner, moral facts into ones that he feels could explain observations, ethical naturalism would be justified and the problem would be solved. However, at this stage he hesitates: "the relevant functions, roles and interests", he says, "can at best be only vaguely indicated, so the proposed analysis is difficult to evaluate. Nihilism remains a possibility."[54]

The Open Question Argument

At this stage, Harman feels he has shown that nihilism and noncognitivism have difficulties, but that naturalism may or may not have worse ones. In order to help

---

[53] *Ibid.*, p.17. Harman is quite correct in recognizing that the vagueness we necessarily run into in describing some fact, or in attempting to determine the truth of some statement, does not make the fact at issue any less a fact or any more a relativistic or subjective one: vagueness, that is, implies neither nihilism nor relativism. However, it is not entirely clear why the question of vagueness is under discussion here at all. The vagueness issue pertains, as Harman rightly illustrates in his discussion of the door left slightly ajar, to those situations where certain sorts of instance (in this case, instances of a door being in a particular position relative to its frame) can be arranged along a smooth continuum between two opposing properties (in this case, the property of being open and the property of being closed) such that any instance along the continuum is meant to be correctly describable by exactly one of the opposing property-terms, and where some real or hypothetical instances (in this case, the slightly ajar door) lie in a location on the continuum where they cannot straightforwardly be assigned either property to the exclusion of the other. However, none of the cases Harman has been discussing – the case of the doctor contemplating an involuntary organ transplant in Room 306, the case of a good meal, the case of a good watch or a good teacher, etc. – has to do with vagueness. It appears that Harman is confusing instances of vagueness with instances in which some term is ambiguous, and both of these with instances in which something is difficult to determine or specify. But these are all different things. It is true that, when we encounter a problem of vagueness in our attempt to determine whether something has a certain property, we find it difficult to determine or specify whether it has that property. However, the converse of this (on which Harman seems to rely) does not hold: the difficulty in determining whether the watch, or the choice to kill the healthy patient, is a good one is not identical with a vagueness issue, and in fact it does not seem that a vagueness issue arises in either case. Furthermore, an ambiguous term need not be vague (the word 'bishop' can mean the chess piece or someone holding that office in the church, but despite this ambiguity neither understanding of the word involves any vagueness); and a vague term need not be ambiguous: while it is unclear whether the slightly ajar door is open or closed, the *meanings* of these two latter terms are clear).

determine whether nihilism should be preferred to naturalism in ethics, Harman turns to a consideration of G. E. Moore's famous 'open question' argument. This argument, Harman says, is offered by "moderate nihilists" (i.e. noncognitivists) in order to show "that ethical naturalism involves a 'naturalistic fallacy'"[55]

The open question argument, as Harman presents it, runs as follows: ethical naturalists[56] hold that moral properties are natural properties. Hence, they seem committed to the view that any moral property (including the property of being morally the right thing to do) can also be defined as some simple or complex natural property N. So the ethical naturalist will hold that, in any given case, the morally right thing to do is just the action that has natural property N. However, the argument goes, it would always be legitimate to ask, in any such case, "While I realize that this proposed action has natural property N, is it morally the right thing to do?". The fact that this is an open question is meant to show that moral rightness is not the same as N, regardless of what natural property is given as the definition of N. Therefore, the argument concludes, ethical naturalism must be false.

---

[54] *Ibid.* Since Harman admits that he has not determined whether there is a strong case to be made against functionalism, it seems to follow that he cannot have presented a very strong one against it. This point is worth bearing in mind in what follows.

[55] *Ibid.*, p.18.

[56] In this connection, the reader should be aware that Harman's characterization of ethical naturalism on p.18 – "An ethical naturalist holds that there are moral facts and that these can be 'reduced' to natural facts" – is not an accurate representation of the position as commonly understood today. There are both reductive and non-reductive versions of naturalism discussed in the contemporary literature, and both types will be dealt with in Chapter 6 of this dissertation. However, it is also true that explorations of non-reductive naturalism came to be developed more fully after the publication of *The Nature of Morality* (and largely, perhaps, in response to it – see, for example, Sturgeon 1986a and 1986b), which makes Harman's failure to consider this possibility more understandable.

Harman points out that there is a well-known objection against this argument: it seems to beg the question against ethical naturalism[57]. An ethical naturalist who holds that ethics can be defined as some natural property would hold that such questions are closed, not open; so the argument simply assumes the falsity of the view it is meant to counter.[58]

Since Harman acknowledges that he has not presented a good reason for rejecting the functionalist version of naturalism[59], and since he also reaches the view that the open question argument fails as an adequate argument against any kind of naturalism, one might assume that he must either make a stronger case at this point or else abandon or weaken his position that there is a problem with morality. However, his own conclusion at this point is quite different: he claims that "we are now in a position to see more clearly the way in which ethics is problematic"[60]. In support of this claim, he attempts to draw a contrast between moral facts and two analogous types of potentially problematic facts: facts about the average American citizen, and facts about colours. Harman feels that Naturalists need have no problem with either of these sorts of facts; but he reiterates that,

---

[57] As Harman notes, this objection was first prominently articulated by William Frankena (Frankena 1939).
[58] *Ibid.*, pp.18-21. Harman closes this section with a strange characterization of utilitarianism as a form of radical 'redefinitional naturalism' which, Harman says, "is actually not a version of ethical naturalism at all" because it holds that "our moral terminology is so vague, unclear and confused that we would do well to replace it with better and more precise terminology" (p.20). Harman feels that radical redefinitional naturalism has difficulties (though not necessarily, it seems, clearly decisive ones) in that the burden falls on the radical redefinitional naturalist to show that ordinary moral terminology is confused; and furthermore, he feels that it would be difficult for the radical redefinitional naturalist to avoid the temptation to "cheat, using 'ought' sometimes as he has defined it and at other times in its ordinary sense." (p. 21). The best way around this difficulty, Harman suggests, is to "dispense altogether with moral terminology in favor of utilitarian terminology and, instead of talking about what people ought to do, talk instead about what would satisfy the most interests. But that would be to give up any pretense of ethical naturalism and reveal that you have adopted extreme nihilism" (*Ibid.*).
Since Harman's interpretation is deeply questionable here (it seems highly unlikely that many utilitarians would characterize themselves as engaging in radical redefinitions of goodness, let alone that they would think of themselves as nihilists!), and since Harman presents no other potential versions of radical redefinitional naturalism, I have chosen not to dwell further on this option.
[59] See, again, footnote 27.
[60] *Ibid.*, p.21

by contrast, "ethics remains problematic" for naturalists because if, as reductive naturalists hold, "moral facts might be reduced to facts about interests, roles and functions", this reduction "would have to be complex, vague and difficult to specify"[61]. However, he recognizes that "reduction of facts about colors is also complex, vague, and difficult (probably impossible) to specify".[62] So what is especially problematic about moral properties?

The crucial difference between colours and moral properties, according to Harman, is that

> [e]ven if we come to be able to explain color perception by appeal to the physical characteristics of surfaces, the properties of light, and the neurophysiological psychology of observers, we will still *sometimes* refer to the actual colors of objects in explaining color perception, if only for the sake of simplicity... We will continue to believe that objects have colors[63] because we will continue to refer to the actual colors of objects in the explanations that we will in practice give. A similar point does not seem to hold for moral facts. There does not ever seem to be, even in practice, any point to explaining someone's moral observations by appeal to what is actually right or wrong...".[64]

Harman concedes that his other example does not fare so well in this regard: "facts about the average American citizen never seem to help explain observations, even in practice...

---

[61] *Ibid.*, p.21. As I hope my discussions in the previous footnotes have made clear, however, none of these things seems to entail any problem for naturalism; and Harman has yet to offer an argument as to why naturalism (or even, more specifically, reductive naturalism) is threatened by this complexity, vagueness, etc. A reduction that is "complex, vague and difficult to specify" is – as Harman's own use of the word 'reduction' here indicates he realizes – still a reduction.

[62] *Ibid.*, p.22

[63] It is interesting that Harman says "we will continue to believe that objects have colors" and not "we will continue to believe that colors are reducible": perhaps this shows that his intended target here in ethics is not reductive naturalism, but naturalism or non-naturalism *tout court*. It is not clear how this could be squared with the rest of what Harman says here, however; so I will assume that this is simply a lapse.

[64] *Ibid.* It is curious that Harman, despite the fact that he has just discussed Frankena's point that the open question argument seems to beg the question against naturalism, nonetheless apparently commits the same error here. It is not clear why any naturalist would accept the premise that "there does not ever seem to be... any point to explaining someone's moral observations by appeal to what is actually right or wrong",

But there is this difference. We can give a *precise* reduction of facts about the average American citizen; we cannot for moral facts".[65] So both of the two alleged failings of moral facts – the difficulty of accounting for them by reduction to simpler natural facts, and their inability to 'explain observations' indirectly – turn out to be shared by two other types of facts that Harman holds to be entirely legitimate (i.e. colour facts and facts about the average American citizen, respectively). What gets them off the hook, for Harman, is that neither of them has *both* of these alleged failings.

Emotivism

Having discarded ethical functionalism as unworkable in his second chapter, Harman returns to what he calls 'moderate nihilism' in his third. The form of this view most widely discussed in the 20th century was emotivism, the view that "to value something is not just to have a belief about something; it is to have an attitude toward something. It is to be in favour of something. To value something is to be in an emotional state, not a cognitive state".[66] According to this view, as Harman presents it, disagreeing on some

---

so this reasoning would be ineffective against a naturalist for precisely the same reason that the open question argument is.

[65] *Ibid.*

[66] *Ibid.*, p.28. One important advantage of emotivism and other forms of expressivism or projectivism (according to which moral utterances and commitments can be seen as, in part, expressing and endorsing some attitude or value) is that they may be able to make more sense of why holding a moral position seems to entail having a certain motivation or reason to act in accordance with it. Many of the criticisms leveled against these views, including those offered by Harman, turn in some way on the difficulties one faces when one denies that moral utterances are statements. But *must* these positions involve such a denial? That assumption turns out to be dubious, as Simon Blackburn points out in *Spreading the Word*: "...it would be wrong to infer that *no* description is given from the fact that an attitude is *also* expressed" (Blackburn 1984, p.169). Harman, as will be seen, strays into the middle of this problem by failing to distinguish these two understandings of emotivism; and his treatment of the position suffers as a result. In the passages just quoted from Harman, for instance, there is a tension between the statement '...to value something is not ***just*** to have a belief about something; it is to have an attitude toward something' (which implies that valuing is having a belief plus something else), and the following statement that '[t]o value something is to be in an emotional state, ***not*** a cognitive state' (which implies that valuing is not having a belief at all).

matter of value is quite different from disagreeing on a matter of fact. Since matters of fact (such as the statement that the Oregon Taxpayers Union kidnapped Sally Jones, to use Harman's example) can be presented as statements, and indeed as the conclusions of reasoned arguments, it is possible to resolve factual disputes through argumentation. However, two people could agree on all the facts and still disagree in their values, so (according to emotivism as Harman understands it) their moral disagreement will persist in spite of their agreeing on all the facts. Conversely, our attitudes and values are not genuinely expressible in statements (we should not be led astray, emotivists tell us, by the fact that 'Sally Jones ought not have been kidnapped' looks as though it were a statement), and "as soon as we agree in attitude, our moral disagreement has ended, no matter what differences in belief remain".[67]

Having set out this view, Harman considers whether an emotivist could maintain that morality is objective, or whether he or she must be a moral relativist. He concludes that an emotivist might resist relativism if he or she believed that there were some uniform human values.[68] Harman spends the rest of his third chapter considering only the version of emotivism that involves the absence of uniform values. He presents some well-known objections against this more narrow view. First, it follows from this sort of emotivism that our moral utterances are neither true nor false, but this "conflicts with ordinary ways

---

Nonetheless, it should be granted in fairness to Harman that he is not the only philosopher who runs into these sorts of problems. Alexander Miller, for instance, actually presents a longer version of the Blackburn quote just given in order to warn readers of the error (Miller, p.38), but then neglects his own warning just two pages later ("According to emotivism, when I sincerely utter the sentence 'Murder is wrong', I am not expressing a belief or making an assertion, but rather expressing some non-cognitive sentiment or feeling, incapable of being true or false" (Miller, p.40)) and then, surprisingly, continues to attack this apparent straw man for the rest of the chapter.

[67] *Ibid.*, pp.27-28.

[68] *Ibid.*, p.30

of speaking".[69] Second, "[t]he emotivist holds that moral judgments are expressions of feeling. But expressions of feeling do not depend on reasoning from general principles, nor do they require defense by appeal to principle, whereas moral judgments do depend on reasoning from general principles and do require defense by appeal to principle"[70]. Third, we sometimes hold that people are mistaken regarding moral issues; but this seems to imply that our moral thinking can involve the possibility of truth or falsehood (i.e. that it can be 'truth-apt').

However, Harman feels that emotivists have powerful responses to these concerns. In response to the first objection, Harman sees the emotivist as starting out by taking a deflationary account of truth, "according to which 'it is true that S' means no more than 'S'"[71]. If the anti-emotivist presses the point and insists that emotivism is therefore implausible, since it implies that there are only moral truths "in a manner of speaking" and not as "part of the order of nature", the emotivist "may have to say that our ordinary ways of talking and thinking are simply mistaken".[72] The second and third objections, by contrast, are dealt with in a fashion that Harman attributes to Hume: our moral feelings can be depicted by emotivists as stemming from general principles, and as being liable to error, in that they sometimes depend on our beliefs (which are unproblematically truth-apt). For instance, "if the revolutionary had initially judged that the Jones kidnapping

---

[69] *Ibid.*, p.34; but as I have explained in footnote 40, it is not clear that this follows from emotivism.
[70] *Ibid.*, p.37. Throughout this discussion, Harman refers to moral utterances as 'expressions of feeling'. The reader might be excused for wondering whether the force of this objection might have been blunted had Harman maintained his characterization of emotivism from the beginning of the chapter, where it was depicted more broadly as the view that "moral judgments express the speaker's emotions, feelings, attitudes, intentions, or more generally, norms and values" (*Ibid.*, p.28)
[71] *Ibid.*, p.35.

was a good thing, because of his beliefs about its consequences, and then came to believe that the consequences were going to be the exact opposite, he would then suppose that his initial moral judgment had been mistaken, because the beliefs on which it had been based had been mistaken".[73]

Since Harman seems to feel that these responses are adequate against the three major objections he considers, one might expect that emotivism will receive a favourable final verdict from him. However, Harman in the end becomes disappointed with the "triviality of emotivism", pointing out that it "appears less radical than it seemed at first".[74] "We may begin to wonder", he says by way of conclusion, "how emotivism differs from simple common sense"[75].

---

[72] *Ibid.* Those who are more familiar with the current debate over noncognitivism might wonder whether the strategies given above for that view are really the best ones. See, for instance, Blackburn 1984, 1993; Gibbard 2003, 2006; and Schroeder 2007, 2008 and 2009.

[73] *Ibid.*, p.38. Whatever the merits of these responses to the anti-emotivist objections Harman raises, they seem to be based on a rather oversimplified version of Hume's ethical view: see, for instance, Kemp Smith, Chapters VI and IX. Moreover, Harman's extensive discussion of Hume in this chapter unfortunately depicts Hume as holding two inconsistent views, though Harman seems not to notice that he has presented him in this way. At some points, Harman says that "Hume thinks that our moral feelings involve, among other things, a general sympathy for others and, consequently, our moral feelings are often based on beliefs about the general tendencies of acts to affect people's interests. We morally approve a course of action if we believe that course of action falls under principles that, if generally followed, would promote human happiness and diminish misery. This is not the only moral concern that we have, according to Hume, but it is one of them..." (*Ibid.*, p.39). At other points in the same chapter, however, Hume is represented as holding that there are no general moral feelings at all: "...suppose that human nature imposes no real constraints on basic values and that practical reasoning creates no new values and that practical reasoning creates no new values but only enables you to pursue your values in the light of whatever information you have... [I]n other words, suppose that practical reason always is and ought to be the slave of the passions. This would be to accept the sort of conception of practical reasoning we find in Aristotle (384-322 B.C.) or Hume" (*Ibid.*, pp.30-31). Harman is not alone among contemporary metaethicists in providing caricatures of Hume instead of the real thing: even those who present themselves as authorities on Hume's moral theory have been caught in serious lapses (see, for example Sturgeon's careful exposure of one such lapse in Sturgeon 1980, p.126). Nevertheless, Harman's case is doubly unfortunate since he has employed two caricatures of Hume that actually contradict one another.

[74] *The Nature of Morality*, p.39

[75] *Ibid.*, p.40. Presumably, this conformity with simple common sense is just what emotivists hope to have achieved. One might wonder just what Harman is looking for in a moral theory: surely, he is a difficult customer to please. As we have seen, he rejects moral nihilism because so much of it goes against ordinary ways of thinking and talking. He then rejects utilitarianism because, while it nearly always conforms to

However, Harman emphasizes that emotivism might look very different when held together with the view that there are universal moral values. He considers emotivism in this light in the fourth chapter of his book.

The Ideal Observer Theory

Emotivism, when combined with the view that there are universal moral values, would (Harman feels) give us the ideal observer theory.[76] The ideal observer theory is a version of constructivism. According to constructivism, moral judgments are truth-apt (i.e. there can be true or false moral judgments), but the moral facts referred to in these judgments are 'stance-dependent': to say that some moral judgment is true is just to say that some real or idealized person or group of persons would, when considering the question at hand in either an actual or idealized setting, arrive at this same judgment. Constructivists

---

ordinary ways of thinking and talking, it strongly varies from these ways in a few notable cases (e.g. the case of the involuntary organ transplants discussed earlier). Finally, he looks askance at emotivism precisely *because* he feels it conforms too closely to ordinary ways of thinking and talking.

[76] Harman's reasons for thinking this are not fully clear. The ideal observer theory, as will be discussed presently, entails the view that there are moral facts, whereas throughout much of his third chapter Harman has identified emotivism as a view that there are, strictly speaking, no moral facts. It is true that, particularly toward the end of that chapter, Harman comes to think of emotivism as consistent with the view that there are moral facts, and as otherwise consistent with basic common sense. But this still does not explain why emotivism has any closer connection with the ideal observer theory than 'basic common sense' does; and if all he means to say is that the ideal observer theory follows from a combination of basic common sense and the view that there are universal moral values, it is not clear why he mentions emotivism at all. The best I can make of this is that Harman sometimes thinks of emotivism as the view that "moral opinions are feelings" (e.g. p.41), whereas he seems to feel that the ideal observer theory is "a version of ethical naturalism that takes moral judgments to involve judgments *about* feelings" (p.42 – emphasis in original). However, this connection seems quite weak.

Nor is Harman's account of the relationship between emotivism and the ideal observer theory necessarily self-consistent throughout his fourth chapter. For instance, his title for the chapter is 'Emotivism *as* the ideal observer theory'; then, on p.41, Harman claims that "emotivism *is not easily distinguished from* a version of ethical naturalism called the 'ideal observer theory'"; then, on p.46, a discussion entitled "Emotion *versus* the ideal observer theory" begins; then, later on the same page, the reader is told that "[e]motivism… *leads to* the ideal observer theory"; etc. (italics mine). Hence, it is not possible for any discussion of Harman's views on the relationship between these two theories to be both accurate and consistent, and I apologize for this.

vary along a spectrum. Some tend toward the more subjectivist side, and hold that any moral claim is true so long as it is sincerely believed by some (perhaps relevantly specified) agent or appraiser (or perhaps by some group of these agents or appraisers). Other constructivists tend toward the objectivist side, and deny that moral judgments can be rendered true merely by being held sincerely by some person or persons. According to these latter constructivists, what renders a moral judgment true is only the fact that it would be assented to by a suitably informed, and suitably objective, moral observer contemplating the matter at hand. These objectivist sorts of constructivists are known as 'ideal observer' theorists.[77]

At this point, Harman seems to feel that he has reduced the field to two major possibilities: subjective constructivism (i.e. moral relativism) and objective constructivism (i.e. the ideal observer theory).[78] All the other options have now been rejected on the grounds that they fail in their attempts either to eliminate moral properties altogether (radical nihilism), to deny their factuality to some degree (moderate nihilism), to reduce them to more simple natural properties (reductive naturalism), or to defend them against the problem of observation and explanation (simple naturalism). In his fourth chapter, Harman considers the strengths and weaknesses of the ideal observer theory. Then, in the following chapters, he begins to construct his positive case for relativism.

---

[77] See Chapter 2 of Shafer-Landau (2003b) for an excellent discussion of the varieties of constructivism and some interesting challenges they must face.

[78] Harman does not, however, use the term 'constructivism' himself or group these two options together as I have done.

The ideal observer theory, as Harman defines it, is the view that "something is wrong if and only if an impartial spectator or ideal observer would disapprove of it; an ideal observer is defined to be disinterested, well informed, vividly aware of the relevant facts, and so forth. If correct, this analysis yields a naturalistic reduction of moral facts".[79]

The reduction Harman has in mind is of moral properties to some dispositional natural properties. A dispositional property is one that stands in a particular sort of relationship to some occurent property – to a tendency of some object to act or be acted upon in a certain way. When placed in water, sugar has the property of dissolving. Dissolving is therefore an occurent property of sugar in water. But we can also say of sugar that it has the property of being the sort of thing that would dissolve *if* placed in water: in other words, we can say of sugar that it is water-soluble. Being water-soluble is hence a dispositional property of sugar. Harman gives another example: redness. Redness can be seen as a dispositional property if one accepts the view that an object is red if and only if a normal perceiver would see it as red under optimal viewing conditions.[80] As Harman's definition of the ideal observer theory in the previous paragraph makes clear, the theory attempts to account for moral properties dispositionally in a way that is very similar to this account of redness. The benefit of a dispositional understanding of morality, Harman thinks, is that it can allow us to account for moral properties in a reductive naturalistic manner, but without (unlike functionalism) running into the ostensible difficulties already presented.

---

[79] *The Nature of Morality*, p.44
[80] *Ibid.*

However, Harman feels that not every ideal observer need have the same reactions to moral issues.[81] This is a peculiar feature of Harman's account of the ideal observer theory, and it seems *prima facie* to be inconsistent with the very idea of an ideal observer: wouldn't any two completely disinterested and impartial observers with access to all the relevant facts, etc., necessarily have the same reactions as one another? But Harman's view (as becomes clear later on) is that radical moral disagreement of a sort that could not be rectified even by giving ideal observer status to the disputants is possible; so there is perhaps no inconsistency in his somewhat unusual characterization of the view.[82] Harman feels that the possibility of disagreement among ideal observers is implausible in some cases (like the case of colour), but importantly plausible in the case of morality. In the morality case, Harman feels, we might have to imagine different ideal observers from Peking, Moscow, and New York, and for that reason end by talking not about wrongness per se but rather about 'Moscow wrongness' and 'New York wrongness', just as we talk about solubility in water versus solubility in oil.[83]

---

[81] *Ibid.*, pp.44-45

[82] However, this still may show up a possible inconsistency in *Harman's* view, depending on how he is to be interpreted. On p.30, he claims that "[i]t is compatible with emotivism to assume that there are moral facts, if you also assume that there is enough uniformity in human nature. An emotivist can even be an ethical naturalist. For example, if an emotivist believes that moral approval and disapproval derives [sic] from a universal sympathy we feel for others, then he might adopt a naturalistic definition of moral wrongness..." Harman seems to suggest by this and other passages that an emotivist can avoid nihilism and be a naturalist only by assuming the existence of this sort of uniformity. Harman is also, as we have seen, clear on the point that the ideal observer theory is naturalistic and non-nihilistic. And the very title of the Harman's chapter of the ideal observer theory makes clear, as we have seen, that he regards the ideal observer theory as a variety of emotivism. It follows from all these claims together that an ideal observer theory implies the existence of some such uniformity; for if it were held together with the view that no such uniformity existed, a contradiction would be entailed. Therefore, if this reading is correct, it seems Harman holds that the ideal observer theory entails the existence of uniform moral values. But this really does appear inconsistent with the possibility of two ideal observers who suffer from deep moral disagreement.

[83] *Ibid.*, pp.45-46. I will discuss this issue, and Harman's intended contrast between the colour case and the morality case, at great length in Chapter Three. To avoid redundancy, I am leaving until then a discussion of some further important points in Harman's treatment of the ideal observer theory.

In the end, Harman gives the same qualified endorsement to the ideal observer theory that he gave to emotivism: he does not think there is a strong argument against it, but he does feel that it is merely trivial in the final analysis. The reason is that the proponent of the ideal observer theory holds that something is morally wrong if and only if an ideal observer would morally disapprove of it, but "[t]o define 'X is wrong' as 'an ideal observer would feel moral disapproval of X' is very much like defining 'X is wrong' as 'an ideal observer would think that X is wrong'".[84]

However, Harman does think there is something interesting to be learned from both emotivism and the ideal observer theory:

> Emotivism is not completely trivial; it does make one good point. To think that something is a good thing, the right thing to do, what ought to be done, and so on, are ways of being in favor of something… The ideal observer theory is not completely trivial either. It brings out the point that moral beliefs involve a claim of impartiality. This is sometimes put by saying that moral judgments are "universalizable" – if you believe that X is wrong, you must also believe that anything relevantly like X would also be wrong, no matter what people were involved, no matter how they were involved… Combining the two points, we see that you can have moral beliefs only to the extent that you are susceptible to a kind of unbiased motivation. For, if you are in favor of something, you are to that extent motivated to try to bring it about, and to be morally in favor of something is to favor it in an unbiased way.[85]

This sets the stage for Harman's discussion of moral motivation, which (as will soon be seen) he connects with the view that morality involves universalizable claims of a sort that could be made by an ideal observer who lacks universal moral values. These two ideas together, as we shall see, form the basis for Harman's own moral theory.

---

[84] *Ibid.*, p.49. However, this reasoning seems somewhat dubious, as I will argue in Chapter Three.
[85] *Ibid.*, p.52

The Diversity of Moral Laws

Harman begins his fifth chapter with an argument supporting the existence of moral laws. He discusses some cases that point to the initial ethical plausibility of the principle of double effect, according to which it is worse to harm someone while intending to do so than it is to harm someone to a similar degree if the harm comes as the unwanted, but antecedently recognized, result of acting with the intention to do something else.[86] From this, Harman argues as follows:

> The evidence that our morality incorporates principles like the principle of Double Effect is just that some such principles are needed in order to account for our intuitive moral judgments. Of course, whether or not our morality actually involves this principle, rather than another, cannot be established by this brief discussion. But this discussion should be adequate to indicate that our moral system is a complex one. Moral reasoning often involves fine distinctions of a sort that have the characteristic of law. So, if morality is not a delusion, there must be a moral law, or something like it.[87]

He also notes that the view that there is something akin to a moral law receives support from common ways of talking.

Harman feels that some important results follow from the existence of moral laws: "Now, it is the mark of law that it has sanctions. Who then enforces the moral law? Is it society? Is social pressure the sanction?"[88] He answers these last two questions

---

[86] An excellent discussion of the doctrine of double effect appears in Alexander and Moore 2007.

[87] *Ibid.*, p.59. The logic underlying this argument seems to be of the form 'Many instances of X seem to have characteristic A, and A is also frequently a characteristic of instances of Y. Therefore, all instances of X must also be instances of Y, or else of something like Y'. If this accurately encapsulates Harman's reasoning here, and it seems that it does, then the argument is fallacious. Counterexamples abound – for example, 'Dogs give a great deal of pleasure to many of their owners, as do coin collections: therefore, dogs must either be some kind of coin collection, or else very similar to coin collections'.

[88] *Ibid.*, p.60. Those readers who are conversant with the difficulties faced by creationist arguments for the existence of God will be particularly familiar with the error made here: Harman has committed the fallacy of equivocating between two senses of the word 'law'. The sort of law Harman feels he has established in his discussion of the doctrine of double effect (as he understands the doctrine) is simply a regularity: we

positively, but with qualifications: there can be, he admits, personally held moral laws that are not the result of social pressure. However, he feels that these laws are best understood as the result of the internalization of the commands of one's parents:

> For consider how morality is learned. At first, "X is wrong" or "X is bad" means something like "Mommy and Daddy do not want me to do X". The child's concept of his parents is of omniscient omnipotent beings whose love is all important to him; he tries to please them so that they will love him and not punish him. But he is not always sure what they want. He needs a way to predict their desires. So he uses his imagination. He tries to imagine how his parents would react to various things he might do... He comes in this way to pretend in his imagination that he is a parent and then he tries to act so as to please this imagined pretend parent. To use the Freudian terminology, the child develops a "superego" or an "ego ideal".[89]

Harman claims that "we do not need to assume anything about psychoanalytic theory"[90] in order for his point to hold. However, it soon becomes clear that Harman's major argument for the diversity of moral laws does indeed depend on the truth of Freudian psychology. His argument is that since different people's parents make different demands, and since moral laws derive from our superegos, and since our superegos are formed by an internalization of our parents' demands, there must be diversity in moral laws. A glance at the way Harman presents his argument should make the dependence of his case on Freudian psychoanalytic theory clear:

---

tend to find actions less blameworthy when the harm they produce is not directly intended, so there is a fairly regular correlation between blameworthiness and intended harm. But this not the same as the sort of law that can only be established by a personal lawmaker or political body. (The creationist version of this is to argue directly from the fact that there are observable regularities in nature to the conclusion that there must be a conscious, divine lawmaker). Nor, as Shafer-Landau points out, does the argument improve when moved from the scientific to the moral arena, where the laws are normative instead of merely descriptive: laws of logic and reasoning are also normative, but we go wrong if we assume from the existence of modus ponens, say, that some person, deity or politically powerful social group must have made up that rule of reasoning (see Shafer-Landau 2003b, p.77)

[89] *The Nature of Morality*, p.60

> [T]he child comes to internalize principles of morality. Having formed a superego, the child then tries to retain its love and avoid its punishment. Thus, "X is wrong" comes to mean something like "X is incompatible with the demands of my superego". Of course, the child himself is his superego. He is his own superego to the extent that he adopts the role of his own parents. In trying to act morally, then, the child is trying to obtain his own love and respect. This continues after the child has grown up and become an adult. Inasmuch as the adult tries to act morally, he tries to act so as to retain the love and respect of that part of himself that plays the role of his imaginary parents... Your superego is an idealized version of your parents, who in the first instance are the people who transmit society's demands to you. Your parents' demands may not coincide directly with the demands other parents make of their children and, furthermore, your idealized version of your parents may differ from the actual version in any number of ways... [91]

However, one's parents' demands, and one's interpretations of their demands, are not the only things that will determine the sort of morality one accepts. There is also the issue of whether one identifies one's internal 'imaginary parent' with God:

> If [one] does, his morality will take the form of a divine law morality. This may seem to have the disadvantage that, if there is no God, his morality will rest on a fiction. But the nonreligious adult seems to be in an even worse position. His morality rests on a fiction whether or not God exists, since his morality rests on the pretend commands of an imagined parent.[92]

Our superegos, and those factors that lead to their development, are "socially useful... For morality is socially useful and the development of the superego is what makes people

---

[90] *Ibid.*

[91] *Ibid.*, pp.60-61. Harman's decision to base this important argument on Freudian theory seems to be an unfortunate one, since Freudianism is under something of a cloud. See, for example, Erwin 1995; Macmillan 1997; Crews 1998; and Cioffi 1999. Serious criticisms of Freud's work, including philosophical ones, were well-known before the publication of *The Nature of Morality* (e.g. Popper 1965; Wollheim 1974), so it is unclear why Harman does not address these concerns. Still, it must be admitted that not everyone agrees that Freudian theory is entirely untenable. See, for example, Elliott 1998 for a contrary view.

[92] *The Nature of Morality*, pp.61-62. The logic behind this last sentence is somewhat unclear: presumably, if God exists, his commandments are not "pretend commands of an imagined parent".

act morally (when they do). The superego is therefore a socially useful fiction. But it is a fiction – a pretense. That is the point. Morality rests on pretense and fiction".[93]

Nor does Harman feel that belief in morality can be justified even for the purposes of social utility: "[I]t is no defense of morality and of the fiction on which it rests to say that morality is socially useful. Given the moral point of view, we have a reason to retain anything that is socially useful. But it is circular reasoning to defend the superego in this way. That is to say no more than that we have a moral reason to preserve morality. Since this moral reason rests on a fiction, it is only a fictional reason".[94]

Since morality cannot legitimately be defended on the grounds of its social utility, and since it rests on nothing more than 'pretense and fiction', one might suppose that Harman plans to do away with morality after all. However, he shies away from this conclusion also, on the grounds that "the habit… of pretending to be an idealized parent is, for

---

[93] *Ibid.*, p.62. It is not at all clear why Harman feels it follows from the reasons he provides that the superego (if it existed), or morality as a whole, must be a 'fiction' or a 'pretense'. His other claims would, however, if true, at least support the view that morality is socially useful. Still, that seems to be a considerably less interesting or controversial claim.

[94] *Ibid.* This seems somewhat difficult to square with Harman's apparent disagreement with radical nihilism. More obviously, though, he seems to have committed a serious error in this argument. He claims that it is circular reasoning to say that morality should be preserved for the purposes of social utility, on the grounds that to claim that something that is socially useful ought to be preserved is already to employ moral values. But these grounds are faulty. The reason why is, interestingly, made clear by Harman himself just three pages before this passage. There (p.59) he contrasts "the 'ought' of reasons, as when we say that a thief ought to wear gloves, meaning that he has a [prudential] reason to do so" with "the moral ought, as when we say that the thief ought not to steal". Clearly, by this standard, one can say that socially useful things ought to be done if one uses what Harman calls "the 'ought' of reasons", since one stands to benefit from the strengthening of society. Indeed, this is the apparently uncontroversial premise used by all social contract thinkers. Therefore, one does not need to use the moral 'ought' to defend the superego, and Harman's argument collapses.

various reasons, not an easy habit to modify or get rid of… There is no easy way to abandon your superego and therefore no easy way to abandon morality".[95]

Regardless of whether one *can* abandon morality, however, one might wonder whether one has good *reason* to be moral. If, as Harman imagines, one were hypothetically able to detach oneself from one's superego, could that be a rational thing to do? Harman's answer is that this question might not make sense: our own standards of judgments arise from the superego, so it seems inevitable that we will approve of them unless they are internally inconsistent.[96]

All this leads Harman to the view that those moral rules we happen to follow and believe to be most important are contingent matters, and depend on some combination of our society, our parents, and our internal representations of these authorities. Since they are nothing more than this, Harman continues, it seems that cases of interpersonal or intercultural moral conflict are in principle unresolvable and there is no objective sense in which anyone's moral view can be objectively better than anyone else's. In that case, supposing that Harman's arguments up to this point are all good ones, moral relativism would follow. However, Harman wishes to give objective morality the benefit of the doubt: he admits that it could also be that the principles we discover in various people's moral thinking "contain universal principles that should be accepted by any rational

---

[95] *Ibid.*, p.62. Interestingly, Harman reinforces this point by tying his argument further to Freudian psychoanalytic theory: "The superego will not disappear quietly. It represents a habit of acting that is embedded in your personality structure, which serves deep psychological needs by allowing you to identify with your parents, *a habit defended by psychological mechanisms of repression that make it very hard to come to grips with*" (*Ibid.*) [Emphasis mine].

[96] *Ibid.*, p.63.

being", even though these principles might "go beyond what is required by rationality alone".[97] But is this the case? It is to this matter that Harman turns next.

Knowledge and motivation

Harman begins his sixth chapter by attempting to bridge the gap between the moral 'ought' and moral knowledge. He argues as follows: "[W]e would not morally blame someone for doing something that he could not have known was wrong. So, for example, we assume that it cannot be true that someone ought morally to have helped you in your time of need if there was no way in which he could have known that you needed help".[98] As will be seen, much of great importance hangs on this brief bit of reasoning.

Harman emphasizes that it is not only non-moral facts (such as the fact that someone is in need of help) that must be known in order for one to be morally obligated: anyone who does not know that some moral principle is correct is not bound by that principle. Harman recognizes that such a strong and startling claim needs argumentative support. He attempts to provide it in the following passage:

---

[97] *Ibid.* The meaning of the term 'rationality' in this discussion is admittedly somewhat obscure, and Harman is far from the only metaethicist who uses this ambiguous word without providing much in the way of explanation for it. I will devote some space in Chapter Four to unpacking what 'rationality', 'reasons' and 'reasonable' could mean in these contexts.

[98] *Ibid.*, p.65. However, it is not clear that this follows: there does not seem (to me, at least) to be anything obviously wrong with saying, for instance, 'I'm very sorry that I didn't help you when you fell and broke your arm: no doubt, I should have helped, and I would have if I had only understood what was going on. But I sincerely believed that what you were doing was a part of your performance art piece, so I didn't realize until later that you needed assistance. I see now that I ought to have helped'. What is particularly interesting about Harman's account is that it implies that one can never learn that one was remiss in one's ethical duties, and hence that one is always fully aware of one's own ethical obligations. That would imply that it is never informative to tell someone what he or she ought morally to do. But that seems (again, to me) to be implausible, and hence at least to require some good argument in its defense.

> If any form of moral nihilism should turn out to be correct and moral knowledge should turn out to be impossible, so that no one could ever know what he ought morally to do[99], it would follow, according to our usual ways of thinking, that no one could ever do anything morally wrong. Even the most moderate form of nihilism, which concedes that there are moral facts but denies that we can know them (see Chapter 3) would therefore completely undermine morality as we ordinarily think of it.[100] Either moral knowledge is possible, or morality as we ordinarily conceive it is, in the words of the German philosopher Immanuel Kant, "a vain and chimerical notion".[101]

Harman also feels that he has established, in his first chapter, that moral knowledge cannot arise from observations (this seems to derive from his contention in that chapter that moral *facts* cannot help *explain* observations); and that he has established, somewhere in his third and/or fourth chapters, that to know some moral fact is to be motivated by it. Harman now combines these views with the claim that the only moral

---

[99] Given that Harman has defined moral nihilism as "the doctrine of no moral facts", etc. (p.11), this is a strange thing for him to say: if some form of it is correct, then there is never *anything* that anyone ought morally to do.

[100] However, Harman seems to misremember what he discussed in Chapter 3 of his book. No discussion of moral skepticism – of the view, that is, that we can never know what we ought morally to do – appeared there, and moral skepticism does not seem to follow from emotivism in any case. In fact, Harman devoted three pages (pp.37-39) of Chapter 3 to defending emotivism against the important objection that it seems to imply that we are never mistaken about morality – in other words, that our moral knowledge is always perfect!

The most charitable interpretation of Harman's reasoning I can concoct is that, since emotivism involves the claim that moral utterances are not statements of fact, there are not strictly speaking facts about morality that one can know. However, as Harman himself seemed to recognize in Chapter 3, there are many ways for emotivists to get around this issue without making the odd claim that nobody knows anything about morality. Neither Ayer (Harman's own example of an emotivist) nor any other prominent advocate of that view seems to have *thereby* felt committed to the radical thesis of skepticism. Ayer, for instance, talks about our relationship with others who "have received the same moral education as ourselves" and with whom we can, therefore, have fruitful moral debates (Ayer, p.111). But what could Ayer mean by a moral education, or a successful moral debate, if he believes that nobody can know anything at all about what he or she should do?

[101] *Ibid.*, pp. 65-66. It does seem true that morality (in any sense in which we could discuss its contents, at least) would be a 'vain and chimerical notion' if *no* moral knowledge were possible; and we seem entitled to the *prima facie* assumption that morality is not just a vain and chimerical notion. However, this does not help Harman make the point he needs to establish, which is that *all* moral knowledge must be immediately and practically available to all moral deliberators.

On a historical note, it is also not clear that this is an accurate rendering of Kant's views. Harman does not cite a section or page number from Kant, nor does he list the edition he consulted. The Cambridge Texts edition of the *Groundwork* contains the following sentence: "Here mere conformity to law as such, without having as its basis some law determined for certain actions, is what serves the will as its principle, and must so serve it, if duty is not to be everywhere an empty delusion and a chimerical concept" (Kant, p.15).

rules binding on someone are those known by that person, and reaches the conclusion that our moral motivations must exist independently of all observation-based knowledge. That is to say, all our moral motivations must be *a priori*.[102]

He next asks us to consider the case in which a "totally self-interested individual"[103] notices another driver, whose car has apparently broken down, trying to wave him down on a deserted road. Being totally self-interested, the driver decides not to stop: he is enjoying his drive too much to interrupt it, even briefly, to offer assistance. Harman wonders whether someone who is not entirely self-interested might succeed, through saying 'How would you have felt if *you* were the stranded motorist?' in trying to persuade the driver that he ought to have stopped. Harman considers the possibility that the driver might insist that only self-interest can give one a reason to do something; but then he reflects that we can sometimes convince such a person "that he has a reason to help, a reason not based on self-interest. It is striking that, by reasoning in this way, we can get a self-interested person to act for such a reason".[104]

Harman feels that an important tentative conclusion follows from this: first, the fact that one would have desired someone else to have stopped for one on the road had the situation been reversed may itself be an adequate non-self-interested reason for stopping. In other words, the selfish driver may, despite the fact that he feels no inclination to stop,

---

However, if, as it seems, this is the passage Harman has in mind, then it is not clear that moral knowledge is in fact under discussion.

[102] *Ibid.*, p.66

[103] *Ibid.*, p.69

[104] *Ibid.*, p.70. Of course, the reader might wonder how Harman knows that those on whom this strategy is successful are in fact *totally* self-interested, or indeed whether the fact that they can be persuaded by non-self-interested considerations might make it analytically true that they are *not* totally self-interested.

be *irrational* in driving past. Hence, if this conclusion is correct (and he will question it shortly), you can have a reason to perform some action even if you do not desire to perform it.[105]

As we have seen, Harman accepts the view that without moral knowledge, morality is vain and chimerical. He also feels, as we have seen, that moral knowledge is necessarily *a priori* and that it invariably provides moral motivation to the knower. He refers to all this, somewhat incorrectly perhaps, as "Kant's earlier argument that either moral motivation derives from reason alone, or morality is a vain and chimerical notion".[106] Should this argument be accepted? Not necessarily, thinks Harman: it all depends on whether fundamental moral principles are binding on everyone. If they are not, then a more relativistic picture emerges.

Morality and Reasons

Are there fundamental moral principles that are binding on everyone? Harman feels that the answer to this question must depend on whether we can have a moral obligation to do something we have no reason to do;[107] or in other words, on whether saying that somone morally ought to do something implies that that person has a reason to do that thing.

There is much of interest here, and many moves I believe to be suspect. However, since I will critically examine Harman's claims and arguments on these points much more fully

---

[105] *Ibid.*, p.70
[106] *Ibid.*, p.77

later on, I will resist the urge to offer my critiques of his treatment now. Instead, I will limit myself now to a statement of Harman's conclusion: he comes to believe, by the end of his seventh chapter, that one can have a moral obligation to do something only if one has a reason to do it. Since he also believes that the reasons people have for doing things are not universally shared, he concludes that morality consists only of principles "that are accepted by a group of people", and that they "are binding only on those who share them or whose principles give them reasons to accept them".[108]

Custom and Relativity

This conclusion naturally suggests the following possibility: since the moral commitments people feel bound by – and the things people take to be morally right and wrong – vary from community to community, perhaps actual moral obligations are similarly community-dependent. In other words, perhaps moral relativism is true. Harman considers this possibility in some depth in his eighth chapter.

He begins by discussing a problem he sees with the ideal observer theory: while the ideal observer theory may tell us *what* we ought to do, it will not tell us (Harman says) *why* we ought to do it. Harman uses a variant of the open question argument to make this point: if someone does not care whether his or her actions would be approved of by an ideal observer, then pointing out to him or her that an ideal observer would think some particular action best would not be helpful in getting that person to see that he or she

---

[107] Again, I acknowledge that the meaning of the phrase 'having a reason' here is unclear; but I will argue in Chapter Four that this lack of clarity plagues many metaethicists' discussions of moral reasons, including Harman's. For now, I regret that the reader must simply make what he or she will of the phrase.
[108] *Ibid.*, p.90

ought to perform it. That is to say, it may still be an open question to such a person whether the action the ideal observer theory endorses ought to be performed.[109]

Harman next sets out the internalist/externalist dichotomy, as he understands it:[110]

> Some say our concern [i.e. about doing the right thing morally] is purely internal to our moral beliefs, whereas others say that it depends at least in part on external sanctions. For example, a divine law theory says that right and wrong derive from God's law, which He enforces with divine sanctions; in this view, we care about right and wrong because we care about whether we are going to go to Heaven or to Hell. Emotivism, on the other hand, tries to explain our concern with right and wrong by identifying the thought that something is right with one sort of concern and the thought that something is wrong with the opposite sort of concern. Emotivism takes moral concern to be purely internal to our moral beliefs, whereas a divine law theory holds that it derives in part from external sanctions. Emotivism is a kind of *internalism*; the divine law theory is a kind of *externalism*.[111]

---

[109] *Ibid.*, p.91.

[110] In the two pages leading up to this quote, Harman raises the internalism/externalism issue in two different ways: first, he says that "[t]he ideal observer theory (for example) is defective in that it does not say why anyone **has a reason** to do what he ought morally to do" (p.91), and later, that "[i]t is true that we are not always **motivated** to do what we think is right. Sometimes we are tempted to do something else and sometimes we just do not care. But normally we do care about right and wrong and an ethical theory must somehow account for our concern" (pp.91-92). (Emphasis mine in both cases). The context suggests that Harman may well think these two matters are identical. Readers familiar with the contemporary meta-ethical scene may note that this appears to be a conflation of the motivational internalism issue (i.e. the issue of whether one can believe a moral claim without being able to be motivated by it) with the reasons internalism issue (i.e. the issue of whether one can have a reason not to do something if neglecting to do that thing would be entirely consistent with all the demands of one's motivational set). However, it should also be noted that the *locus classicus* for the reasons internalism issue – Bernard Williams' 'Internal and External Reasons' – was published slightly after *The Nature of Morality*. Regardless of where the fault for these terminological discrepancies lies, it is problematic to run together the ideas of reasons internalism and motivational internalism, since they are two different things. Some problems arising from this conflation will become clear shortly.

[111] *Ibid.*, p.92. The reader will note from these definitions that Harman not only runs together the issue of whether someone has a reason to do something with the issue of whether that person can be motivated to do that thing, but also runs together both of these with the issue of whether that person is 'concerned with' that same thing (or perhaps with the issue of whether the person is concerned with doing the right thing in general). As will become clear, Harman also seems to run together the issue of whether someone has a reason to do something with the issue of whether that person ought to do that thing (which, depending on the way we are meant to understand the phrase 'has a reason', might be legitimate – more on this in Chapter Four). Taken together with Harman's definitions of internalism and externalism (which surely are meant to be mutually complementary), all this makes clear that Harman is begging the question by defining his terms in this way. For these definitions, for the reasons mentioned above, already imply that the only things one morally ought to do must be either things that arise from our own personal concerns (if one is an internalist) or else things that we would be punished for not doing by some force external to us (if one is an externalist). And it follows *trivially* from this – whether one is an internalist or an externalist, by Harman's

Should we be internalists or externalists? To answer this question, Harman again appeals to the criterion of preserving common ways of thinking and speaking. Accordingly, he tentatively rejects internalism: the main problem with it, he feels, is that it entails that anyone who constructs and follows any set of moral imperatives that arises from his or her private set of concerns is perfectly moral and reasonable, so long as this set of imperatives is neither internally inconsistent nor inconsistent with anything else in that person's set of concerns. In other words, one would need to accept a sort of subjectivism in order to accept internalism. But "subjectivism like this is counter-intuitive, because we are ordinarily inclined to suppose that a person's intentions, aims, goals, plans and projects are one sort of thing and morality is another... We are inclined to think that morality has an external source, not an internal one".[112]

With internalism having been dealt with in this way, Harman concludes that externalism must be true. According to externalism, as Harman understands it, "morality does depend, at least in part, on external sanctions. If these are not divine, they must be human: that is, they must be social. Morality must, then [sic] depend essentially on some sort of social law enforced by custom and social pressure".[113] With that, Harman has arrived at his goal: morality must for the most part be a matter of custom and social pressure, and only sometimes merely a matter of one's personal concerns. In other

---

lights – that absolutely everything one ought to do must follow either from one's personal concerns or from the dictates of one's religion and society. Since this is precisely what the non-relativist denies, it is illegitimate for Harman to take it for granted in setting out the terms of the argument he uses to establish relativism.

[112] *Ibid.*

[113] *Ibid.*, p.93.

words, moral relativism is true. All that is required now is to articulate this relativism more fully and respond to the best of the anticipated objections.

The Social Custom Theory

Harman begins his articulation of relativism by setting out one formulation of it and considering its strengths and weaknesses. This tentative formulation gives rise to a number of distinct theories he entitles "social custom theories". The social custom understanding of morality is "that morality derives from the rules or customs that society enforces in a certain way, a way that may be specified differently in different theories."[114] Harman is quite clear that, on this view, there are absolutely no constraints on the moral rules that societies ought to enforce, save for the two requirements that a) the set of rules be internally consistent[115] and b) any factual assumptions that underly the society's decision (however this is brought about) to enforce a moral rule cannot be based on faulty non-moral beliefs.[116] Harman's statement of this point is very clear: "In a social custom theory of morality, then, the enforcement of custom is basic to morality. Morality is constituted by the rules, *whatever they are*, that society enforces. *There is no prior morality: if society were to enforce different rules, what is right and what is wrong would change*."[117] [Emphasis mine].

---

[114] *Ibid.*, p.93
[115] *Ibid.*, p.94
[116] *Ibid.*, pp.94-95. To take Harman's example, a society that endorses slavery on the basis of an erroneous belief that the slaves are innately inferior to humans in some way is objectively wrong in doing so. However, it seems clear from what Harman says that a society that endorses slavery on the simple grounds that it is pleasant to oppress people is doing nothing wrong. This interpretation of Harman, uncharitable though it may sound, is almost certainly correct: it receives strong support in his discussion of Hitler, as the reader will see shortly.

Harman feels that the social custom theory of morality has an important benefit over subjectivism: whereas subjectivism leaves it unclear why societies should have the moral rules they do, the social custom theory explains this convincingly. For example, he explains, "we normally suppose that there are moral restrictions on stealing, lying and cheating and that there are moral duties to be honest, to help others, and to show gratitude when you are helped. It is obvious that society has an interest in such restrictions and duties… Quite apart from any assumptions about a prior morality, we can understand why there might be social pressure to act in accordance with these rules…"[118]

Harman recognizes that not everyone will be inclined to agree with the radical view that, prior to a society setting out morality by fiat, there is no such thing as moral right or wrong. He borrows from social contract theorists the notion of a 'state of nature' (i.e. one in which no social rules are enforced). Many people feel, he admits, that

> …even in a state of nature, you should not kill other people or harm them in any way (except in self-defense, etc.); you should not steal from them; you should not try to mislead them; you should

---

[117] *Ibid.*, p.94. Admittedly, the 'whatever they are' description may be in tension with the two requirements just mentioned.

[118] *Ibid.*. These remarks seem to be in tension with the quotes just given above: if it really is the case that a society can put together just *any* set of self-consistent rules, and if there is no external morality, and if right and wrong change whenever social rules change, then it would seem that no consistent set of moral rules is intrinsically better than any other. That being so, it is not easy to understand why the social custom theory would make the rationale behind some society's choice of rules any clearer. The most charitable interpretation I can find for Harman here is this: while one could *in theory* find a society where those in power have a love for violent disembowelment (say) that transcends any other desire, and while in *this* society it would be morally permissible (or even obligatory) for everyone to follow the rule that one should violently disembowel oneself and others whenever possible, in *fact* this is a purely academic possibility since everyone actually does have the more familiar sorts of desires.
If Harman wishes to make this move, however, then it may become too difficult for him to distinguish this view from the social *contract* theory in an interesting way. This is particularly so since it may well be that the very idea of a society of disembowelment fanatics might run into problems of consistency, particularly since there could not (as Harman insists) be anything in the motivational set of such people that would disagree fundamentally with the motivation to disembowel. But the social contract theory, unlike the social custom theory, does not have the relativistic features Harman desires.

> try to help them if they need help. Other people feel that a state of nature would be more like a condition of all-out war of each one against the others in which moral constraints would play no real role. If the first group of people are right, no social custom theory of morality can be adequate to our ordinary conception of morality. But it is difficult to get a clear picture of what a state of nature would be like, so it is not easy to be sure whether you think that moral constraints would apply in such a state. Furthermore, it is cetainly not *clearly* essential to our usual conception of morality that moral constraints should apply in a state of nature. [119]

Since these facts about a state of nature cannot clearly be known by us, Harman suggests that we consider more directly the possibility that this sort of relativism will have counter-intuitive results, which (since Harman relies on conformity with common ways of thinking and talking as an important criterion for the acceptance of a view) would not be good news for the social custom theory. It is true, he admits, that our ordinary conception of morality appears inconsistent with the social custom view that "what is right and what is wrong will not always be the same in different societies".[120] However, he feels that this apparent inconsistency is lessened when we consider the matter more carefully. He draws our attention to the fact that, when faced with some very trivial discrepancies concerning moral rules between societies, we are not inclined to think that there is any objective right or wrong to be found (one can imagine, as an example of this, the question of how much eye contact one should engage in while speaking with a

---

[119] *Ibid.*, pp.95-96. This passage contains an important fallacy. The plausibility of the social custom theory, Harman concedes, depends on there being a good response to those who feel that people in a state of nature *should* avoid stealing from others, misleading them, etc. He does this by arguing that whether they should is unclear. But he switches, two lines into the passage, from the discussion of whether those in a state of nature *should* follow these rules to the quite different discussion of whether they *would*. The question about what *would* happen is not at issue, and those who find the social custom theory doubtful (as I do) can reasonably and consistently maintain that people would behave badly in a state of nature despite the fact that they should not.

[120] *Ibid.*, p.96. Incidentally, this casts some doubt on my attempt, in footnote 93, to interpret the social custom theory charitably. There, I assumed that Harman intends us to understand that all societies will *in fact* hold very similar sets of rules, even though very different sets of rules would be justified by the theory for some logically possible, but nonexistent, societies. But if this is not what Harman means, then it

friend). He also points out that, in some cases of cultural differences in morality, we judge a moral stance to be wrong because we feel it is based on false beliefs (one can imagine here a case where a society endorses the persecution of homosexuals in the mistaken belief that homosexuals are likely to be pedophiles). The social custom theory is not threatened by any examples of these sorts, Harman says: the true test comes when we consider cases where we judge some other (real or hypothetical) culture's morals to be objectively wrong even where we do not believe them to be based on a mistaken belief.

To illustrate this crucial question, Harman suggests that we consider a society that endorses the practice of cannibalism. He is careful to describe the cannibalistic society in such a way that we are much more likely to find it disgusting (or perhaps even unhygienic) than inhumane:

> To keep matters simple, we should not suppose that they actually raise humans expressly for the purpose of eating human flesh, as we raise certain animals. Nor should we suppose that they hunt people in the way in which our ancestors used to hunt animals. The people that are eaten, then, are people who have died natural or accidental deaths, people who have been legitimately executed as criminals, and people who have been killed in wars. In our society, we think that it is wrong to eat human flesh even in these cases. The issue is whether it is wrong for those in the cannibal society to do so.[121]

---

remains unclear why the social custom theory would help us to understand why we have the moral rules we do; and this explanatory power was presented as a major point in its favour.

[121] *Ibid.*, pp.97-98. However, since the whole purpose of our considering this society of cannibals is to see whether the most difficult instances of intercultural moral disagreement can be made to conform with what we take to be common sense morality, why not consider a much more difficult (and obvious) example than this one? Let us imagine, in addition to Harman's cannibal example, one of my own: in my version, the society endorses the eating of foreigners who are still alive and in great pain. The sole reason for doing this is that the society finds human flesh tasty, and finds the screams of the victims to be pleasant dinner entertainment: nobody in this society bases the moral rule that one is ethically permitted (or better yet, obligated!) to perform such acts on any false beliefs. Moreover, there is nothing else in this society's moral rules that conflicts with the moral imperative to cannibalize others: at the heart of this society's moral

We now have to consider 'whether it is wrong' for those in the cannibalistic society to eat human flesh in the situations Harman describes. He feels that our intuitions are not clear either way: while we may deem it morally wrong to eat humans, nonetheless "we have some inclination to think that, when cannibalism is a custom of a society, the members of that society who eat human flesh cannot be blamed for doing so. We feel uncomfortable about saying that it is morally wrong of *them* to eat human flesh."[122]

Harman tries to reinforce this point with a second example. This time, we are to imagine a society in which the people accept that it is morally wrong to harm one another, but not that it is morally obligatory to help one another. Members of this society feel it is perfectly acceptable to abandon to their fates those who have been in serious accidents, even if these accident victims could easily make a full recovery at almost no inconvenience to anyone else by being taken a short distance to the doctor, say. At first, Harman thinks, we might feel intuitively that the members of this society are acting

---

values is a complete lack of concern for any outsiders. I hope the reader will consider my cannibalistic society, and not just Harman's, in what follows.

[122] *Ibid.*, p.98. Do we also feel uncomfortable saying that it is wrong of the cannibals in *my* scenario to eat human flesh? I certainly do not feel that discomfort; and if there is nothing odd about saying so, then this argument of Harman's fails.

Nor is it clear that one can move as freely as Harman does from the premise that one cannot be blamed for doing something to the conclusion that one's doing it is not morally wrong. Imagine, for instance, that a team of brave philosophers proposes to journey to my land of cannibals to try to stop them (whether by reasoning, propaganda, or force – it doesn't matter) from eating the peaceful nearby islanders whose flesh is considered a delicacy in the finest cannibal restaurants. Before the journey, someone asks the philosophical missionaries, 'But can you really blame them? After all, everything they've been taught suggests that this is morally upright behaviour. Try to see it from their side!'

'You miss the point', one of the philosophers replies. 'We don't blame them, and that isn't the point of our trip. The point is that, regardless of whether they can be blamed, they're doing something wrong. And we want to stop that'.

Is this second-last sentence unintelligible? Is it a contradiction in terms, given ordinary ways of speaking, to say that someone has done something wrong even though it might not be right to blame the person for it? If not, then Harman seems mistaken in making this move.

wrongly. But when we think about it more carefully, he suspects, we will find it less clear that this is so: the reason is that "given the social practices of the society, it is not obvious that anyone in it can be blamed for not helping others".[123]

Harman tries to reinforce this last point with a further consideration:

> it is not obvious that if you were to visit that society and were to see someone in need of help, which you but no one else could provide, you would have any duty to help. For, if the situation were reversed, he would not suppose that he had any reason to help you; and it is not clear that you are obligated to help someone whom you know would not help you if the situation were reversed.[124]

From all this, Harman feels, it becomes evident that our moral intuitions neither clearly favour nor clearly oppose the results of the social custom theory. Hence, though there is not (Harman thinks) anything clearly wrong with this view, we need to find a more 'precise' form of relativism so that we can test our intuitions more adequately. He turns to this next.

The Tacit Convention Theory

---

[123] *Ibid.*, p.98. As I argued in the last footnote, however, it is not clear that we can move fairly from the belief that the people cannot be blamed (if indeed they cannot, which is far from clear in this case) to the conclusion that they are not acting wrongly or that their moral views are not wrong.

[124] *Ibid.*, p.98. Some points in this passage are less clear than one might wish. For instance, it seems that Harman is invoking the higher-level moral rule that one need not follow a moral rule to someone else's benefit if that other person would not follow that same rule for you. Since he has also made clear earlier in the chapter that he assumes there to be no morality outside of the moral rules a particular society happens to endorse, it is not clear what sort of status this higher-level rule is meant to have here. Is it just a personal principle of the member of our society? Does that individual need to follow his or her moral principles while visiting the other society? If not, then the basis for invoking the higher-order rule in this case is mysterious. If so, then (since the person from our society presumably also accepts the moral principle that we should help those in need when this can be done without great inconvenience) Harman's claim collapses.

We can see morality as a set of tacit conventions that we engage in for self-interested reasons. As Harman points out, Hume accepts that some – but not all – of morality can be explained in this way. According to this 'tacit convention' theory, members of a society engage in processes of "implicit bargaining and mutual adjustment" to arrive at these conventions, as when two people rowing the same boat adjust their rowing speeds to match one another's.[125] We engage in such bargaining and adjustment from a "self-interested motive."[126] Just as it is in each rower's personal interest to cross the river, it is in everyone's personal interest to live without fear of violence, theft, and so on; and these self-interested aims are most easily achieved by our agreeing to follow tacit rules on the condition that everyone else does.[127]

Harman claims that "Hume's tacit convention theory of morality is a more specific version of the social custom theory", but that it also has some advantages over the more general version. One of these advantages is that it explains more clearly why considerations of social utility are so important in moral rules: "certain rules are conventionally adopted because each person benefits from everyone else acting in accordance with those rules. We therefore expect rules to be adopted if they promote social utility in the sense that they are beneficial to all."[128]

---

[125] *Ibid.*, p.104.
[126] *Ibid.*
[127] See Skyrms 2004 for a more recent discussion of this idea.
[128] *Ibid.*, p.105. How is the claim that this is a variety of the social custom theory to be squared with Harman's earlier comment that the social custom theory is the view that right and wrong change whenever social rules change, and that morality consists in whatever rules a society happens to enforce? This would appear to be a major inconsistency unless Harman means that there are core moral rules (the ones that lead to social utility and arise out of self-interest) to which we agree for prudential, and not moral, reasons; while it is only other rules (such as, perhaps, the rule that men need to wear neckties to formal social proceedings) that are rationally optional but nonetheless binding if a society happens to endorse them. So following the principle of charity, it seems we should understand Harman in this way.

Harman suggests that seeing the tacit convention theory as a more specific variant of the social custom theory sheds light on our intuitions regarding his cannibalism example. He explains the matter as follows: in our society, we have it as a tacit convention to treat one another with respect and dignity, and our "conventional habits of expressing that respect" involve our finding it wrong to eat other humans.[129] However, those living in Harman's cannibalistic society "see nothing wrong with eating people, *and there is no obvious reason why they should*"[130] [emphasis mine].

Harman feels that, once we have considered the cannibalism situation carefully with these points in mind, we will recognize certain facts about moral judgment and its limitations. He presents many of these as sets of contrasts, all of which are meant to illustrate the linguistic implications of 'ought' and 'should':

However, it then becomes even less clear how what Harman suggests is any different from the more standard (and, according to its proponents, non-relativistic) versions of the social *contract* theory. For the standard social contract views of morality (not to mention more or less all other normative views) also are silent on specific issues like necktie-wearing, leaving these matters to the idiosyncracies of culture.

[129] *Ibid.*, pp.105-106.

[130] *Ibid.*, p.106. The part of this quote I have italicized is all-important, of course: if there is some reason why they should find it objectionable, then it would follow by Harman's own account that it would (other things being equal) be wrong of them to perform their acts of cannibalism; and then Harman would no longer be talking about a version of the social custom theory. Further, since moral conventions arise from implicit, self-interested bargaining (where one adheres to moral principles only when everyone else will and only when this agreement seems to lead to one's own benefit), it must *not* be against the wishes of Harman's cannibals that they should be eaten, in their turn, when they have 'died natural or accidental deaths' or 'been killed in wars', say. Hence, the only way in which Harman's cannibal example can work for his purposes is if the members of that society have no overall opposition to being eaten themselves. But in this case, it seems to fail for a different reason: the social practice we would then have to imagine would, however disgusting we might find it, not be the sort of thing many of us would find *ethically* objectionable. So the example seems to fail either way, now that Harman has set out his view more clearly.
For that reason, I suggest that a better test would be (again) my own cannibal scenario, where those who are eaten do not wish to be *and* where they do not belong to the same society. Is it morally right for the cannibals to perform their cruel actions then? This is the sort of question Harman needs to deal with.

> It does not seem right to say that each of them ought morally not to eat human flesh or that each of them has a moral duty or obligation not to do so. At best we might say that it ought not to be the case that they eat human flesh, but… that is not the same sort of judgment at all…
>
> From our own point of view we can judge their acts and their situation, even their society and morality; but we cannot, it seems, judge *them*…
>
> We can make certain moral or evaluative judgments about the cannibals; for example, we can call them "ignorant savages". But we cannot correctly say of them that they are morally wrong to eat human flesh or that they ought morally not to do it…[131]

He recognizes, however, that the reader may have difficulty understanding the sorts of distinctions he intends by these examples, since it is "not always clear when someone has a reason to do something" on Harman's view. He seeks to clarify these points by means of two extended discussions: one about Hitler and Stalin, and the other about the acts/omissions dichotomy.

---

[131] All these quotes are from p.106. Regrettably, Harman does not back up his claims about the proper usage of 'ought' and 'should' by reference to any dictionary. I have surveyed the entry for 'ought' in the Oxford English Dictionary and found no evidence for the distinction Harman makes. It is an onerous process to comb through the classics of literature for instances where an author has to decide whether to use the words 'ought' or 'should' to refer to the actions of those who do not seem to accept the moral views of the speaker; but I hope that one such discovery will suffice. The passage in question is from Genesis 34, which describes the encounter between Dinah, Jacob's daughter, and Shechem, son of the Hivite prince Hamor. While Jacob and his family are sojourning in Hivite territory, Dinah is raped by Shechem, who then asks his father to arrange for Dinah's hand in marriage. Hamor attempts to negotiate with Jacob and his family as though the matter were some sort of business arrangement, offering them property, free trade, and the prospect of peaceful cohabitation. When Jacob's family insists, as a condition of accepting the deal, that all the Hivite men be circumcised, Hamor agrees and circumcises them. The accepted rules concerning rape and marriage in the place and time when the story was written, together with the biblical description of these events, give the reader no reason to suspect that Shechem or the rest of the Hivites were violating any of their *own* social conventions in acting as they did. Nonetheless, when Jacob's sons hear about the rape, they become "indignant and very angry, because he had wrought folly in Israel by lying with Jacob's daughter, for such a thing **ought** not to be done." (Revised Standard Version, Genesis 34:7; the King James Version renders this in the same way -- emphasis mine). Later, when Jacob finds fault with the wisdom of his sons Simeon and Levi (who have just massacred the Hivites in retaliation), they reply: "**Should** he treat our sister as a harlot?" (Genesis 34:31, Revised Standard Version. the King James Version also uses 'should' -- emphasis mine). Unless the committees overseeing these and other

Hitler, Stalin and the Obligation to Assist

As we have seen, Harman's case for some form of relativism (specifically, the tacit convention theory) depends throughout on appeals to our moral intuitions and to everyday ways of speaking. The greatest obstacle he faces at this point in his argument for the tacit convention theory is that our intuitions may not agree with his own: as the above quotes show, he relies on a sharp distinction between discussions of a) what people ought to do or should do and b) what it ought to be the case that they do. In order to persuade his reader of the crucial importance of this distinction, he asks us, first of all, to accept some assumptions about Hitler and Stalin: we are to assume that Hitler saw nothing wrong with any of his horrible actions, whereas Stalin felt moral remorse about doing what he did but nonetheless deemed that he was doing the best he could in the circumstances.[132] Once we have done this, Harman feels confident that we will agree with him on the following carefully-worded points, which I quote here at length so that the reader might be able to judge the arguments fairly:

> [T]he following remarks are weak and even in some way odd: "It was wrong of Hitler to have ordered the extermination of the Jews". "Hitler ought morally not to have ordered the extermination of the Jews".
>
> One might suppose that it is the enormity of Hitler's crime against humanity that makes such remarks seem too weak. He killed so many people; it would have been wrong of him to have killed only one. To say simply that it was wrong of him to have ordered the extermination of the Jews suggests that it was *only* wrong – that it was wrong only in the way in which murder is wrong. And, given what Hitler did, that is as if one were to say that it was *naughty* of Hitler to have ordered the extermination of the Jews.

---

biblical translations are guilty of strange lapses in usage, Harman is mistaken on these important points of usage on which his argument rests.

> This explanation, however, is not completely satisfactory. First of all, there are things we can say about Hitler without the same sort of oddity. Although it would be odd to say that it was wrong of Hitler to have acted as he did, it is not equally odd to say that what Hitler did was wrong. Similarly, there is no oddness in the remark, "What Hitler did ought never to have happened." That is not odd in the way that it is odd to say, "Hitler ought morally not to have ordered the extermination of the Jews." But, if the enormity of his crime makes the one remark odd, why doesn't it make the other remark as odd?"[133]

Harman attempts to answer this question by denying that it is odd to say that *Stalin* ought not have ordered the deaths of the millions of people *he* ordered, given the assumptions we are making about Hitler and Stalin. The difference is that, unlike Stalin, "Hitler was not just immoral, he was amoral, he was evil."[134] In other words, Harman says, he was "beyond the reaches of morality or at least that part of morality that we invoke in judging him to be an evil man."[135] It follows from this that anyone who is evil on Harman's view must be unable to be made to care that he or she is evil.[136] Harman is clear that we *cannot* judge Hitler: "We can judge his acts with reference to our morality, but not Hitler himself, since that would imply that Hitler accepted the relevant moral conventions".[137] Hence, to judge that Hitler is evil (which Harman feels is correct) is, surprisingly, not to make a moral judgment about him. Nor is this all. Those who think there was some reason for Hitler not to have ordered the deaths of millions turn out to be mistaken, as becomes clear when we recognize that Hitler was evil: "If he was willing to exterminate a whole people, there was no reason for him not to do so".[138]

---

[132] *Ibid.*, p.108.
[133] *Ibid.*, p.107.
[134] *Ibid.*, p.108.
[135] *Ibid.*, p.108.
[136] *Ibid.*, p.108.
[137] *Ibid.*, p.109.
[138] *Ibid.*, p.109.

At this point, Harman feels he has said enough to establish that his version of relativism actually squares quite well with ordinary ways of speaking and thinking. However, he feels there one further selling-point of the tacit convention theory worth mentioning: it helps us make sense of the moral distinction we draw between the obligation to assist and the obligation not to harm. The first of these obligations is weaker than the second: for example, while it is bad to allow the needy to die through our failure to intervene, it is morally much worse to kill someone. This fact can be explained quite readily, Harman feels, if we adopt the tacit convention theory; for we can then imagine a sort of bargaining process between the rich and powerful and the poor and weak. The poor and weak would prefer that everyone respect both the obligation to assist *and* the obligation not to harm, since they stand to benefit from both; but the rich and powerful would prefer that we accept only the obligation not to harm, since they stand to receive much less (and to have much more expected of them) if society were to recognize an obligation to assist. The result, according to Harman's story, is a compromise between the two, where both obligations are accepted but where the obligation to assist is much weaker. The fact that this aspect of our morality could perhaps be explained by this story of our bargaining our way toward morality from where we originally stand (and not from where we ideally might stand, as Harman makes clear)[139] gives *some* support to the tacit convention theory, Harman feels.[140]

---

[139] *Ibid.*, p.111.

[140] *Ibid.*, pp.110-112. This is at least somewhat more persuasive than Harman's discussions of Hitler and the cannibals, though of course its persuasive power depends among other things on the difficulties other theories have in accounting for this same fact, which Harman does not discuss. Still, it is not without its other problems. George Sher, in his review (Sher 1980) of an article of Harman's in which a similar argument is made, points out that we are not given enough information about the initial bargaining situation to know whether what we have now would be a reasonable compromise between the powerful and the powerless; and that, moreover, we have good reason to think that self-interested parties would *not* end up in the position Harman sets out. For as Harman sees it, the powerful ended up agreeing to answer to both the

The Tacit Convention Theory and Internalism

Now that Harman has established relativism, and in particular the tacit convention theory, a further question arises: is it a form of internalism or externalism? Harman sees it as a combination of the two: "[t]he tacit convention theory… sees moral principles as principles for which the source is both internal and external. They are legislated by others and by yourself."[141] Moreover, both of these sources are necessary. As we have seen before, according to any version of the social custom theory, there is no such thing as a merely personal moral rule: "morality is essentially social."[142] At this point, Harman adds the internal requirement: you yourself must accept a principle in order for it to be binding upon you. "The principles that apply to you, according to this theory, are not simply whatever principles are conventionally accepted by the surrounding group; you must accept the conventions too. Otherwise they could not give you reasons to do things, and judgment about what you ought morally to do or about what it would be right or wrong of you to do could not be made with reference to those conventions".[143]

---

obligation to assist *and* the obligation not to harm, while the powerless end up agreeing to answer only to the latter (as can be seen, for instance, in the fact that the poor do not have to pay income taxes). In other words, the powerful are portrayed as agreeing to give up something for nothing. But surely, if they were rational and purely self-interested in addition to being powerful, they could manage something better than that. Moreover, as Sher points out, Harman's assertion that we have a 'weaker' principle of mutual aid than of no mutual harm actually fails to account for our everyday morality: it implies only that our duty to assist can be overcome by our duty not to harm when we find ourselves affected by both obligations, but it leaves unexplained the fact that we often take ourselves not to be compelled by the duty to assist when there is *no* obligation-not-to-harm to overcome it in the other direction.
These points seem to rob Harman's argument of its force.

[141] *The Nature of Morality*, p.112.

[142] *Ibid.*

[143] *Ibid.*, pp.112-113. This seems to imply, strangely, that one's fear of punishment (which, again, Harman feels must be implicit in any form of external morality) cannot by itself give one a reason to follow a rule. But why not? If a person who respects a moral rule for reasons of pure self-interest is not thereby acting with a reason, then (as the reader will by now see) Harman's ethical project collapses immediately.
The most charitable interpretation I have been able to devise for Harman here is as follows: my knowledge that there is a moral rule, with its corresponding threat of punishment, cannot by itself provide me with a reason to follow it *if* I can be sure that I will avoid detection in breaking it. For if I am able to do that, then

One difficulty Harman anticipates is that we all seem to be subject to a variety of external moralities we care about: those of our family, our business, the terrorist organization of which we are a member,[144] etc. What should we do when these demand different things of us? Harman feels that in cases like that, we should just follow the rules of whichever group happens to be most important to us.

There is, however, a "limiting case of morality in which the relevant 'group' contains only one person."[145] In these cases, Harman thinks it may be that the person in question will have reasons to do certain things, but that nobody else will. He gives a peculiar illustration of this: "For example, a pacifist may think that he morally ought not to participate in wars, although he will not make the same judgment about other persons. He will not say that it is wrong of them to participate, although he will certainly think that it is bad for everyone that they engage in wars."[146]

Immediately after discussing this unusual case of what he takes to be morality with no objective external constraints, Harman clarifies his position by saying that "[w]ithout

---

I am not doing anything I have an external reason not to do (since I do not fear the sanction) nor that I have an internal reason not to do (since I do not personally care about following this rule). However, this would make it difficult to understand why most people ought not break whatever rules they can get away with, unless we posit, with Harman, a Freudian superego that internalizes the social rules our parents give us, demanding that our ego obey them; and as I have already noted, there are problems with going this way. Also, even if we accept Freudian theory, problems would arise if people's parents gave one moral rules that disagreed with those of the rest of society; for in that case we would have only internal reasons for following some rules and only external reasons for following other ones. But Harman's whole case for relativism in his fifth chapter, as we have seen, depends on there being important differences between the rules given by different parents.

[144] Admittedly, Harman does not actually use this particular example; but it is not clear from anything he does say why this should not be a legitimate possibility.

[145] *Ibid.*, p.113. This conflicts, of course, with Harman's previous claims that morality is 'essentially social' and that it must have an external component.

objective external constraints, there would be no such thing as morality, as we ordinarily understand it, even if people adhered to their own personal principles."[147]

At this point (the end of the ninth chapter), Harman has come to an end of his defense of his version of relativism. The next two chapters, which are meant add some clarifications to what is meant by having a reason, I will discuss in Chapter Three; so I will skip over them here. I will also pass over the final two chapters of *The Nature of Morality*, since they are concerned only with discussions of egoism and utilitarianism and do not contribute to Harman's case for relativism.

While *The Nature of Morality* is Harman's fullest defense of his moral views, and while much that he says there is repeated by him elsewhere, he makes a few points in his other ethical writings that strengthen his case for relativism. I will spend the rest of this chapter summarizing these points.

Moral Relativism and Moral Objectivity

Harman's section of this book (Harman and Thomson 1996) is mostly spent recapitulating arguments that appear in *The Nature of Morality*. There are, however, some new points worth mentioning.

In his first chapter, Harman attempts to support relativism by emphasizing the extent of moral diversity: it "occurs not just between societies but also within societies and in a

---

[146] *Ibid.*, pp.113-114.
[147] *Ibid*, p.114

way that leads to seemingly intractable moral disagreements."[148] Harman feels that these disagreements (he gives as examples the abortion issue, the moral vegetarianism issue and the issue of our obligation to assist the needy) "survive extensive discussion" even between rational and fully-informed people.[149] Moral relativism, Harman feels, is "a reasonable inference from the most plausible explanation" of this moral diversity.[150]

However, a more important theme in this chapter is an extended comparison between morality and physics. Just as physicists now accept (Harman says) that questions of whether a given body is objectively at rest or in motion are pointless given the truth of Einstein's theory of relativity, the most thoughtful ethicists now accept that questions of whether a given act is objectively right or wrong are also predicated on an error.[151]

In his second chapter, Harman promotes his view of morality as the result of a bargaining process. This time, in addition to making his point by referring to the contrast between the strong obligation not to harm and the weak obligation to assist, he invokes our moral attitudes toward nonhuman animals. Why do most people not accord to such animals the same moral status given to humans? The answer, Harman explains, lies in the fact that "animals cannot get together to put pressure on people" in the way that disenfranchised people can.[152] However, this leads to an obvious difficulty: it would seem to suggest that, since nonhuman animals can *never* turn up at the bargaining table, it would never be wrong for anyone in any society to harm animals, on Harman's view. But this would

---

[148] Harman and Thomson 1996, p.10.
[149] *Ibid.*, p.11.
[150] *Ibid.*, p.8.
[151] This will be discussed at some length in Chapter Two.

lead to two unfavorable conclusions: first, many of us would find the view that animals (not to mention severely disabled humans, who cannot easily get together and pressure people for their rights) have no moral worth a counterintuitive one; and second, if it did indeed follow from Harman's view that everyone – regardless of society – were morally permitted to harm unowned nonhuman animals, then his view would quite clearly not be a relativistic one.

Harman seems to recognize that problems are lurking here, for he attempts an account of our moral feelings toward animals. First, he admits that there "may well be something to" the sociobiological view that we are "genetically constructed so as to feel such concern and respect for others and that makes morality possible"[153]. However, he also entertains the alternative possibility that we bargain ourselves into being concerned for the well-being of others because it is advantageous for all of us to be cared about by other members of society. On this view, our feelings of concern for others would arrive after our agreement to act *as though* we were concerned for them, and would be inculcated through force of habit. Harman defends this possibility with a curious argument:

> One develops a habit of taking an interest in others. This is possible because taking an interest in something or someone is the kind of thing one can do. One watches a game on television and decides to root for one or another of the teams; one does this to make the game more interesting to watch; doing so structures one's perception of the game. Similarly, one takes an interest in a conversation at a party because that makes it less boring to stand there and listen. In playing a

---

[152] *Ibid.*, p.26.

[153] *Ibid.*, p.26. There had better *not* be much to it, though, or else Harman's relativism is in trouble. In this book, he defines relativism as the view that "[t]here is no single true morality. There are many different moral frameworks, none of which is more correct than the others" (p.8). But – unless we attribute to Harman the unlikely view that different societies are composed of people with radically different hereditary emotional traits – it would follow from the sociobiological thesis he discusses, and from his view of

> game with friends, one takes an interest in winning; otherwise one's friends will get mad. In the same way, one might take an interest in the welfare of other people. This is something one is encouraged to do as a child and one can develop the habit of doing it.[154]

Our concern for animals might arise in the following way, Harman conjectures: we start by training ourselves to feel concern for our family members, for the reasons just described; then, we extend this to all the members of our society. After this, "there will be a tendency for the disposition to apply to similar cases, so one may find oneself with concern and respect for outsiders, too. One might even [?] come through stimulus generalization to be disposed to feel a certain concern for animals."[155]

Other Writings

In addition to these two books, Harman has argued for moral relativism in a number of articles. While some of these predate his longer works, he has corrected them[156] for a recent compilation (Harman 2000). As his writings after 2000 have not dealt with the

---

morality as a bargain struck between our various selfish interests and moral sentiments, that relativism as described here is false.

[154] *Ibid.*, pp.26-27.

[155] *Ibid.*, p.27. I have doubts about some of Harman's odd claims about our reasons for doing the things he describes, but pass over them in order to raise the following question: why is it, on this view, that one would experience stimulus generalization rather than stimulus discrimination in regard to animals? It is well known that appropriate co-operation between these two processes is needed in order for a person to function effectively: as one psychology textbook puts it, "Stimulus generalization ensures that we expand our learning beyond the immediate conditioned stimulus. However, stimulus discrimination operates to guarantee that we do not generalize too broadly or inappropriately." (Matlin 1998, p.194) One example of this is that "a child who has been burned by a gas stove will not show a fearful emotional response to a photograph of a fire" (*Ibid.*). Now, according to Harman, morality is a contract reached by bargaining from self-interest and from a concern for the things we care about. But according to Harman's present suggestion about animals, we do not have a self-interested reason for caring about them, and we do not have a natural concern for them. Therefore, having moral rules limiting our dealings with them is simply a losing proposition for us, and contrary to Harman's idea of the nature of morality. Hence, if stimulus generalization were to lead us toward concern for animals, we would be generalizing too broadly and inappropriately. Therefore, one would expect that stimulus discrimination would come into effect at that point, and stop us from caring about them. Why has this not happened? The reader is left in the dark at this point, for Harman does not take his speculations any further than the point reached in the passage I have cited.

[156] Harman 2000, p.xi

relativism issue in any new way[157], this latest collection seems to offer an up-to-date presentation of Harman's views on the matter.

The first of these essays, 'Moral Relativism Defended', begins with an attempt to separate moral relativism from what Harman takes to be the straw man version he finds in the work of some of its prominent critics. These critics represent relativism as being committed to the following inconsistent triad of premises:

> (a) There are no universal principles.
> (b) One ought to act in accordance with the principles of one's own group.
> (c) Principle (b) is a universal moral principle.[158]

Instead of this, Harman suggests that we should view moral relativism as a "soberly logical thesis". Harman articulates this thesis, and defends it, by drawing an analogy with the sizes of dogs: just as a dog cannot be simply large or small, but only large or small relative to some comparison class (such as the class of other dogs), so actions can only be good or bad in relation to some set of social standards.[159] In the same way, Harman holds, morality depends on what we can be made to care about:

> Consider this imaginary example. Intelligent beings from outer space land on Earth, beings without the slightest concern for human life and happiness. Some of us seem to be able to imagine that the fact that a certain course of action on their part will injure inhabitants of Earth might mean nothing to them and might in particular give them no reason to avoid the action. Those of us who

---

[157] The one possible exception I have found is Harman 2008a, in which the claim is made that "Any absolutist (non-relativist) reduction of morality faces the epistemological problem of showing how that conception of morality is better supported than its competitors. The problem is that there are competing moral frameworks and no way to test them against the world." (Section 2.3). However, it appears that the intended point is the same as that made in Harman's famous example of seeing the proton vs. seeing the wrongness of burning a cat (see the first chapter of *The Nature of Morality*); so it will not receive special treatment in this dissertation.

[158] Harman 2000a, p.3.

[159] *Ibid.*, pp.3-4.

> think we can imagine this find it odd to go on to say that the beings ought nevertheless to avoid injuring human beings or that it would be wrong of them to attack human beings. Of course we would want to resist them if they did such things and we would certainly make negative judgements about them; but we would judge that they are dreadful enemies to be repelled and even destroyed, not that they morally ought not to act as they do.[160]

Much of the rest of 'Moral Relativism Defended' deals with arguments I have already dealt with in my discussion of *The Nature of Morality*. The one further point of interest in that article is a curious discussion about moral bargaining. In the course of arguing that morality is the result of bargaining, Harman makes some claims that seem at least difficult to reconcile with the relativism he seeks to defend (such as "that each person has an inalienable right of self-defence and self-preservation"[161]). But he also considers an important objection against the bargaining theory of morality: "It will, obviously, not be enough to say that one has implicitly agreed to keep agreements, since the issue would then be why one keeps *that* agreement."[162] Harman's response to this objection is that to accept an agreement in this second-order way implies intending to keep it. Further, he claims that "our moral understanding contains or implies an agreement to be open and

---

[160] *Ibid.*, p.5. This argument seems fallacious in one of two ways, depending on what Harman takes 'reason' to mean in the fourth line. On the one hand, by including as a premise the assumption that the aliens have "no reason to avoid the action", Harman might mean to deny that it would be morally wrong of them to perform it. In that case, the claim he makes is trivially true, since it amounts to no more than saying that those who think the aliens would be doing nothing wrong think that the aliens would be doing nothing wrong; and no interesting conclusion (let alone a version of interesting moral relativism) could follow from it. On the other hand, Harman's claim that some people can imagine aliens who have no reason to avoid harming us could just be his way of reiterating that some people can imagine aliens that do not *care*, and perhaps could not be made to care, about us. In that case, the strength of his argument hinges on his later claim that those who can imagine this find it odd to say that it would be wrong of the aliens to harm us maliciously. But this claim would, in that case, be false: to take one counterexample, *I* can imagine the situation and I do not feel that making this judgment would be odd in the least. I will discuss this further in Chapter Four.

[161] *Ibid.*, p.12.

[162] *Ibid.*, p.13.

honest with others."[163] He then attempts to clarify his view as follows: "Not everything generally agreed on is right, and sometimes courses of action are right that would not be generally agreed to be right. But this is no objection to my thesis. My thesis is not that morality derives from an agreement in moral judgement; the thesis is rather that moral judgements make reference to and are made in relation to an agreement in intentions."[164] Sadly, the natural questions of just how this agreement in intentions is arrived at, how we are even to imagine it, how one could determine what such an agreement would involve for any given set of people with given initial intentions, and how what people agree upon and take to be right is meant to be different from what *is* right on Harman's view (for these, we are now meant to understand, are two different things) do not seem to be explained by Harman either here or elsewhere.

The second, third and fourth articles of this collection ('What is Moral Relativism?', 'Relativistic Ethics: Morality as Politics' and 'Justice and Moral Bargaining') contain some points that will be worth considering later: I will save discussion of them until then.

---

[163] *Ibid.* This move seems to lead to difficulties. For one thing, it appears to imply that we have *non-relative* moral duties to be open or honest with others: since we need to be open or honest in order to be party to these agreements *regardless* of the society in which we find ourselves, and since we need to be party to these agreements in order to have morals at all, then it seems that anyone who has any morals at all must hold it to be morally obligatory to be open or honest with others. But it follows from this that there are some moral obligations that apply not only to all actual sets of morals, but even to all *possible* sets of morals. Hence, it follows that Harman's brand of moral relativism is false. For another thing, it leaves unanswered the difficulty it was meant to deal with – the difficulty that we already need a prior set of moral obligations in order to be morally obliged to keep the promises that we bargain our way toward, whereas Harman's views seem to entail that this bargaining is the *origin* of all our moral obligations. For someone might, for all Harman says here, merely pretend to have the right intentions at the hypothetical bargaining table while secretly intending to break the promises made there. Harman tries to avoid this problem by saying that "such a person may or may not have… a reason [for taking the agreement seriously]. If someone does not already accept something of our morality, it may or may not be possible to find reasons why he or she should." (p.13-14). But this cannot be satisfactory even on Harman's account, since it (together with Harman's other views about morality) implies that anyone who feels he or she can escape detection does nothing morally wrong even when breaking a rule to which he or she holds *others*.

[164] *Ibid.*, p.14

The fifth article ('Is there a Single True Morality?') is, in my view, Harman's most sophisticated and original contribution to the relativism discussion: I will discuss it further in the Conclusion. The final article of the collection ('Moral Philosophy and Linguistics') is interesting in that it makes a proposal – that we might make progress in morality by searching for a set of universal moral constraints, just as linguists search for a universal grammar – which, if successful, could mean trouble for relativism.[165]

What Remains?

The reader who has patiently worked through the commentary offered in the footnotes will now see that, if my remarks have been relevantly free from error, not everything in Harman's case for relativism can survive critical scrutiny. Of what remains, some of Harman's points may still be powerful enough to count as plausible arguments for relativism; some can at best stand as arguments against certain forms of realism; and some may only survive as more general issues that nearly any metaethical position must deal with. To help ensure that my answer to relativism is a viable one, I take no chances here: everything surviving this preliminary criticism of Harman will be treated as a potential component of a new, strengthened argument for relativism, and will be dealt with at some point in the dissertation.

---

[165] Harman has, more recently, promoted this idea elsewhere. See, for example, Harman 2008b.

I present here what I take to be a full inventory of what has not yet been dealt with adequately. I have arranged these into thematic categories, each of which will receive its own chapter. Here, then, is a more detailed outline of the next two chapters:

Chapter Two: Scientific Analogies. Harman makes his case for relativism by invoking three scientific analogies that have not been dealt with: the analogy between moral relativity and Einsteinian relativity; the analogy between moral frameworks and comparison classes; and (most famous and most important) the analogy between seeing the wrongness of burning a cat and seeing a proton. The first two of these, I will argue, commit the fallacy of faulty analogy; but the exposition of these fallacies deserves fuller treatment than could have been offered in Chapter One. The third analogy will be dealt with using the parallel arguments approach.

Chapter Three: Moral Disagreement and Moral Reasons. I will begin this chapter with an outline of Harman's version of the argument from moral disagreement, which I will subject to the parallel arguments approach. Next, I will consider Harman's discussions of the relationship between morality and our reasons for action. In particular, I will examine what Harman could mean by the phrase 'having a reason', which plays a key role in his strongest arguments for relativism. I will then respond to these arguments using (*inter alia*) the parallel arguments approach. However, an interesting point arises from this discussion: it seems, for reasons that I will make clear, that Harman is not alone in using the term 'having a reason' in a way that leads to difficulties. Many anti-relativistic metaethicists seem to be guilty of similar lapses.

Chapter Two: Scientific Analogies

Introduction

I divide this chapter into two sections. In Section I, I deal with a) Harman's analogy between moral relativism and a different sort of relativism involving comparison classes, and b) his analogy between moral relativism and Einstein's theory of relativity. In Section II, I deal with Harman's famous contrast between the roles of moral and scientific facts in their respective theories.

Section I – Relativism, Comparison Classes and Einstein's Theory of Relativity

These first two comparisons of morality with science are meant to be analogies rather than disanalogies.[166] The first of these analogies appears in his important early paper, 'Moral Relativism Defended'. The analogy reads as follows:

> My moral relativism is a soberly logical thesis – a thesis about logical form, if you like. Just as the judgement that something is large is true or false only in relation to one or another comparison class, so too, I will argue, the judgement that it is wrong of someone to do something is true or false only in relation to an agreement or understanding. A dog may be large in relation to chihuahuas but not large in relation to dogs in general. Similarly, I will argue, an action may be wrong in relation to one agreement but not in relation to another. Just as it is indeterminate

[166] It is, admittedly, not entirely clear whether Harman intends these comparisons to serve as *arguments* for moral relativism or merely as clarifications of what he means when he says that morality is relative. If the former, the following discussions will (if successful) show that Harman is guilty of the fallacy of faulty analogy in both his comparisons. But even if his analogies are not intended as arguments, I hope that my discussions will show that Harman has misconceived the relationship of ethics with science in both instances.

> whether a dog is large, period, apart from any relation to a comparison class, so too, I will argue, it is indeterminate whether an action is wrong, period, apart from any relation to an agreement.[167]

The second analogy I consider appears most importantly at the beginning of *Moral Relativism and Moral Objectivity* (Harman and Thompson,1996), where it sets the stage for Harman's case for relativism in that book. The analogy invokes the theory of relativity as follows:

> Motion is a relative matter. Motion is always relative to a choice of spatio-temporal framework. Something that is moving in relation to one spatio-temporal framework can be at rest in relation to another. And no spatio-temporal framework can be singled out as the one and only framework that captures the truth about whether something is in motion.
>
> According to Einstein's Theory of Relativity even an object's mass is relative to a choice of spatio-temporal framework. An object can have one mass in relation to one such framework and a different mass in relation to another. Again, there is no privileged spatio-temporal framework that determines the real mass of an object.
>
> I am going to argue for a similar claim about moral right and wrong. That is, in my contribution to this book I am going to defend *moral relativism*. I am going to argue that moral right and wrong (good and bad, justice and injustice, virtue and vice, etc.) are always relative to a choice of moral framework. What is morally right in relation to one moral framework can be morally wrong in relation to a different moral framework. And no moral framework is objectively privileged as the one true morality.[168]

In her book-length debate with Harman, Judith Jarvis Thompson spends over seven of the thirty pages in which she is permitted to respond to Harman's remarks drawing out some implications of his analogy with Einstein's theory of relativity that she suspects Harman

---

[167] Harman 2000a, pp.3-4

would be loath to accept. But she makes no mention of the fact that Harman's analogy breaks down when one considers what relativity in physics actually involves. There are also, to my knowledge, no criticisms of 'Moral Relativism Defended' that take issue with Harman's comparison class analogy.

In fact, both of these analogies have serious problems.

Ia: Comparison Classes

Let us begin with the "soberly logical thesis" Harman advances in 'Moral Relativism Defended'. Harman's claim, as we have seen, is that "[j]ust as the judgement that something is large is true or false only in relation to one or another comparison class, so… the judgement that it is wrong of someone to do something is true or false only in relation to an agreement or understanding". But this very statement of the intended analogy makes clear where it falls apart: a comparison class is not the same thing as an agreement or understanding. For the analogy to work, Harman's position would have to be that the judgement that someone ought not to do something can be true or false only in relation to one or another *comparison class*.

But this would get Harman nowhere. To see this, let us continue his analogy while avoiding this confusion. We start with the notion that, just as no dog is large, period, but only large in relation to a certain comparison class, no action is bad, period, but only bad

---

[168] Harman and Thomson 1996, p.3. Harman offers a similar account in Harman 2008a.

in relation to a certain comparison class.[169] From there, the correct continuation would be to say that just as a certain dog might be smaller than most dogs but larger than a chihuahua, a certain moral action might be morally better than most actions but worse than others. For instance, one might say that shoving an innocent stranger against a wall is morally worse than giving the stranger a dirty look, but morally better than killing that stranger. But this, of course, is something that everyone accepts and that non-relativists in particular would have no reason to deny.

A defender of Harman might, in reply, point to recent efforts to articulate a more general (i.e. not merely ethical) relativism that is actually consistent with objectivism. These forms of relativism assert only that all truths are relative to someone's actual or possible perspective; they leave open the possibility that there are objective truths. Objective truths, on these relativistic accounts, are simply true in *all* perspectives, and hence still true relative to each perspective individually.[170] Similarly, one might attempt to interpret Harman's relativism more charitably by assuming that he means only to say that every sincerely held moral statement is true given someone's perspective, regardless of whether or not there are also some objectively true moral statements (i.e. statements that are true on every actual, or indeed every possible, perspective).

However, there are excellent reasons for rejecting this attempt at charitable interpretation. First, since the resulting view would be indistinguishable from anything a moral non-relativist would say, it would be surprising if anyone would bother to propose it as an

---

[169] Incidentally, it is not even entirely clear (to me) that a dog cannot just be large, period. But my criticism does not rest on this point.

alternative view (let alone an interesting or important one) of morality. Second, Harman has said nothing to support this interpretation of his view, and many things that blatantly contradict it. For example:

> According to meta-ethical relativism, there can be conflicting moral judgements about a particular case that are both fully correct. The idea is that two people with different moralities might reach conflicting moral judgements concerning a particular case – for example, one saying the agent was morally right, the other saying the agent was morally wrong – where both opinions are correct.
>
> The two judgments must really conflict and therefore must be judgements about the same particular case. It is not enough to contrast a judgement about an act in one society with a judgement about the 'same act' in another society.
>
> Furthermore, it is not enough for meta-ethical relativism that there should be such conflicting judgements. Both judgements must be correct. In particular, neither judgement can rest on mistakes about the facts of the case, a failure fully to appreciate these facts, incorrect reasoning, or anything else that might distort someone's judgement.
>
> Finally, both judgements must be fully correct. Borderline cases are not enough, cases in which one might be equally justified in saying one thing or the opposite. According to meta-ethical relativism, the one judgement is fully correct for the one speaker, the conflicting judgement is fully correct for the other speaker, and it would not be correct for either speaker to make the other judgement.[171]

If there could be two such speakers making two such judgments, both correct, then it would follow that there is no objectively correct answer to that issue. Hence, it follows that, on Harman's account, there are no objectively correct judgments on those issues that give his relativism its character.

---

[170] See Hales 1997, Hales 2006 and Kolbel 2007 for examples of this attempt.
[171] Harman 2000b, p.24

To take this clarification further: either Harman thinks that *some* moral claims are objectively true or false while others are not, or else he thinks that *no* moral claims are objectively true or false. We have already seen (in Chapter One) that on Harman's account, all morality is dependent on both socially enforced norms *and* personal principles, both of which must be in agreement. Hence, the only way that Harman could accept the existence of any objectively true moral statements would be if he accepted that there are some universally socially enforced norms *and* some universally-held personal principles *that coincide with the universally socially enforced norms*. However, Harman actually holds that there are no instances of universally held norms or principles. In his 'Moral Philosophy and Linguistics', he argues that "It might be thought that there is a universally recognized central core of morality involving at least prohibitions against killing and harming others, against stealing, and against lying to others. But whether this is correct would seem to depend on what counts as acceptance of the same principle in different societies."[172] And even these things, Harman says, change from society to society: rules against harming, stealing from, lying to and even killing others depend (among other things) on whether the 'others' in question belong to a specially-privileged group, the constitution of which changes from society to society. Harman concludes that "[i]t may be true that rules against killing, harm, lying and cheating occur in all societies. But there is a sense in which these are not exactly the same rules, since the protected group changes from one society to another."[173] He goes on to emphasize the presence of similar disagreements over individually-held principles.[174]

---

[172] Harman 2000f, pp.217-218

[173] *Ibid.*, p.218. See also Harman and Thomson 1996, pp.9-10.

Now, since Harman holds that even those moral rules that are the most likely to be objective turn out (by his standards, and on his evidence) not to be objective, it follows that he holds that there are no objective, first-order moral facts.[175] The moral version of the Hales/Kolbel interpretation of relativism, insofar as this interpretation can help Harman in the way I have mentioned, entails that there *are* objective, first-order moral facts. Therefore, Harman cannot correctly be understood as endorsing an anemic relativism of this type. Hence, this attempt to defend his analogy would not be successful.

Ib: Einstein's Relativity

The comparison with Einstein's theory of relativity is also flawed. It is true that, according to the theory of relativity, there is no specially privileged physical framework. It is also true that, in alternative but equally acceptable frameworks, the same body can be seen as at rest or in motion, and the mass of an object may vary. But Harman unfortunately breaks off his comparison here, after having endowed his relativistic view with the aura of scientific credibility but before having examined just what a truly parallel view of morality would look like.

What Harman seems to miss about relativity theory is that – to misquote Richard Dawkins – however many ways there are of describing a given set of physical events correctly, it is certain that there are vastly more ways of describing those same events

---

[174] *Ibid.*, pp.218-219

[175] I.e. facts of the form 'X ought to do Y', 'It would be wrong for V to do Z', etc., as opposed to facts *about* these first-order facts, such as the fact (if it is one) that morality is not wholly relative.

incorrectly.[176] Regardless of which spatio-temporal framework one employs in describing a particular set of physical events, there must be a precise and direct mapping of every detail from any spatio-temporal framework onto any other, or else at least one of those frameworks will be wrongly constructed according to relativity theory. This is not at all parallel to what moral relativism tells us, since the sorts of moral views licensed under moral relativism are not so tightly restricted.

Further, once we have chosen a spatio-temporal framework different from our own and said, for instance, that this sack of flour is at rest and weighs 7 kilograms, everything else in that spatio-temporal framework will be fixed. Again, if Harman's moral relativism says nothing more than the parallel to this[177], then it has nothing to say aside from what non-relativists already accept.

Finally, every spatio-temporal framework that correctly accounts for the facts (and this correctness, according to relativity theory, is an entirely *objective* matter) will have many things in common with every other viable spatio-temporal framework. For instance, *every single one of the successful spatio-temporal frameworks implies precisely the same descriptions and predictions as one another*, when adjusted for framework differences. If a rock is thrown toward a window, then either that window will smash in every single one of the successful frameworks or else it will smash in none of them. If Harman wishes his moral relativism to work along these lines, he will have to say parallel things

---

[176] Dawkins 1996, p.9

[177] I.e. that from the correct assignment of a truth-value to any first-order moral claim in any moral framework, together with any complete assignment of truth-values from any *other* moral framework, the truth-values of all other moral claims in the first moral framework will be fully determinable.

about morality. But of course he does not want to say these things, since that would rob his moral relativism of any interest. Nor, for all the reasons already given in the discussion of Harman's 'contrast class' analogy, can it be held that it was Harman's *intention* to articulate a form of relativism that is compatible with the claims of moral objectivists. Hence, this analogy fails also.

I turn now to what I consider to be the most interesting of Harman's analogies with science, which lies at the heart of what some have called the 'Harman/Sturgeon debate'.[178]

Section II: The Relationship of Fact and Theory in Science and Morality

This third parallel has received the most attention from current philosophers, but some of those who discuss it seem unaware of just how (or even whether!) Harman intends it to contribute to his overall argument for moral relativism. In this section, I will first explain Harman's argument in some detail, making clear how he sees it as pertinent to his case for relativism; then, I will set out and critically examine the most prominent responses to the issues Harman raises. The writers of these responses, interestingly, tend to apply the approach solely to the topic under discussion in this section. My own 'parallel arguments' strategy differs not only in this respect, but also in two other features already noted: a reluctance to commit myself to any particular form of moral realism, and a

[178] This debate consists of four exchanges, two from each philosopher: Harman 1977, Sturgeon 1985a, Harman 1986, and Sturgeon 1986.

limitation of the approach to the context of the *practical* sort of philosophy.[179] As I will argue later on, these features seem to increase the power and range of the approach.

To get to the parallel arguments, however, we need to take a trek through the sometimes labyrinthine forests of the Harman/Sturgeon debate. This debate begins with an argument of Harman's that I skimmed over in Chapter One, and proceeds along paths that can quickly lead one astray in the darkness. When one crosses a thick and unfamiliar forest in the dark, one's best bet is to look skyward periodically and to use the North Star to navigate. We will have a chance, in good time, to find a North Star of our own to guide us through these pathways. But first, let us examine Harman's argument more fully.

Harman's Argument

Harman asks the reader to consider the following two cases:

1) A scientist looks into a cloud chamber and sees a vapour trail. The scientist understands that, given the way the cloud chamber is set up, it is most reasonable to assume that this vapour trail was caused by a proton. So the scientist says to himself, 'There goes a proton!'[180]

[179] See pp.11-15 of the Introduction.
[180] Harman 1977, p.6

2) Walking around a corner, you come upon some juvenile delinquents pouring gasoline on a cat and lighting it on fire. You think there is something morally wrong with this. You say to yourself, 'That action is immoral!'[181]

Harman claims that these two cases, which he clearly thinks are parallel to one another, show an important difference between the relevance of ethical observations to ethical theories and the relevance of scientific observations to scientific theories. He explains this difference as follows:

> Observational evidence plays a part in science it does not appear to play in ethics, because scientific principles can be justified ultimately by their role in explaining observations.... Apparently, moral principles cannot be justified in the same way. It appears to be true that there can be no explanatory chain between moral principles and particular observings in the way that there can be such a chain between scientific principles and particular observings. Conceived as an explanatory theory, morality, unlike science, seems to be cut off from observation.[182]

The problem, that is to say, is that we can tell an unbroken causal story connecting the existence of the proton with the confirmation of a physical principle, but cannot tell one linking the existence of the immorality of mistreating the cat with the confirmation of an ethical principle. The scientist comes to believe that some physical principle was confirmed *because* he saw a vapour trail, and he saw a vapour trail *because* there was a proton whose activity caused it; so the proton (partly and indirectly) caused the confirmation of the principle. But the immorality in harming the cat does not cause anything: we could just as easily explain your belief that the act is immoral by referring

---

[181] *Ibid.*, p.4
[182] *Ibid.*, p.9.

to your inner psychological states and nothing else. Hence, "the explanatory chain from principle to observation seems to be broken in morality."[183] And if there is no unbroken explanatory chain extending from principle to observation, then we cannot ever be justified in accepting some principle. Therefore, assuming there are no cases of principle-confirmation in ethics that somehow get around this problem, we are never justified in accepting ethical principles.

Harman never makes clear how this argument is meant to tie in with his overall argument for relativism (a fact that has led some of his critics to suppose that he is arguing instead for moral nihilism[184] or moral scepticism[185]). Presumably, however, the following is a fair reconstruction of what he has in mind: the above considerations leave us to choose between a) accepting the view that we are not justified in accepting any of our moral principles and b) finding some way to confirm moral principles that gets around the problem Harman takes himself to have found. We have good reason to reject the first of these options, because it would conflict with our everyday ways of thinking and speaking. Therefore, we need to find some other acceptable way to confirm our moral principles. Noncognitivism, functionalism, and the other options Harman rejects in Chapters 2-7 of *The Nature of Morality* fail for the reasons given in those chapters. That leaves only relativism, which can overcome this problem because it entails that the ultimate basis of morality lies in a combination of social customs and personal concerns (which we can learn about quite easily through observation or introspection, and which

---

[183] *Ibid.*, p.8

[184] See, e.g. Koons 2000, p.247: "Harman infers from this that there are no moral facts, and that some form of moral nihilism must therefore be true."

[185] See, e.g. Miller 2003, p.140: "Harman's challenge thus threatens us with a version of moral skepticism."

therefore leave room for an unbroken explanatory chain to extend from our knowledge of moral facts to our moral views, as Harman wishes).

Sketching out Harman's overall case in this way helps us to see that his 'broken chain' argument, even if successful, would only get us to moral relativism if the other premises of his overall case were acceptable. But as I hope to have shown in Chapter One, the arguments Harman relies on to establish the rest of his case for relativism are seriously flawed in crucial places. Therefore, what remains is at best an argument for moral skepticism. Now, the question of moral skepticism lies beyond the scope of this dissertation: while I feel the parallel arguments approach will be equally useful in answering moral skepticism as moral relativism, there is too much literature on the moral skepticism issue for one to deal with both it *and* relativism in a comprehensive fashion in a work of this size. Nonetheless, although the above considerations show that a decisive response against Harman's observation problem is not needed to answer his moral relativism, I will outline how his 'broken chain' argument for moral skepticism can be overcome by the parallel arguments approach.

Before doing this, however, I would like to point out a third independent way in which Harman's argument goes wrong. Harman's prominent critics, for all the strength of their responses, seem to have overlooked the fact that the initial comparison that motivates his argument rests on yet another faulty analogy.

Problems with Harman's analogy

There are, in fact, two different problems with his analogy to be dealt with here. The first problem is that the observers in the two stories are doing two quite disanalogous things. The scientist, looking at the cloud chamber and seeing the vapour trail, *concludes* that there is a proton in the cloud chamber. You, on the other hand, round the corner and say to yourself – *but do not 'conclude' on the basis of your observation* – that it is immoral for the hoodlums to burn the cat. The scientist, before looking at the cloud chamber, does not know that there was a proton there: his observation is what gives him that information. But in the story, you already know that it is immoral to burn a cat before making the observation:[186] all you observed was an instance of something you already knew had the property of being immoral.

We can correct this defect in the analogy by adjusting the scientific case to match the ethical one. To do this, we need a case where someone who knows some science observes something that he or she recognizes on the basis of that scientific knowledge to be an instance of something else, just as you, when rounding the corner, recognize the burning of the cat to be an instance of an immoral action on the basis of your moral knowledge. And it would be question-begging, at this point, to assume that non-relativistic moral knowledge is unattainable. Let us suppose, for instance, that a different scientist (mine is female, to avoid confusion with Harman's) rounds a corner and observes a cat crossing the street. She says to herself, 'There goes a mammal!'

---

[186] At least, *I* already know that it is immoral to burn a cat, despite never having seen such a thing done.

We can now consider whether there is an unbroken explanatory chain extending from the cat's being a mammal to the scientist's judgment that it is one. Perhaps the scientist's young relatives want an explanation of why she thinks that the cat is a mammal. "How do you know?", they ask. Now, will it do for the scientist to respond, "I believe it, in part, because it *is* a mammal"? Perhaps, or perhaps not: certainly, this would be an unusual way of talking. But I see no reason to think that using the claim that the cat is a mammal as part of the explanation of the scientist's belief is any less strange than using the claim that burning a cat is wrong as part of an explanation for your *moral* belief.

We can now see that Harman's 'problem' depends on the truth of the following claim: *If I ask someone why she believes that a cat is a mammal, then it would be legitimate for her to cite the fact that a cat is a mammal as part of her explanation; but if I ask someone why she believes that burning a cat is immoral, it would not be legitimate for her to cite the fact that burning the cat is immoral as part of her explanation.* Now, Harman does not offer a reason why anyone should accept that claim, and clearly one who does not accept that claim should not be persuaded by Harman's argument (once it has been made relevantly parallel in the way I have suggested). His argument only works, therefore, if his audience already has an independent reason for thinking that moral facts are irrelevant to our moral beliefs in ways that scientific facts are not irrelevant to our scientific beliefs. Since he never goes on to explain why this irrelevance is a problem, but merely assumes that his audience will agree that it is one, it follows that his argument only works if his audience already believes that moral facts are irrelevant to moral beliefs *and* that this is a problem for ethics. In other words, it follows that Harman's argument for the conclusion

that there is a problem with ethics only works if his audience already believes that there is a problem with ethics. But this suggests that Harman's argument commits the fallacy of begging the question: it only works if his audience accepts the very assumptions the argument is meant to establish.

The second problem with Harman's analogy has to do with the way in which he wishes the conclusions arrived at (i.e. 'There goes a proton'/'That action is immoral') to serve as possible confirmation-tests for principles. The problem is that the term 'principle' is ambiguous, and Harman seems to use it in two different ways. This can be seen if we consider what Harman thinks of as a moral principle and a scientific principle, respectively.

Harman never defines 'moral principle', so we can best understand what he means by the term by considering the single example he provides to illustrate his point:

> Consider the principle that, if you are given a choice between five people alive and one dead or five people dead and one alive, you should always choose to have five people alive and one dead rather than the other way around. We can easily imagine examples that appear to confirm this principle...[187]

What Harman means by the term appears to be some sort of maxim or rule (whether a hard and fast rule or a rule of thumb seems to be left open) that tells us what we *should* do in at least very specific situations (e.g. in situations where we are "given a choice between five people alive and one dead or five people dead and one alive"). Perhaps

---

[187] Harman 1977, p.3

Harman would also include among these moral principles those that give more general advice, like "Act always according to that maxim which you could at the same time will to be a universal law of humanity."

Since Harman's question here is "Can moral principles be tested in the same way [i.e. as scientific ones], out in the world?" his analogy must involve a parallel understanding of the term 'scientific principle'. Harman never defines that term and never gives an example of something he calls a scientific principle. Instead, he talks about "scientific hypotheses"[188] and "a scientific theory".[189] The context suggests he may be using all three of these terms interchangeably; but even so, he does not make clear what sort of thing he is referring to by any of them. The scientist who examines the goings-on in the cloud chamber is presumably testing *some* hypothesis, but it is never made clear which hypothesis is at issue or how it is being tested. Still, it seems fair to assume that he is thinking of something along these lines: 'Protons must have a greater range than current physical theory assumes'.

With these ideas in mind, let us consider the rest of the analogy Harman wishes us to accept:

> ...in both science and ethics, general principles are invoked to explain particular cases and therefore, in both science and ethics, the general principles you accept can be tested by appealing to particular judgments that certain things are right or wrong, just or unjust, and so forth; and these judgments are analogous to direct perceptual judgments about facts.[190]

---

[188] *Ibid.*, p.4
[189] *Ibid.*, p.6
[190] *Ibid.*, pp.5-6

But is it really at all clear that "general principles are invoked to *explain* particular cases" in ethics? Why would this be, if ethical principles (as Harman's own example of the 'five versus one' principle makes clear) tell us what we *should* do rather than what *will* happen? Just what would it mean for a statement of the form 'Someone in situation Y should do X' to be "invoked to explain particular cases"? It also seems unlikely that Harman means that the ethical and scientific principles we hold explain, psychologically, why we make the judgments that we do. For on that interpretation, the explanation in both the ethical and the scientific case would be *trivially* explanatorily adequate.

Perhaps, if Harman's analogy were appropriately corrected and clarified, these two objections could be met and the argument could stand. However, regardless of whether some defender of Harman's view could restate his analogy more adequately, it remains true that it is not adequately put *now*. And since Harman is the one who wishes to present a major problem for ethics (indeed, *the* problem with ethics)[191], he bears the burden of presenting his argument in a way that makes clear that there is a problem at all. As the above reasoning seems to show, this has not yet been done. Thus, a careful examination of Harman's faulty analogy seems to provide us with all we need to get around his 'problem with ethics'.

However, it also seems that Harman's clear, non-technical style of writing, conjoined with a certain freedom in the way he sets out his arguments, has succeeded in throwing many otherwise thorough philosophers off the scent. Rather than consider the

assumptions behind Harman's arguments, which seem too simple to be problematic in their own right, it seems tempting to accept them as given and seek out some adequate response in the territory as mapped out *by* those assumptions. But as we will continue to see, Harman's innocent-seeming assumptions are often inaccurate, unjustified, or even inconsistent; and so the territory quickly becomes unnecessarily difficult to navigate if one follows his map. In some cases, as will soon become evident, these illusions befall even those philosophers who have succeeded in finding some, but not all, of the faults in Harman's metaphors.

It is within that illusion-plagued landscape, however, that the parallel arguments I have mentioned have been presented. So, while we already have an adequate response to Harman, let us venture deeper into the Harman/Sturgeon debate. Within a few more pages, we will be able to catch the glimpse of the North Star promised earlier, which will allow us to ignore Harman's map and navigate the terrain more reliably.

Sturgeon's Reply

Theodore Sturgeon's 1985 paper 'Moral Explanations' provides a different sort of response to Harman's argument from that which I have given above.

Sturgeon's Preliminary Clarifications

Sturgeon begins by pointing out that Harman's argument would not be very impressive or important if it only reasserted familiar verificationist worries. There seem to be good reasons, Sturgeon feels, to believe that ethical statements will do just as well as scientific

---

[191] *Ibid.*, p.1

ones against typical verificationist attacks;[192] and at any rate, if Harman's concerns collapse back into verificationist ones, there will turn out to be no *distinct* problem to be dealt with in Harman's argument.[193]

Sturgeon also points out that "there are several ways in which Harman's argument invites misunderstanding"[194] due to the unclear way in which much of it is presented, and that the issue needs to be clarified before it can be dealt with properly. One of the sources of misunderstanding, he points out, lies in Harman's "focusing at the start on the question of whether there can be moral *observations*", whereas in fact "this question turns out to be a side issue, in no way central to his argument that moral principles cannot be tested against the world".[195] In fact, the real issue Harman raises has to do with whether there is an unbroken explanatory chain extending from moral facts to our moral beliefs, whether or not the belief-formation prominently includes observations.

Another way in which Harman's account invites misunderstanding arises from his claim that the existence and accessibility of moral facts are not *needed* to explain moral beliefs. Surely, Sturgeon admits, it is true that one can always find some way to account for the fact that someone has a particular moral belief without invoking the existence or accessibility of moral facts. But that is not very interesting – after all, it is also (trivially) true that we can account for *any* beliefs someone happens to have without invoking the existence or accessibility of the corresponding facts. Nor, although Sturgeon does not

---

[192] Sturgeon 1985, pp.231-232.
[193] *Ibid.*, p.233
[194] *Ibid.*, p.234
[195] *Ibid.*

make this point, is this only true of explanations about beliefs: one could give an explanation for absolutely anything while scrupulously avoiding the invocation of any particular set of facts, *provided that one does not care whether the explanation is any good*. In order for Harman's claim to be an interesting one, he must say more – he must say, in this case, that the best explanations that involve reference to non-relative moral facts are no *better* than alternative explanations that make no reference to non-relative moral facts. Nor, as Sturgeon points out, does Harman feel it is a matter of the latter beating the former by a negligible amount: Harman claims that moral facts are "completely irrelevant" to any explanation of our observations and beliefs. In order to defeat Harman's claim, therefore, it is not necessary to show that no explanation of moral beliefs that does not involve reference to moral facts can be successful. One need only show that, at least in some cases, reference to some moral facts in explaining someone's holding a moral belief is *relevant*.[196]

The third of Sturgeon's preliminary clarifications is a particularly important one.[197] To respond fairly to Harman's argument, we must extend the principle of charity to him. Now, if Harman commits the fallacy of begging the question in presenting his argument, then it collapses easily, and no further discussion is needed. So following the principle of charity, we should imagine (if possible) that he is not begging the question. If his argument is not begging the question, then it must be the sort of argument that would be effective even against those who do not already accept its conclusion. The conclusion of Harman's argument is that there is no unbroken explanatory chain extending from moral

---

[196] *Ibid.*, pp.236-237

facts to our moral beliefs. Those who deny the existence of moral facts will, quite obviously, already doubt whether there is an unbroken explanatory chain extending from moral facts to our moral beliefs; and this should also be apparent to them. Therefore, such people are in no need of Harman's argument; and if Harman's argument is only meant to persuade them, then it begs the question. Therefore, we must interpret Harman's argument as answering the concerns of those who believe there *are* moral facts: as Sturgeon puts it, he "can see no way of evaluating the claim that *even if* we assumed the existence of moral facts they would still appear explanatorily irrelevant, without assuming the existence of some, to see how they would look."[198] However, Sturgeon makes a point to remind his readers that he recognizes that this assumption is only a provisional one: "Does this prove that there *are* such facts? Well of course it helps support that view, but here I carefully make no claim to have shown so much."[199] We may, ultimately, be entitled to reject this supposition, if we find that assuming it leads to a contradiction; and it may be that Harman has some other argument in mind against the existence of moral facts, on the basis of which we would not be justified in assuming their existence in the first place. But the first of these would have to be *shown*, which has not been done; and the second would, Sturgeon claims, render Harman's argument pointless anyway since we would then be persuaded of its conclusion before it begins.

---

[197] Sturgeon discusses this point in *Ibid.*, pp.237-238. His way of expressing the idea is quite different from mine: I hope that my summary does it justice.
[198] *Ibid.*
[199] *Ibid.*, p.238

For reasons of space, I will not summarize all of Sturgeon's fourth preliminary clarification here[200], but only one point in it that seems especially relevant. The point is this: Harman, as we have seen, feels that "there could be moral facts if these facts were reducible in some way or other to other facts of a sort that might help explain observations."[201] Sturgeon himself is a non-reductive naturalist: he feels that moral facts *are* natural facts, despite the fact that they may not be reducible to other natural facts. To him, these moral/natural facts do not *need* to be reduced to anything else in order to be able to explain observations. But he is careful in his response to Harman not to assume anything about whether moral facts will be reducible to facts of other kinds:

> A naturalist… will certainly want (and a critic of naturalism will likely demand) a fuller account than this of just where moral facts are supposed to fit in the natural world. For all I have shown, moreover, this account might even provide a scheme of reduction for moral discourse: My argument has been not that ethical naturalism could not take this form, but only that it need not.[202]

Putting all these preliminary clarifications together, we can see what it is that Sturgeon needs to do in order to defeat Harman's argument head-on: he needs to find just one case where, in explaining someone's moral belief (whether or not some moral observation was involved in the formation of that belief), it is relevant to refer to some moral fact (whether or not that moral fact is in turn reducible to nonmoral natural facts); and furthermore, Sturgeon is to be entitled to the preliminary, tentative assumption that there are moral facts.

---

[200] I will, however, touch on some of the points Sturgeon makes there later on.
[201] Harman 1977, p.21

Sturgeon's Counterexample

To accomplish this task, Sturgeon suggests that we consider a rather despicable character from history, one Selim Woodworth. Woodworth was a young officer who volunteered his services as overseer of a rescue operation in the snowy mountains. Having received this posting, however, Woodworth apparently did almost nothing toward fulfilling the important duties that came with it: instead, he "spent time arranging comforts for himself in camp, preening himself on the importance of his position; and as a predictable result of his cowardice and his exercises in vainglory, many died who might have been saved, including four known still to be alive when he turned back".[203] Not only did he fail to lead the rescue parties himself; he even "had to be 'shamed, threatened and bullied' even into organizing the efforts of others willing to take the risk".[204] In spite of his having done almost nothing toward the fulfillment of his duties, even under pressure, Woodworth was unable even after the fact to understand the role his professional negligence played in the disasters that befell those he was duty-bound to help as the following story suggests: One John Stark, who saw the deficiencies in Woodworth's character and actions and stepped in to do the job Woodworth was paid to do, saved a Mrs. Breen. "When after great hardships Stark brought Mrs. Breen to Woodworth's camp, she mentioned how they had suffered. Woodworth said to her you may thank me Mrs. Breen for your safe delivery. Thank you I thank no boddy but God and Stark and the Vergin Mary she said. Putting Stark second best and I think he deserved it. [*sic*]"[205]

---

[202] Sturgeon 1985, p.241
[203] *Ibid.*, p.244
[204] *Ibid.*

Sturgeon attempts to put this example to work as follows: Bernard DeVoto, a historian who discusses this episode, claims that "Passed Midshipman Woodworth was just no damned good."[206] Sturgeon holds that DeVoto's conclusion is a perfectly reasonable one, given the evidence:

> It isn't, for example, that Woodworth was a basically decent person who simply proved too weak when thrust into a situation that placed heroic demands on him. He volunteered, he put no serious effort even into tasks that required no heroism, and it seems clear that concern for his own position and reputation played a much larger role in his motivation than did any concern for the people he was expected to save. If DeVoto is right about this evidence, moreover, it seems reasonable that part of the explanation of his believing that Woodworth was no damned good is just that Woodworth *was* no damned good.[207]

But if Sturgeon is correct about this, then – he argues – he has found a case in which, in explaining someone's moral belief, it is relevant to invoke the existence of some moral fact; and this is just what he needed to do in order to defeat Harman's argument.

Sturgeon also provides a second example: there was greater and more widespread opposition to slavery in Britain, France, and English-speaking North America in the eighteenth and nineteenth centuries than in other places and times. What explains the belief that motivated that opposition? Sturgeon suggests that these beliefs may well have been due to the fact that the slavery really was a great deal worse in those places and times than in others, and that many opponents of slavery came to notice this fact.[208] This,

---

[205] These last details are from Johnson 2006. In the end, Woodworth managed to get himself elected to the state senate, while Stark perished in relative obscurity
[206] Sturgeon 1985, p.244
[207] *Ibid.*
[208] *Ibid.*, p.245

then, would be – Sturgeon thinks – another case of a moral fact being relevant to a plausible explanation of a moral belief.

Despite his having been careful up to this point, Sturgeon has made a misstep here. He has read, entirely correctly, Harman's statement that "morality, unlike science, seems to be cut off from observation"[209], so he offers examples of ways in which a moral belief or theory can be informed by an observed fault or virtue of some action, institution, or character. However, as Sturgeon himself noted earlier, the observation issue is a red herring. The above statement by Harman is inconsistent with many of his other claims, which (I will argue presently) locate the problem not as lying between the observation and the theory, *but rather as between the moral properties of some object and its non-moral properties*. Moreover, it is only the second of these interpretations of Harman's concern that has any bite: the first would actually render his overall case for relativism seriously inconsistent since (as I will now show) Harman *himself* gives a strong response to the first observation-related concern. Therefore, the principle of charity demands that we interpret Harman as meaning something other than what Sturgeon presently takes him to mean; and, as will be seen, Sturgeon himself comes to recognize this point later on in his paper. However, while Sturgeon ultimately returns – for this paper, at least – to the correct path, the wrong path he begins to take seems to have been a tempting error for many other important philosophers writing on the Harman/Sturgeon exchange.

What I take to be the error should become clear when we recall that, in *The Nature of Morality*, Harman admitted that there can be "a good farmer, a good soldier, a good

teacher, a good citizen, [or] a good thief."[210] In addition to saying that our use of words like 'good' and 'bad' in non-ethical contexts avoids what he has identified as 'the problem of ethics', Harman clearly endorses as correct and unproblematic the use of other terminology that is sometimes used in ethical contexts: "the words 'ought' and 'wrong' are relevant as before. During battle, we say, a soldier ought to obey his superior officers without question. It is wrong for a teacher to play favorites. A thief ought to wear gloves."[211] The difficulty, Harman thinks, only arises when we try to link up non-moral properties (including the property of 'goodness' in the sense described above) with moral properties, like *moral* goodness. Harman's own statements on pp.14-17 of *The Nature of Morality* cannot, it seems to me, plausibly be construed in any other way.

Hence, Harman has a powerful response available against Sturgeon's counterexamples: it is most consistent with what Harman has said about teachers, soldiers and thieves for him to hold that Woodworth's having been "no damned good" is an unproblematic *non*-moral property. The same goes for Sturgeon's second attempt at a counterexample: Harman can safely agree that the slavery in question was much more *cruel*, more *onerous*, and more *troubling* to both slaves and others in these places; and he can even agree that it was 'worse' in the sense that it was all these things. Again, all this is consistent with Harman's way of disambiguating such terms as 'good' and 'wrong'. The explanatory chain extending from Woodworth's being no damned good to DeVoto's judgment that he was no damned good is unbroken, Harman can admit: the question for Sturgeon to

---

[209] Harman 1977, p.9
[210] *Ibid.*, p.15
[211] *Ibid.*

answer is how this chain extends from the fact that Woodworth was no damned good to the fact that he was *immoral*.[212]

Once this point is recognized, it becomes difficult to see how Harman's argument differs from J. L. Mackie's famous 'argument from queerness'. In his *Ethics: Inventing Right and Wrong*, published in the same year as Harman's *The Nature of Morality*, Mackie sets out that argument as follows:

> [One] way of bringing out this queerness is to ask, about anything that is supposed to have some objective moral quality, how this is linked with its natural features. What is the connection between the natural fact that an action is a piece of deliberate cruelty – say, causing pain just for fun – and the moral fact that it is wrong? It cannot be an entailment, a logical or semantic necessity. Yet it is not merely that the two features occur together. The wrongness must somehow be 'consequential' or 'supervenient'; it is wrong because it is a piece of deliberate cruelty. But just what *in the world* is signified by this 'because'? And how do we know the relation that it signifies, if this is something more than such actions being socially condemned, and condemned by us too, perhaps through our having absorbed attitudes from our social environment?[213]

Given that Harman is happy to accept the existence of an unbroken explanatory chain between the non-moral goodness of a teacher and our belief in that goodness, it seems impossible that he could mean to point to a break in the chain anywhere *but* where

---

[212] Admittedly, this is in some ways a strange question to make Sturgeon answer: Sturgeon, after all, holds that "moral facts are nothing but natural facts" (Sturgeon 1985, p.239), so to ask him to establish the link between moral facts and these sorts of natural facts is to ask him to establish the chain between something and itself. Perhaps Sturgeon could make the complaint that, in this context, that request is illegitimate. The point, however, is that Sturgeon does not do this with his counterexamples: the counterexamples simply assume that the cruelty of slavery and the cowardice of Woodworth are straightforwardly immoral, despite the fact that it is just this assumption upon which Harman means to cast doubt. For that reason, Sturgeon's counterexamples beg the question.

[213] Mackie 1977, p.41

Mackie does: in the link between non-moral goodness and moral goodness, or (to use Mackie's example) in the link between deliberate cruelty and moral wrongness.

This is the key point of Harman's challenge; but it seems (perhaps, as Sturgeon says, because Harman writes in a way that seems to invite misinterpretation) very easy to confuse it with something else. So as to ensure that *we* do not go astray, let us keep our eye on this point, which I will hereafter call the 'North Star'.

The North Star: The 'problem with ethics' Harman raises is the apparent problem of a missing link between the moral and non-moral features of some action, person, character trait or thing.

Quinn's Objections against Sturgeon

Later on in his first paper, as we will see, Sturgeon attempts to solve Harman's 'problem with ethics' – the correct one – by means of a parallel argument concerning science. It is unfortunate that, as just noted, Sturgeon's attempt to deal with Harman's problem head-on by providing a counterexample or two misses the target. To ensure that we are not missing some important problems by jumping across the explanatory gap with Sturgeon, let us consider the response of Warren Quinn, who cries foul at Sturgeon's leap.[214] Quinn's paper, like Sturgeon's, at times seems to lose the thread of discussion,[215] but at other times deals with the central issue quite clearly.

---

[214] Quinn, 1986

[215] For instance, on p. 534, Quinn temporarily mistakes Harman's issue for the issue of whether "it was the abolitionists' verdict on slavery, and not ours, that explained its demise." Clearly, though, neither Harman nor Sturgeon is interested in whether anyone's verdict – whether ours or the abolitionists' – *explains* the

Quinn wonders what could possibly stand in as a proper link between Woodworth's cowardice and Woodworth's immorality.[216] The answer to the question – and its plausibility – depends, he suggests, on whether the term 'coward' should be taken as a psychological or a moral term. The explanatory gap between the moral and non-moral properties of some event, person, action or thing (which I will hereafter call the 'Harman-Mackie gap') cannot be bridged by a specifically moral term, since that would leave unanswered the question of *how* such a moral term can connect with a non-moral one; and the mirror-image of this problem arises if we try to bridge the gap with a specifically non-moral term. A third possibility that suggests itself is that some term, like 'coward', might serve as both a psychological *and* a moral one, so that to attribute cowardice to someone would automatically be to bridge the Harman-Mackie gap. But in this case, Quinn argues, the property referred to by the term would no longer be a *necessary* part of the explanation. For suppose that Sturgeon claims that Woodworth is cowardly, and that the term 'cowardly' attributes both a moral and a psychological property to Woodworth. Someone else – let's call him Schmurgeon – might agree that Woodworth had the psychological property that Sturgeon attributed to Woodworth, but might feel nonetheless that there is nothing particularly immoral about being a coward, despite the fact that (as Sturgeon and Schmurgeon both agree) it was Woodworth's cowardliness that led to the disasters that befell those trapped in the mountains. In that case, Quinn suggests, Schmurgeon's explanation of DeVoto's judgment that Woodworth was immoral will

---

demise of slavery; but rather on whether someone's verdict can be explained *by* some moral fact. It is also puzzling – though less relevant – that Quinn seems to think anyone is implying that our verdict today can explain the success of some social action in the past.

[216] *Ibid.*, pp.534-535

seem just as good as Sturgeon's. Schmurgeon, like Sturgeon, will cite a) Woodworth's cowardliness and b) DeVoto's prior conviction that cowardliness is immoral as part of the explanation for DeVoto's judgment; and the only thing Schmurgeon will omit – the moral claim – will not be missed. Therefore, the use of some term that implies both a moral *and* a non-moral property will not help here, because its explanatory power will go no further than that of a purely non-moral term.[217]

So what else could bridge the Harman-Mackie gap? It may be that it cannot be bridged. This is the line Quinn takes, though he is neither a moral skeptic nor a moral nihilist, let alone a moral relativist: he feels that the best option for the moral realist is to reject the explanatory requirement for morality altogether.[218] I will have more to say about this move later in the chapter. To complete our picture of this issue, however, it is worth considering Quinn's reasons for rejecting the other plausible option for bridging the gap: that of keeping moral and nonmoral properties on opposite sides of the chasm, and allowing some sort of relationship to connect them.

Quinn argues against the viability of this option as follows:

> …[w]e have a special reason to be skeptical about the possibility of reducing the moral to the nonmoral. Such reductions would not, presumably, be offered as analytic. And it is far from obvious that they can be modeled on familiar synthetic property reductions in the physical sciences, for example, that of heat to molecular energy. In these familiar examples, the informativeness of a reduction is explained by claiming that our everyday concept identifies the property in question by way of its nonstructural phenomenal features. We locate the presence of heat by the way it feels to us, a manner of identification that does not reveal anything of heat's

[217] *Ibid.* I am paraphrasing rather freely.
[218] *Ibid.*, pp.540-544

> fundamental structure. The parallel claim for moral badness would be that we recognize its presence by the way it feels or appears to us, its fundamental nature lying elsewhere. But what could these phenomenal features of moral badness be?[219]

There is much of interest here. Let us bypass a number of technical disputes by assuming with Quinn that this relationship *must* be a type-reductive one in the sense in which Quinn conceives it, and that it cannot be one of token identity or (more broadly) supervenience. If we give Quinn the benefit of these assumptions, would his argument *then* give us sufficient grounds for thinking that the gap cannot be bridged?

If so – and remember that this is as good as it gets for Quinn – then the parallel arguments strategy can be employed here, as follows: Quinn begins by claiming, in general, that the two points cannot be linked by a moral term, a non-moral term, a term with 'dual citizenship', or by any relationship connecting two terms on opposite sides of the gap. Harman – though, as we have seen, not Quinn – uses this to argue that there is a 'problem with ethics'. This provides an opportunity for the deployment of an initial parallel argument to the effect that, by that same token, there is a 'problem with psychology', a 'problem with biology', etc. For in these cases also, it is unclear how one could move from psychological to non-psychological facts; from biological to non-biological facts; etc.

Quinn (and Harman, whose case for there being a problem with non-relativistic ethics seems to depend on the success of something like Quinn's reasoning) could overcome this application of the parallel arguments approach by pointing to a disanalogy between

---

[219] *Ibid.*, pp.535-536

the two arguments: this is what Quinn attempts to do at the end of the passage just quoted. He suggests that any reduction of moral terms to non-moral ones (this reduction being, for the moment, what we are charitably assuming to be needed here) must be along the lines of the reduction of heat to molecular energy, which he takes to be straightforward in a way that the reduction of moral terms cannot be. But there are problems with Quinn's move. For one thing, there are many different scientific definitions of 'heat'. First, it can mean "hotness, i.e. relatively high temperature".[220] Second, it can mean "energy that is transferred from one body to another as the result of a difference in temperature".[221] These two definitions are not coextensive: for instance, heat can be transferred to a container of ice from some other body as the result of a difference of temperature, but the temperature of the contents of the container will not increase until *all* the ice has melted.[222] Hence, there can be an increase of heat under the first scientific definition without a corresponding increase under the second. Nor are these the only proper scientific definitions: there are at least three others.[223] This is a problem for Quinn because our "everyday concept" of heat cannot correspond to five incompatible things unless there is something problematic in the concept itself. We might try to avoid this problem by taking a more lenient approach to reduction, and allowing for such incompatibilities. But then the parallel complaint against ethics would seem to be lost.

---

[220] Denker 2005, Section 15.1
[221] *Ibid.*
[222] *Ibid.*
[223] *Ibid.*

Quinn's best bet at this point would be, perhaps, to avoid this difficulty by identifying the everyday concept of heat with exactly *one* of the scientific definitions – for instance, as Quinn has it, with the definition of heat as 'molecular energy'. But this, also, will have its costs. For one thing, it will have the result of showing that some seemingly straightforward applications of our "everyday concept" of heat (including the application of that concept to the objects of our sensations according to their 'nonstructural phenomenal features') will be mistaken. This may be something Quinn – and Harman – would accept as a necessary evil. But if they do, their point against (non-relativistic) moral reduction would seem to be lost. For in that case, in order to avoid an uncomfortable parallel argument, their argument against the explanatory power of moral properties would need to imply an acceptance of the view that our everyday concepts may not survive a legitimate process of reduction fully intact. And in that case, they would lose their ability to complain that utilitarians (say) go wrong in identifying moral goodness with the not-always-intuitive formulation of 'the greatest happiness for the greatest number', which would rob their argument of all its force.

And that is only the beginning. If Quinn's argument is a good one, then it must follow that for *all* acceptable 'everyday concepts' that are also dealt with in the physical (and, presumably, social!) sciences, there must be a precise and unproblematic reduction from the properties found in the 'nonstructural phenomenal features' of our experience to the properties identified by whichever sciences are deemed acceptable. However, there are

many everyday concepts meeting that description that would seem to have an even worse time meeting this challenge than the concept of heat: colour concepts, for example.[224]

To summarize the situation so far: given the problems with the rest of Harman's overall case for relativism (as seen in Chapter One), his sub-argument concerning the explanatory chain from moral property to moral belief seems unlikely to succeed in bringing us closer to moral *relativism*: taken by itself, it leads (at best) to moral skepticism. If one nonetheless chooses to explore this argument, one finds that it is founded on a faulty analogy between *learning* that a free proton is in motion and *understanding* that burning a cat is wrong. However, one can charitably imagine another version of this argument – which would then be very (indeed, probably exactly) similar to Mackie's. The crux of this argument is the claim that there is a problematic, and unfulfillable, need to build an explanatory bridge between moral and nonmoral properties in order to avoid a problem with morality. Quinn and others get around this problem – quite successfully, it seems, as I will argue later on in the chapter – by accepting that the bridge cannot be built, yet maintaining that this is not problematic. The other option splits into two parts: first, one could (like Sturgeon) maintain that moral terms span the gap; or second, one could maintain that there is a non-relativistic reductive relationship that joins terms on both sides. In either case, it seems possible to use the parallel arguments approach to counter the worry that an explanatory bridge cannot be built. To do this, one only has to show that the prospects of building a bridge connecting moral and non-moral properties are at least as good as those of building a parallel one between

---

[224] See, for example, Hardin 1988

observable and unobservable entities or properties in science. And neither of these responses to Harman seems to have any difficulties at this point.

With this fuller sketch of the terrain in mind, let us return to Sturgeon's paper.

Sturgeon on Moral Counterfactuals

Though Sturgeon veers off the correct path in his discussions of Woodworth and slavery, he does, as I mentioned, return to it again later in his essay: he at one point recognizes that the break in Harman's explanatory chain must reside (if anywhere) in the Harman-Mackie gap. His discussion of the point begins on p.246, where he tries to determine what Harman could mean by suggesting that the moral wrongness of an action is not relevant to our judgment that it is morally wrong. Sturgeon puts the central question as follows:

> Is it true, as Harman claims, that the assumption that the children are really doing something wrong is 'totally irrelevant' to any reasonable explanation of your making that judgment? Would you, for example, have reacted in just the same way, with the thought that the action is wrong, even if what they were doing *hadn't* been wrong, and could we explain your reaction equally well on this assumption?[225]

As Sturgeon points out, the clause 'even if what they were doing hadn't been wrong' is ambiguous: it could mean 'even if they had been doing something *else* that wasn't wrong', or it could mean 'even if they had been doing the same thing in the same circumstances and for the same reasons, without this being wrong'. Sturgeon considers both of these possibilities in turn.

---

[225] Mackie 1985, p.246

First, let us suppose that the teenagers are doing something other than setting a cat on fire (or anything else that actually is morally wrong): perhaps they are simply playing an innocent game or doing their math homework. Would someone *then* react in the same way as in the scenario in which they were burning the cat instead? Presumably not; and if this is the way in which Harman means us to interpret him, then it is clear that the fact that they are doing something morally wrong is anything *but* irrelevant to our moral assessment of the teenagers.[226]

But Sturgeon, now following the guidance of the North Star, recognizes that this is not a plausible or consistent reading of Harman: the break in the explanatory chain, again, can only lie between the wrongness of an action and the non-moral characteristics *of just that same action*. Therefore, Harman's claim

> is not that if the action had not been one of deliberate cruelty (or had otherwise differed in whatever way would be required to remove its wrongness), you would still have thought it wrong. It is, instead, that if the action were one of deliberate, pointless, cruelty, but this *did not make it wrong*, you would still have thought it was wrong.[227]

This insight is what makes it possible for Sturgeon to present, at last, his own 'parallel argument'.

Sturgeon's Parallel Argument: Ethics and Science

With all this clarified, Sturgeon points out, we can see that there is something odd in Harman's worry. Moral realists of most stripes (including Sturgeon's) endorse the view

---

[226] I am simplifying somewhat from *Ibid.*, pp.246-249.

that moral properties supervene on non-moral properties. A supervenience relationship can be understood as follows: "A set of properties *A* supervenes upon another set *B* just in case no two things can differ with respect to their *A*-properties without also differing with respect to their *B*-properties. In slogan form, 'there cannot be an *A*-difference without a *B*-difference'."[228] So to say that moral properties supervene on non-moral properties is just to express the commonsensical view that, given any two actions performed in exactly similar circumstances and with exactly the same results, intentions, etc., either both are immoral or neither is immoral. But it seems clear from Sturgeon's analysis of Harman's claim that this entailment is just what Harman denies. Sturgeon feels that he can resist Harman's claim by showing in detail what a *truly* parallel case from science would look like:

> [C]onsider what follows from the parallel claim about microphysics, in particular about Harman's example in which a physicist concludes from his observation of a vapor trail in a cloud chamber, and from the microphysical theory he accepts, that a free proton has passed through the chamber. The parallel claim, notice, is *not* just that if the proton had not been there the physicist would have thought it was. This claim is implausible, for we may assume that the physicist's theory is generally correct, and it follows from that theory that if there hadn't been a proton there, then there wouldn't have been a vapor trail. But in a perfectly similar way it is implausible that if Hitler hadn't been morally depraved we would still have thought he was: for we may assume that our moral theory also is at least roughly correct, and it follows from the most central features of that theory that if Hitler hadn't been morally depraved, he wouldn't have done what he did. The *parallel* claim about the microphysical example is, instead, that if there hadn't been a proton there, but there *had* been a vapor trail, the physicist would still have concluded that a proton was present. More precisely, to maintain a perfect parallel with Harman's claim about the moral cases, the antecedent must specify that although no proton is present, absolutely *all* the nonmicrophysical facts that the physicist, in light of his theory, might take to be relevant to the question of whether or not a proton is present, are exactly as in the actual case. (These macrophysical facts, as I shall for convenience call them, surely include everything one would normally think of as an observable

---

[227] *Ibid.*, p.250

> fact.) Of course, we shall be unable to imagine this without imagining that the physicist's theory is pretty badly mistaken, but I believe we should grant that, *if* the physicist's theory were somehow this badly mistaken, but all the macrophysical facts (including all the observable facts) were held fixed, then the physicist, since he does accept that theory, would still draw all the same conclusions he actually does. That is, this conditional claim, like Harman's parallel claims about the moral cases, is true.
>
> But no skeptical conclusions follow; nor can Harman, since he does not intend to be a skeptic about physics, think that they do.[229]

Sturgeon also presents a different parallel argument, this time as a defense of his view that moral properties are natural but non-reducible.[230] He begins by considering what Harman's demand that moral properties be reducible to natural ones might amount to. One possibility is that Harman feels it would be necessary for moral properties to be reducible to the properties of physics. But in this case, Sturgeon has a simple parallel argument to use against him: there are at present no available reductions of biological, psychological, chemical, economic, etc. properties (nor even, as Quinn suggests, of a great number of everyday properties, such as the property of being a chair)[231] to the properties of physics, either; and it seems doubtful whether there ever *will* be such reductions. In particular, it seems no less doubtful than that there will ever be adequate reductions of *ethical* properties to the properties of physics. Therefore, if reducibility to physics is what Harman demands in order for us to have moral knowledge, then parallel

---

[228] McLaughlin and Bennett, 2005

[229] Sturgeon 1985, pp.251-252. I agree with the last statement about Harman (against whose arguments these points seem strong), but I think the question whether skeptical conclusions do in fact follow depends on a number of issues. Foremost among these is the issue of how exacting one's epistemic standards ought to be. A committed realist, of any stripe, seems compelled to mark out and defend the appropriate place for such a standard, and it would be very difficult to do this without controversy. For the practitioner of the parallel arguments approach as *I* have explained it, however, none of this will be necessary: the mere demonstration that both arguments are equally good is all that one needs, *so long as we conceive of philosophy as a practical discipline* (as I am, tentatively, in Chapters Two and Three).

[230] Sturgeon 1985, pp.240-241.

arguments will show that, by the same token, we cannot have knowledge of biological, psychological, chemical, etc. facts (or even knowledge of the fact that we are sitting in a chair). Following the standards of my parallel arguments approach, this would be an effective counter against Harman whether he accepts or rejects knowledge in all these cases.[232]

The other possibility is that he feels it is merely necessary to show that moral properties can be reduced to some combination of properties from a large grab-bag that includes all the properties of all of the natural and social sciences, plus a presumably vast number of other properties (like the property of being a chair). This gets around the difficulty faced above – it no longer matters whether *these* properties are reducible to physical ones. But the problem should be clear: it seems arbitrary, at this point, to say that all the properties in the grab bag are legitimate in their own right, while moral properties are not. What sort of justification can be given for this discrepancy? Perhaps one can be given: perhaps it is possible to find some fair reason to believe that all these properties have some important quality that moral properties lack. But it is not clear what this reason would be. Since the burden of proof lies with him to establish this if he thinks it can be established, his argument achieves nothing unless and until he is able to come up with such a reason. Sturgeon's parallel arguments provide a convenient stopping-point in our journey through the Harman/Sturgeon debate. However, we are only halfway through the forest. In the next section, I will consider the third and fourth papers in the exchange.

---

[231] Quinn 1986, pp.537-539

[232] See, again, the Introduction to this dissertation for an explanation of why this is so.

Harman and Sturgeon, Part II

Let us begin by reviewing the situation so far and clarifying what Harman needs to accomplish in his response to Sturgeon. Sturgeon capitalizes on the apparent strangeness of Harman's worry that *all* the non-moral properties of some action, etc. could be identical if the moral properties were different. Though he finds this suggestion very odd, Sturgeon offers to accept it on the condition that Harman hold microphysics to the same standard. Since we can only learn about subatomic particles like protons indirectly, by consulting non-subatomic instruments, it seems to follow from Harman's own reasoning that if all the non-subatomic facts were the same, but the subatomic facts were different, we would not notice the difference. Therefore, Harman must either maintain his view that there is a problem with ethics *and* add to it the view that there is a similar problem with microphysics, or abandon both views. Either way, the parallel argument strategy succeeds.[233]

Harman's challenge, then, is to show that these two cases are not parallel after all. Rather than do this, however, he ignores the challenge (apparently, as I will argue, because he misunderstands it) in his response to Sturgeon. Strangely, Harman's ducking of Sturgeon's central problem is apparently missed not only by all the commentators on the

[233] Actually, it *doesn't* succeed for *Sturgeon* either way: since he is arguing for moral realism, his presentation of the parallel arguments only works if he is able to use the threat of scientific skepticism, nihilism or relativism to convince Harman of the errors of his ways. Harman, however, can consistently decide not to worry about this threat; and if he bites the bullet in this way, Sturgeon's plan is unsuccessful. This illustrates a major advantage of what I have called the parallel arguments approach in the Introduction: following this approach, which does not aim at establishing any kind of realism, Harman's willingness to accept the existence of both 'problems' would count as a success, in a sense, for both Harman's side and mine. More on this later.

Harman/Sturgeon debate, but even by Sturgeon himself! This, also, will be shown in what follows.

The two main points in Harman's response[234] are a) an objection against Sturgeon's use of moral counterfactuals, in which Harman explores the possibility and consequences of 'moral epiphenomenalism', and b) a restatement of Harman's original problem with ethics, supplemented now by an objection against the practice of reflective equilibrium. I will deal with these two points separately, discussing Sturgeon's reply to each in turn.

Harman's Response, Part A: Moral Epiphenomenalism

Sturgeon, in his discussion of Woodworth, argues for his view that moral facts can be relevant to the explanation of moral knowledge as follows: first, he has the reader consider a story involving someone acquiring a moral belief in the normal manner; then, he has the reader consider a second story that is exactly the same as the first, except that the moral facts are different. If, for any such story, the same moral belief would not have arisen in the same way had the moral facts been different, then it follows (Sturgeon claims) that the presence of those moral facts is relevant to the explanation of the moral belief; and since there are such stories, Harman is mistaken in claiming that moral facts are *never* relevant to moral beliefs.

Sadly, the most important point here seems invariably to be overlooked; but (thanks to the guidance of the North Star) we can avoid this error now. Sturgeon, it will be

remembered, points out that Harman could mean two different things in saying that moral facts are irrelevant to moral knowledge: there is a consistent and an inconsistent interpretation of him.[235] Harman could be interpreted as saying that it would make no difference if the children had been doing something *else* that was not morally wrong, or he could be interpreted as saying that it would make no difference if they had been doing nothing wrong even though they were doing *exactly the same thing*. And as we have seen, Sturgeon points out[236] that the first of these counterfactuals cannot be what Harman means, so he focuses on the second (after having discussed, and dismissed, the first). This is why Sturgeon expresses some puzzlement regarding what we are to do in imagining that the same non-moral facts apply without the same moral facts applying, since he takes it that moral facts supervene on non-moral ones. It is therefore against this second sort of moral counterfactual, again, that a critic of Sturgeon must focus his or her attention. Let us call the problem of the second sort of moral counterfactual, in which all the non-moral facts are identical and only the moral facts are different, 'the Actual Problem', and the problem of the first sort of moral counterfactual, in which some non-moral facts are different, 'the Erroneous Problem'.

Harman attempts to deal with Sturgeon's discussion of moral counterfactuals by reference to a new story, which I will call 'Cat Abuse II' so as to keep it separate from the prequel, 'Cat Abuse' ('Cat Abuse', again, was the story of the hoodlums pouring gasoline on a cat and burning it). In Cat Abuse II, the main character (Albert) beats a cat with a stick while a secondary character (Jane) looks on with moral abhorrence.

---

[234] Harman 1986
[235] Sturgeon 1985, p.246

Referring to this new story, Harman attempts to articulate what Sturgeon's test for explanatory relevance must amount to: "To supply [apply?] Sturgeon's test, we must therefore ask whether Jane would have believed Albert's act was wrong if Albert had not done something wrong."[237]

Harman seems to agree that any moral fact that could pass this test is explanatorily relevant; but he denies that every fact that passes the test is relevant to the right *kind* of explanation. He argues for this point by invoking the possibility of what he calls 'moral epiphenomenalism':

> A good and conclusive way to see that this sort of explanation is irrelevant to full empirical testing is to notice that a moral epiphenomenalist can accept the relevant counterfactual judgment without having to suppose that moral features of actions ever explain any nonmoral facts. A moral epiphenomenalist takes moral properties to be epiphenomenally supervenient on natural properties in the sense that the possession of moral properties is explained by the possession of natural properties and nothing is influenced or explained by possession of moral properties.[238]

What Harman's 'conclusive' argument amounts to is this: Sturgeon tries to show, in his discussion of the Erroneous Problem – but not in his discusion of the Actual Problem! – that in a counterfactual story about Woodworth (say) involving different moral and nonmoral facts, the moral assessment would be different. But Harman, seemingly forgetting Sturgeon's later point that all this pertains only to the Erroneous Problem, objects that this test does not suffice, because it can yield a false positive: according to

---

[236] *Ibid.*, p.250

[237] Harman 1986, p.63. Sadly, this articulation maintains precisely the double meaning that Sturgeon was so careful to disambiguate; but it should now be clear that, to read Harman consistently, the question can only be whether Jane would have believed Albert's act was wrong *if abusing the cat had not been wrong despite its having been done under the same circumstances, for the same reasons, by the same person, and with the same feelings and intentions in mind as in the story.*

the test, moral facts would be explanatorily relevant even if moral epiphenomenalism were true. And then – here the premise is merely implied, but all-important – *if* moral epiphenomenalism were true, then moral facts would never explain moral knowledge.

In what follows, I will try to show two things: first, that the moral epiphenomenalism objection *cannot* be applied to the Actual Problem, and hence is irrelevant; and second, that even as an attack against Sturgeon's discussion of the Erroneous Problem, it seems to fail. Let us begin by seeing more clearly what assumptions lie beneath Harman's moral epiphenomenalism by spelling out in more detail what it would mean for that position to be true. We are to imagine that, in any supervenience relationship between a moral and a non-moral property, the moral property is the cause – and never the effect – of the non-moral property. We are also to imagine that the moral property has no influence on anything else.[239] Since moral epiphenomenalism, as Harman himself says in his definition above, entails the supervenience of the moral on the non-moral, it follows that there can be no difference in moral properties without a difference in non-moral properties. Putting all this together, and substituting in the details of the case we have been considering as an illustration, we get the following picture: Woodworth's cowardice is the cause of two things. It is the cause, first, of the events that lie in the unbroken explanatory chain extending from the cowardice to DeVoto's moral assessment of Woodworth; and second, of the immorality of Woodworth, which follows (due to the supervenience relationship) from Woodworth's cowardice.

---

[238] *Ibid.*

With this picture in mind, it should be clear that the first point I wish to make must follow: the Actual Problem involves consideration of counterfactual scenarios in which all the non-moral properties are the same, but the moral properties are different. But the supervenience relationship specified by Harman makes this an impossibility. Hence, the moral epiphenomenalism objection fails to address Sturgeon's discussion of the Actual Problem.

Though not quite so obviously (or relevantly), it also seems to follow that the moral epiphenomenalism objection fails even to address Sturgeon's discussion of the Erroneous Problem. Why does Harman think that this objection is effective in this connection? It seems that he is relying on the view that explanatory chains extending from moral facts to moral knowledge must be one-way causal ones, with each point on the chain serving as a cause of the next point in the series. We can chart out the two possibilities Harman has in mind in the following way (with the arrows representing causation):

The non-epiphenomenal option:

(NE1) Moral fact ⟶ (NE2) Nonmoral fact ⟶ (NE3) Observation

The epiphenomenal option:

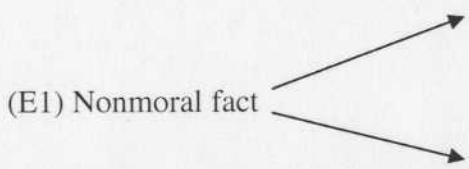

[239] I leave out Harman's other condition that the moral property cannot *explain* anything, since in the context of my analysis, that – as will soon become clear – would be question-begging.

(E2b) Observation

Harman assumes that NE1 can help explain NE3 because it helps (indirectly) to cause NE3, but that E2a cannot help explain E2b because E2a is causally inert. But is this a reasonable thing to assume? Perhaps, but it will have to be *shown*, since there seem to be many plausible counterexamples to it. For example, suppose that a roommate asks me to bring him some milk out of the refrigerator so that he can pour it on his cereal. I take the milk out of the refrigerator, smell it suspiciously (we can imagine that the expiration date is too blurry to read), and say "I wouldn't use that, if I were you. It's gone bad: it'll taste awful and make you sick". Now, surely, I would be justified in thinking these things if the milk really smelled particularly bad. But why? Nobody has tasted the milk, and it hasn't made anyone sick. There is a single cause – the status of the milk – of three effects – its bad smell, the bad taste, and the health-threatining properties. Let us call the single cause M2, and the three effects M3a, M3b, and M3c. Of these three effects, only M3a lies on an unbroken causal chain leading to my knowledge of M3b and M3c. But there seems to be no problem in my saying, in spite of this, that part of the explanation of my belief that the milk is now unhealthy to drink is that it really *is* unhealthy to drink Otherwise, we would have to allow the existence of a much more general epistemological problem, and the parallel arguments approach would be successful here also. Of course, this is not to say that moral epiphenomenalism is true – we are given no reason to think that it is – but only that, even if it *were* true (or even *could be* true), that would not present any problems for Sturgeon, *even if* his point had to do with the Erroneous Problem, which it doesn't.

Sturgeon's Response to Part A

Sturgeon, in his response to what he calls Harman's "brief argument" from moral epiphenomenalism, strangely – and unfortunately – stops navigating by starlight and instead follows the map provided by Harman. Sturgeon concedes Harman's point that

> [t]he moral epiphenomenalist who accepts the counterfactual about Hitler will understand it as follows. Hitler had certain *nonmoral* features that did two things: they *made it true* that he had another property, depravity, and they also *caused* him to behave in certain ways; but unlike these nonmoral properties, **the depravity that supervened on them was causally inert, *so* it is a mistake to cite it in explanation of his behavior.**[240]

Since he overlooks the fact that Harman's attack deals with the Erroneous Problem instead of the Actual Problem,[241] Sturgeon responds to the moral epiphenomenalism objection only by trying to show that it entails an implausible view.[242] This plan of attack seems less likely to work than those more straightforward ones I have suggested, however, because the method Sturgeon employs to demonstrate the unlikelihood of epiphenomenalism is nothing more than a second recitation of the familiar cases in which people tend to attribute causal efficacy to moral properties.[243] Later on, however, he shifts to a more interesting line of attack: he points out that, in morality (as in biology, chemistry, psychology, etc.), the best account for a supervenience relation might not be epiphenomenalism but rather "a kind of causal constitution" of the supervening facts out

---

[240] Sturgeon 1986a, p.74. Emphasis mine.
[241] See, again, Sturgeon 1985 p.250.
[242] Sturgeon 1986a, pp.74-75
[243] *Ibid.*, p.75

of the more basic ones".[244] One could, of course, reject this possibility, holding instead (as Kim does[245]) that all supervenient facts, whether moral or otherwise, are epiphenomenal facts. But then, as Sturgeon points out, there is no distinct 'problem for ethics'[246] – and again the parallel arguments approach – in my sense, at least – succeeds.

To conclude: Sturgeon's treatment of the moral epiphenomenalism challenge may well work, but it does not seem as straightforward as a move that would have been available to him earlier on had he recognized that the assumptions behind Harman's argument are problematic.[247]

Harman's Response, Part B: Traditional Worries about Verification

The more promising of Harman's responses to Sturgeon begins with a reference to 'reflective equilibrium'.[248] One begins to move toward such a state by beginning to inquire into some area of knowledge (or, in the case of 'wide reflective equilibrium', into

---

[244] *Ibid.* This point is made most clearly in a footnote Sturgeon appends to his original paper: "High unemployment causes widespread hardship, and can also bring down the rate of inflation. The masses and velocities of two colliding billiard balls causally influence the subsequent trajectories of the two balls. There is no doubt some sense in which these facts are causally efficacious *in virtue of* the way they supervene on – that is, are constituted of, or causally realized by – more basic facts, but this hardly shows them *in*efficacious. (Nor does Harman appear to think it does; for his *favoured* explanation of your moral belief about the burning cat, recall, appeals to psychological facts (about your moral sensibility), a biological fact (that it's a cat), and macrophysical facts (that it's on fire) – supervenient facts all, on his physicalist view and mine)" (Sturgeon 1985, p.250)

[245] Kim, 1984

[246] Sturgeon 1986a, p.75

[247] Sadly, this point seems not only to have been missed by Harman and forgotten by Sturgeon, but also to have been overlooked by many – and as far as I know, perhaps all – other commentators on the Harman/Sturgeon debate. In his well-read book *An Introduction to Contemporary Metaethics*, for instance, Alexander Miller deals with the Harman/Sturgeon debate as though the sole point of the last half of the exchange pertained to this epiphenomenalism argument; but rather than emphasize that the argument attacks the wrong part of Sturgeon's 'Moral Explanations', Miller takes the relevance of the argument as given and spends 27 pages (pp.150-177) considering the merits of different forms of explanations insofar as they could strengthen Sturgeon's second response to Harman. Quite possibly, it is treatments like Miller's that preserve the impression among metaethicists that Harman has been somewhat successful in making his case for relativism. See Miller, 2003.

all areas of knowledge)[249] with one's own particular intuitions, beliefs, theories, and even inference patterns, but leaving one's mind open to the possibility that one is mistaken in any of these things. One comes, through life experience and the criticism of others, to build on and add to these components of thought; and in that process, one comes to notice internal tensions and inconsistencies in one's thinking. At such times, one needs to resolve the tension in order to return to consistency, but this can be done in a number of ways: in ethics, for instance, one can choose between abandoning a theory to preserve a particular intuition, or in overriding a particular intuition for the sake of maintaining a general theory.

Harman argues that a problem with ethics lies in its employment of reflective equilibrium:

> A moral philosopher who tries to find general principles to account for judgments about particular cases, hoping eventually to get principles and cases into "reflective equilibrium," is studying commonsense ethics. The results of such a study are results in moral psychology. Just as a study of commonsense physics is not a study in physics, the study of commonsense ethics is not an investigation of right and wrong, *unless* what it is for something to be right and wrong can be identified with facts about moral psychology.
>
> Two physicists with different theories and correspondingly different physical explanations and sensibilities will get different results from their thought experiments. It is not evidence for theory *A* that physicist *A*'s thought experiments accorded with that theory. Theorist *A* could not plausibly say, "The evidence for theory *A* is that, when I imagine particles interacting, I imagine them interacting in accordance with the principles of theory *A*." That is not evidence for *A*'s theory, because it is explained by *A*'s holding the theory and having the corresponding physcal sensibility, quite apart from whether or not theory *A* is true. To test theory *A* against the world, the theorist needs to find some result in the world that is most plausibly explained by that theory's being true.

---

[248] The term is from Rawls 1971, p.48
[249] See Hales, pp.39-45

> Such tests are possible in physics, because physical facts can have observable effects on people who do not believe in those facts ahead of time. The issue is whether such tests are possible in ethics.[250]

It is, to begin with, difficult to imagine how someone as familiar with moral discourse as Harman surely is can believe that the advocates of preference utilitarianism, say, or capital punishment have nothing to offer in defense of their views aside from claims like 'When I imagine actions being good, I imagine them being good along the lines of preference utilitarianism' or 'When I think of capital punishment, I think of it as being morally permissible'. As Harman points out, these kinds of exclamations do not count as evidence or arguments for the views they support; they are nothing other than *assertions* of those views. And if Harman is to be taken at face value here, then his position is the rather implausible one that there never are, or could be, any arguments in normative or applied ethics; specialists in those areas, on his view, simply show up to the debate and assert whatever they feel like. But even a casual survey of the literature in normative and applied ethics seems to show a very different picture indeed. So this charge against the reliance on reflective equilibrium and thought experiments in ethics, as stated at least, is difficult to understand.

Nor are such thought experiments even absent from science. Consider, as just one example, how Galileo argued that heavier bodies do not fall more quickly than lighter ones through a *reductio ad absurdum*:

---

[250] Harman 1986, p.61

> Salviati: If then we take two bodies whose natural speeds are different, it is clear that on uniting the two, the more rapid one will be partly retarded by the slower, and the slower will be somewhat hastened by the swifter. Do you not agree with me in this opinion?
>
> Simplicio: You are unquestionably right.
>
> Salviati: But if this is true, and if a large stone moves with a speed of, say, eight while a smaller moves with a speed of four, then when they are united, the system will move with a speed less than eight; but the two stones when tied together make a stone larger than that which before moved with a speed of eight. Hence the heavier body moves with less speed than the lighter; an effect which is contrary to your supposition. Thus you see how, from your assumption that the heavier body moves more rapidly than the lighter one, I infer that the heavier body moves more slowly.[251]

However problematic Harman's statement of this problem may be, it would perhaps be unfair to deny that there is a genuine worry along something like these lines. After all, while there are *some* thought experiments in science, there is also some reference to the mind-independent world. In ethics, however, this is not as clear – we must, if we are honest, sincerely attempt to compare ourselves with Harman's crude sketch of physicist *A* to see how, if at all, we are really so different. And, since we have now come to the end of Harman's moves in the Harman/Sturgeon debate, with all the others having been found to be irrelevant or otherwise unsuccessful, I wish to consider this final problem of Harman's in the best possible light. Let us imagine it, then, free of the problems in Harman's expression of it (as described above); and let us not quibble with the fact that Harman does not really present an *argument* so much as a *challenging question*.[252] The crux of Harman's contribution to the Harman/Sturgeon debate is, on this most charitable

---

[251] Galileo 1991, p.63

[252] Harman's own thought experiments, like those of physicist *A* in his caricature, *assume* rather than argue for their conclusion: to construe Cat Abuse and Cat Abuse II as presenting an *argument* for the uninformative nature of moral facts would be to interpret them, uncharitably, as begging the question.

reading, simply the demand for moral realists to justify their view that we can have moral knowledge. This is by no means an original demand[253], but it is still an interesting one. Let us therefore consider how Sturgeon responds to it, and how others have done so differently.

Sturgeon's Response, Part B

Like most moral realists, Sturgeon believes that there are adequate, non-subjective ways of genuinely expanding one's knowledge of non-relativistic moral facts. Two ways in which we might accomplish this augmentation of our knowledge are, first, by means of a prior trust in some particular moral theory; and, second, by means of a prior trust in the verdicts of some other moral judge.[254] Do these suggestions, given adequate fleshing-out, sufficiently respond to Harman's demand? Perhaps – but to err on the side of caution, let us assume that they are not. What then? In what follows, I will discuss the two-part worry that remains: a) the Mackie/Harman worry that there is no acceptable connection between the moral and non-moral facts of some action, and b) the further worry, seemingly based on a), that progress in normative moral theory is an illusion because we can only ever behave like physicist *A*.

Beyond the Harman/Sturgeon Debate: A Different Sort of Parallel Argument

Up to this point, I have considered only one sort of parallel for ethics: the parallel between ethics and science. It may well be that Sturgeon's set of parallel arguments

---

[253] As Sayre-McCord points out in a discussion of Mackie, these worries go back at least to Moore, and probably to Hume. (Sayre-McCord 1988, p.13)

along these lines, supplemented as needed with other parallels with science (some of which I have provided above), adequately defend ethics against all Harman's attacks considered here. For the sake of argument, however, I will now assume that these parallels break down in the face of Harman's question about the possibility of moral knowledge. Perhaps, indeed, these parallels *must* break down if they are meant to provide a scientific parallel for *every* aspect of ethics. This would, in fact, not be very surprising: there are some plainly important differences between moral and scientific claims. Most importantly, perhaps, moral claims tell us what we *should* do, whereas scientific claims do not. Is there a way to overcome this imperfection in the analogy?

It seems that there is. A closer parallel can be drawn between first-order claims in ethics and first-order claims in *epistemology*: that is, between claims about what we should do and claims about what we should believe. It is important, before proceeding, to make clear the distinction between what I mean by first-order epistemology and the more theoretical second-order epistemology which is most frequently discussed under the name 'epistemology' in contemporary philosophical discourse. The distinction is made clear by Jeremy Randel Koons:

> The term 'epistemology' refers to two rather different practices. Construed narrowly, it is the attempt to construct a set of formal conditions for the justification of beliefs (e.g., explanatory coherence among beliefs; proper inferential relation to foundational beliefs; proper causal genesis, a la reliabilism, etc.). Construed more broadly, it is the practice, common in every field from chemistry to philosophy, of evaluating beliefs as justified or unjustified, scientific methods as rational or irrational, etc... Epistemological discourse is the very thing philosophers spend their time employing; it is their stock in trade. After all, what philosophers do is advance and evaluate arguments; philosophers try to tell us what it is rational to believe about ethics, about the meaning

---

[254] Sturgeon 1986a, pp.72-73

> and reference of words, about mental states. Without epistemology, philosophy as we know it grinds to a halt. Science does the same: how can science succeed if there is no distinction between science and pseudoscience, rational and irrational research methods, justified and unjustified scientific beliefs, and so on?[255]

First-order ethical and epistemological (in the broader sense above) statements, properties and knowledge are normative: they seek to tell us what we *ought* to think or do, which forms of reasoning or actions are *good*, etc. And as Koons points out, we can disqualify epistemological properties, facts and discourse only at the expense of eliminating the possibility of philosophy and science. Therefore, the ethics-epistemology parallel seems ideal for the purposes of the parallel arguments approach in this context.

With this in mind, let us turn to the two problems that still linger, in some sense, from the Harman/Sturgeon debate. The first, it will be recalled, was the problem of the Harman-Mackie gap: what can connect the moral with the non-moral properties of some action?

Let us look, again, at Mackie's clear statement of this problem:

> [One] way of bringing out this queerness is to ask, about anything that is supposed to have some objective moral quality, how this is linked with its natural features. What is the connection between the natural fact that an action is a piece of deliberate cruelty – say, causing pain just for fun – and the moral fact that it is wrong? It cannot be an entailment, a logical or semantic necessity. Yet it is not merely that the two features occur together. The wrongness must somehow be 'consequential' or 'supervenient'; it is wrong because it is a piece of deliberate cruelty. But just what *in the world* is signified by this 'because'? And how do we know the relation that it signifies, if this is something more than such actions being socially condemned, and condemned by us too, perhaps through our having absorbed attitudes from our social environment?[256]

---

[255] Koons 2000, pp.246-248
[256] Mackie 1977, p.41

And now, let us consider a suitably parallel statement of a possible problem from epistemology. The idea now is that there is something 'queer' about the relationship between the soundness of an argument, say, and the non-epistemological features on which this soundness is meant to supervene:

> [One] way of bringing out this queerness is to ask, about any argument that is supposed to be good, how this is linked with its non-epistemological features. What is the connection between the logical fact that an argument's formal structure is logically unsound – say, an instance of denying the consequent – and the epistemological fact that it is a bad argument? It cannot be an entailment, a logical or semantic necessity. Yet it is not merely that the two features occur together. The badness must somehow be 'consequential' or 'supervenient'; it is bad because it is a fallacious argument. But just what *in the world* is signified by this 'because'? And how do we know the relation that it signifies, if this is something more than such argument forms being socially condemned, and condemned by us too, perhaps through our having absorbed attitudes from our social environment?

It may seem strange to distinguish the logical and epistemological elements of an argument: how could anyone see that a chain of reasoning is illogical without seeing that he or she ought not be persuaded by it? But by the same token, it seems strange that anyone could see that some action is a piece of deliberate cruelty (and even, for anything Mackie says, deliberate cruelty *just for the fun of it*) without seeing that this gives him or her a reason not to perform it. Further, there really are some people (including some I have actually met!) who admit that particular sets of beliefs are logically inconsistent but nonetheless believe in them, and who claim not to feel that their logical inconsistency provides any reason not to do so. Finally, Harman himself, in his non-ethical writings, explicitly distinguishes logical from epistemological facts:

> It is important to notice that deduction and induction are not two kinds of reasoning. In fact, induction and deduction are not two kinds of anything. Deduction is concerned with certain relations among propositions, especially relations of implication and consistency. Induction is not concerned with those or any similar sort of relation among propositions. Induction is a kind of reasoning. But, as we will see, deduction is not a kind of reasoning… Logic, the theory of deduction, is not by itself a theory of reasoning. In other words, *it is not by itself a theory about what to believe (or intend); it is not a theory concerning how to change your view.*[257]

Since logical and epistemological facts are not related analytically on Harman's account, the gap between the two must be bridged in some other way. But is this bridge easier to build than one across the Harman-Sturgeon gap? This is what must be shown by those on the side of Harman and Mackie if they wish to overcome the parallel arguments response. But, it seems, this has not been shown; and it is far from clear how it *could* be shown. Until it is, this parallel argument seems to defend ethics from Harman and Mackie's problem at least as well as Sturgeon's scientific parallel does.

Moreover, this new parallel allows us to deal more clearly with the question of how moral facts help to explain moral beliefs, as many philosophers have noted.[258] Harman takes it for granted that moral facts cannot explain anything of which they are not the cause, and therefore (since he also holds that non-reducible moral facts cannot be the cause of moral beliefs) that they cannot explain those beliefs.[259] Let us grant him provisionally, in the spirit of charity, the premise that non-reducible moral facts cannot

---

[257] Harman 1999, pp.27-28. Italics mine.
[258] e.g. Sayre-McCord 1988a, Koons 2000, Pust 2000, Shafer-Landau 2003b, and Nobis 2005.
[259] That Harman takes this for granted is clear, *inter alia*, from his claim that he feels that if moral properties were causally inert (as in his moral epiphenomenalism objection), it follows *without further argument* that they could not explain our beliefs. It is also evident in many of his statements of the problem (e.g. "Could the wrongness of causing needless pain to animals have an observable *effect* on the world that could serve as evidence that this is wrong…?" (Harman 1986, p.60: emphasis mine)

help to cause moral beliefs. The success of his argument then rests on the following line of reasoning: irreducible normative facts, including irreducible moral facts, are causally inert; causally inert moral facts cannot help to explain our beliefs (or anything else); properties that only figure in facts that cannot help to explain anything should not be allowed into our ontology; therefore, moral properties, since they figure only in moral facts, should not be allowed into our ontology.

Is this line of reasoning successful? As Shafer Landau points out quite concisely, there is something peculiar about it, since epistemological facts are also normative:

> If we are unprepared to admit the existence of normative facts [into our ontology], then the claims of those who deny them cannot be *justified, warranted, credible,* or *well-supported.* Yet if we let normative facts back in, then we have a dilemma. Either we insist on the causal inefficacy of normative facts, in which case, having allowed them entry into the ontology, we must give up the strong causal test [i.e. the principle that only those properties or entities with causal powers are permitted to have explanatory powers or a place in our ontology]. Or we can take their admission into the ontology, coupled with the causal test, as an indicator that normative facts can, after all, possess the relevant sort of causal power. But if they do possess such causal power, then causal considerations no longer pose any threat to their existence.[260]

In this way, the parallel arguments approach seems to deal with this argument quite handily.

These last two parallel arguments help to make clear how to deal with the last remaining complaint of Harman's from the last section – the complaint that any attempt to gain knowledge about non-relativistic, non-reductive ethical properties must rely on reflective equilibrium. The parallel here, of course, is that (first-order) epistemology – not to

mention second-order epistemology, metaphysics, and perhaps everything else in philosophy – is equally dependent on some form of reflective equilibrium.[261] Hence, as in Shafer-Landau's dilemma quoted above, Harman must choose once again between accepting the *prima facie* legitimacy of non-relativistic ethical practice, abandoning the legitimacy of his own grounds for rejecting it, or being caught in an inconsistency.[262] Hence, this complaint, also, is dealt with successfully by the parallel arguments approach.

A Final Objection

In this chapter, I have endeavoured to represent the strongest of Harman's arguments in the very best light, and to ensure that each is dealt with adequately. In the process, I have at times had to extrapolate arguments from small suggestions of Harman's, and to present them in their most powerful form. So as not to mar this treatment of the Harman/Sturgeon debate by an omission at the last moment, I anticipate one last possible objection from Harman. The germ of it is to be found in the last two paragraphs of his response to Sturgeon:

> …[T]he options seem to be these: (A) The option of "poisoning the well": deny that scientific claims can be tested in the relevant sense. (B) The math option: claim that there is an analogy between ethics and mathematics, deny that mathematics can be tested in the relevant sense, and deny that this is a problem for mathematics. (C) Noncognitivism. (D) Naturalistic reductionism.

---

[260] Shafer Landau 2003b, pp.113-114

[261] See e.g. Pust 2000 and Hales 2006 for more on this point.

[262] The only other possibility is, of course, for Harman to present a good argument that establishes a relevant distinction between moral and epistemological normativity and makes clear why we should grant special privileges to epistemological normativity. But it is not clear what such a *good* argument would consist in; and the burden of proof must, surely, lie with Harman to present such an argument, if there is one.

> As I have said, I favor (D) with a touch of (C). But my aim in this paper has not been to argue for that. Rather, I have argued that there is a real issue here and one or another of these options has to be taken seriously.[263]

The point of interest lies in the first two lines of this quote. Remember that Sturgeon, as we have seen, points out in his discussion of the Actual Problem that by Harman's standard, even *scientific* facts (such as microphysical ones) will seem irrelevant to explaining our beliefs. Sturgeon takes this to be a powerful response to Harman's concern, since he can see that Harman is unwilling to bite the bullet and accept skepticism about scientific facts. But this sort of argument – which is typical of the sort used in the parallel arguments approach – is just what Harman seems to describe as fallacious[264]. Is he right in accusing Sturgeon of committing a fallacy?

Let us consider this charge in its strongest form. To begin with, let us consider what the fallacy of *tu quoque* consists in. One commits this fallacy when one deals with a criticism not by showing that the criticism does not apply, but by showing instead that the person who offered the criticism is guilty of the same fault. For instance, if Eloise accuses Frida of not being careful to check for passing bicyclists before opening her car door, Frida would commit the fallacy if she responded by pointing out that Eloise is just as bad *and* intended this as a complete response to Eloise's accusation (this last part is important). This is a fallacy because the point that Frida is making – that Eloise is a hypocrite – may well be true, but that has no bearing on whether what Frida is doing is wrong.

---

[263] Harman 1986, p.67

[264] Though Harman calls it an instance of "poisoning the well", it seems from his discussion that he has this fallacy confused with the very different fallacy of *tu quoque*.

Now, is this the sort of thing that Sturgeon and Shafer-Landau are doing? There seems to be an important difference. Frida, in my stereotypical *tu quoque* story, seems to recognize that what she is doing is wrong. She is merely using the *tu quoque* tactic to divert some of the blame from herself, or to lash out at Eloise out of embarrassment. Sturgeon and Shafer-Landau, by contrast, appear to be arguing sincerely for the view that there is *nothing* wrong with allowing non-reductive moral facts into our ontology: they seem to be using the parallels from science and epistemology not to deflect charges they know are strong, but rather to argue that Harman is mistaken about which standards are correct. So the situation is really similar to one in which Jane, in Cat Abuse II, attempts to persuade Albert that he ought not beat someone else's cat with the stick by pointing out how upset he became a week before when someone abused his dog. While Jane is pointing out to Albert that he is acting hypocritically, this does not seem fallacious: it seems, instead, to be an entirely legitimate and rhetorically effective technique to help Albert see that there is something wrong with what he is doing.

But even if the parallel arguments used against Harman in this chapter are innocent of that fallacy, their use may still be, in some way, unsatisfactory. Sturgeon argues that Harman is not consistent in arguing for skepticism about non-relativistic moral properties while denying skepticism about non-relativistic microphysical properties, while Shafer-Landau takes this one step further and argues that Harman could *never* be consistent in rejecting the legitimacy of normative properties: to reject their legitimacy is already to employ them. But even though one cannot consistently say or hold that normative

properties are problematic, it still may be the case that they are not; and the parallel arguments strategies employed by Sturgeon and Shafer-Landau have not dealt with this possibility.

But here it should be remembered that the parallel arguments approach is not *meant* to deal with such 'profound' possibilities: they will be dealt with in the Conclusion. It is enough for the purposes of this chapter if I have shown that all the arguments Harman makes on these points can be answered in a way that would satisfy a 'practical' philosopher.[265] And at this point, that seems to have been accomplished.

[265] See my Introduction for more on the distinction between 'profound' and 'practical' philosophy, and my method of dealing with it in this dissertation.

Third Chapter: Moral Disagreement and Moral Reasons

Introduction

Up to this point, I have followed Harman's practice of discussing two types of relativism – appraiser relativism and agent relativism – interchangeably.[266] However, there are significant distinctions between these two forms of moral relativism, and it is important to keep them in mind.

Appraiser relativism is the view that the moral status of an action or inaction is relative to the moral views of the appraiser of that action or inaction. This is the sort of relativism Harman discusses when he draws an analogy between the variability of the mass of objects relative to different inertial frames and the variability of the moral status of acts relative to different moral views. To an appraiser relativist, there is no objectively correct answer to the question "Should George W. Bush have bombed Afghanistan?" According to some moral views, the answer is yes; according to others, the answer is no; and no moral views (and hence no answers) can be better or worse, objectively, than any others.

Agent relativism, by contrast, is the view that the moral status of an action is relative to certain features of the particular agent performing (or neglecting to perform) that action. These features vary depending on the version of agent relativism in question, but they tend in all accounts to be psychological features such as desires, interests, concerns, etc.

---

[266] At one point in Harman and Thomson 1996, Harman does briefly mention the existence of such a distinction. However, on the first new sentence on the following page, he lapses back into using the

To illustrate this with one of Harman's examples, we can imagine (again) that Hitler cannot be made to feel remorse about killing the innocent people he killed, but that Stalin can (or does) feel remorse about his own killings of innocents. An agent relativist could hold that this discrepancy is sufficient to make it the case that, even though the actions performed by Hitler and Stalin may have no morally relevant differences in themselves, the differences between the agents make it the case that Hitler did not act wrongly, and that Stalin did, in performing those actions.

Not only are these views distinct, but neither one entails the other. One can consistently hold that appraiser relativism is true and that agent relativism is false. For instance, one can imagine a world in which everyone is either a specific sort of ethical egoist (to be explained more fully below), a preference utilitarian, or a divine command theorist who accepts the teachings of some particular branch of Pentecostalism to be correct. Let us assume as true, for the sake of argument, the not unlikely claim that these three moral theories are incompatible with one another. Let us also assume, for the sake of argument, that appraiser relativism is true. Now, we can imagine that opportunities for a perfect crime exist from time to time in this world. Among these opportunities are those in which an old-age pensioner can be defrauded of all his or her savings. Some agents who have the opportunity to commit these crimes can be made to care about their victims, and some cannot. Regardless of these differences between agents, however, the Pentecostals and utilitarians all hold that none of the agents would be acting morally if they were to commit the crimes. And regardless of the differences between the agents, the egoists all

---

general term 'moral relativist' without mentioning which form of moral relativism is indicated. He does not discuss that distinction again. See pp.62-63 of that work.

hold that the agents are morally obliged to commit the crimes (since the particular brand of egoism they espouse entails that anyone who feels moral concern for a helpless victim is morally weak and needs to overcome that weakness by committing the crime in question). In this scenario, then, everyone's moral view is correct (since appraiser relativism is true) and yet none of these moral views entail that agent relativism is correct (since there is at least one act whose moral status is independent of any psychological facts about the agent).

One can also consistently hold that appraiser relativism is false and that agent relativism is true. One might, for instance, believe that the ultimate test of whether an action would be ethical is whether or not one could be made to care about the consequences, and one might simultaneously believe that anyone who thinks otherwise is objectively mistaken.

Appraiser relativism and agent relativism are motivated by different concerns, and hence different arguments tend to be advanced for each. Appraiser relativists tend to be puzzled by the presence of seemingly intractable moral disagreement. If there are objective moral facts, they wonder, how can moral disagreements persist among rational appraisers who seem open to entertaining one another's arguments? Appraiser relativists are typically motivated by such puzzlement to abandon the view that there are objective moral facts. When challenged by those who believe in objective morality, therefore, they tend to support their form of relativism by arguing that it is the best explanation for persistent disagreement among real or idealized disputants.

Agent relativists, by contrast, tend to be puzzled by the claim that one can be objectively obliged to do something that one could not possibly be motivated to do for the sake of those one could not possibly be made to care about. They therefore deny that there can be any such obligations. When pressed, they tend to argue for their position by pointing to apparent difficulties in the view that one can have a moral reason to do something that one cannot be made to care about doing. Hence, while appraiser relativists are fond of the various versions of the argument from moral disagreement, agent relativists are fond of the various versions of what I will call the argument from moral reasons.

It should be noted that neither of these two arguments is effective at establishing the other version of relativism: the argument from moral reasons is not a good argument for appraiser relativism, and the argument from moral disagreement is not a good argument for agent relativism. Consider, first, the argument from moral reasons. If that argument were effective, it would establish that (for instance) Stalin acted wrongly, and that Hitler did not, on Harman's portrayal of them; but what seems to follow from this is merely the *objective* moral fact that one is always justified in doing what one had no psychologically-driven reason not to do. Far from supporting appraiser relativism, this seems to undermine it.

Next, consider the argument from moral disagreement. Even if that argument were made persuasively, such that it followed from it that disputes over such issues as abortion, vegetarianism, and religious toleration are unresolvable even in principle, that conclusion would still be irrelevant to the issue of whether one can be morally obliged to do

something one can't be made to care about. So the argument from moral disagreement is not effective support for agent relativism.

In this chapter, I will critically examine the last two of Harman's arguments for relativism: the argument from moral disagreement and the argument from reasons. For reasons that should now be clear, I will disambiguate Harman's moral relativism into its appraiser and agent strands, and consider his final arguments in the best possible light by applying each to the form of relativism it supports most effectively.

I. Appraiser Relativism and the Argument from Moral Disagreement

I will begin this section by presenting what I take to be the most important problems with appraiser relativism. This will show the hurdles that any argument will need to clear in order to establish this form of relativism. Next, I will present Harman's argument from moral disagreement. Following that, I will present some objections against Harman's argument along the lines of the parallel arguments strategy.

Problems with Appraiser Relativism

Robert Streiffer, in his excellent but neglected work, *Moral Relativism and Reasons for Action*, has compiled (and in some cases strengthened) a number of the traditional objections against appraiser relativism.[267] Here are some of the best:

1) Suppose that Smith claims that stealing money is immoral, and that Jones replies, "That's true: stealing money is immoral". Normal speakers of English would

assume that Jones' claim that stealing is immoral is just a spelling-out of his earlier claim, 'That's true'. However, if appraiser relativism is correct, then this assumption is mistaken. According to appraiser relativism, Jones must mean, in saying 'stealing is immoral', that stealing is contrary to Jones' moral view; but that, in saying 'that's true', he is endorsing the truth of what Smith said (namely, that stealing is immoral according to *Smith*'s moral view).

2) Further, normal speakers of English tend to suppose that Smith and Jones, in the above scenario, are both stating the proposition that it is wrong to steal money. But according to appraiser relativism, this is false. If both Smith and Jones were stating that proposition, then they would both be stating the same thing, which appraiser relativism denies (since it entails that Smith says stealing money is wrong *relative to Smith*, etc.). Furthermore, there seems to be no reason to think that Smith is stating the proposition while Jones is not, or vice versa. Hence, the natural position for an appraiser relativist to take is that Smith and Jones are both alike in that *neither* of them is stating *the* proposition that that stealing money is wrong.

3) Also, and perhaps most famously, appraiser relativism implies that one person can claim that stealing money is wrong while another can claim that stealing money is not wrong, without either of these two speakers being mistaken. This, too, flies in the face of the basic linguistic intuitions of the great majority of competent speakers of English.

---

[267] Streiffer 2003, Chapter 1. I am taking a few liberties with Streiffer's examples here for ease of exposition.

As Streiffer goes on to note,[268] some relativists have responded to these sorts of criticisms by attempting to show why our basic linguistic intuitions might be mistaken in these cases; but these responses seem inadequate. David Wong admits that it is *generally* true that speakers who utter the same moral sentences are in agreement, etc., since people from the same culture (broadly speaking) will have common moral values. However, he claims that this general trend makes us tend to overlook the fact that, when significantly different cultures come into contact and their members take the time to discuss their moral views in depth, it can be seen that these intercultural moral differences are too fundamental to allow for rational resolution, so that our linguistic intuitions are based on faulty overgeneralizations from this limited set of cases.[269] But as Streiffer points out,[270] this response fails. The attraction of appraiser relativism is that it is meant to help explain intractable moral disagreement, but there are vast numbers of moral disagreements that take place within a single society, a single culture, and even a single family. Moreover, even Wong's own examples of intractable moral disagreements – disagreements about the morality of abortion, welfare, and taxes, for example – regularly occur between members of the same society.

It might be thought that this rebuttal merely limits the explanatory power of appraiser relativism to disagreements between cultures, so that Wong's form of moral relativism at least applies between pairs of cultures with radically different moral views (if such pairs of cultures exist). But Streiffer preempts this retreat quite effectively:

---

[268] *Ibid.*, pp.13ff
[269] Wong, 1984

> I do not think that narrowing the range of disagreements will help the Appraiser Relativist. Even if it were rare for an assertion of a moral sentence to be consistent with an assertion of that sentence's syntactic negation, our linguistic intuitions would nonetheless reflect these rare cases, and we would not be tempted to overgeneralize. Consider the sentence 'Dogs dogs fight fight.' Upon canvassing the possible contexts in which that sentence might be used, many people find it intuitively obvious that there are no contexts in which that sentence is syntactically acceptable. But once you point out a context in which the speaker is using that sentence to say that dogs that dogs fight, also fight, the intuition goes away… In general, an intuition that something is impossible is much more sensitive to counterexamples than it is to confirming instances.[271]

Therefore, since our linguistic intuitions about moral sentences survive what Wong claims are counterexamples, we have good reason to maintain that our intuitions are correct and that the supposed counterexamples are no such things. All in all, as Streiffer points out, we have excellent *prima facie* grounds for suspecting that moral disagreements are better explained in some other way – for instance, by pointing out that we lack an adequate and generally accepted notion of personhood (in the case of abortion), of justice (in the case of capital punishment), etc.[272]

It follows from all this that any successful argument from moral disagreement will need to be very powerful if it is to meet the difficult challenge of establishing such a *prima facie* unlikely view as appraiser relativism.

---

[270] Streiffer, pp.14-16
[271] *Ibid.*, p.15
[272] *Ibid.*, p.26

In setting out the problems with appraiser relativism, I have followed Streiffer's analysis quite closely. In what follows, I will move in a somewhat different direction, and make use of the parallel arguments approach.

Harman's Version of the Argument from Moral Disagreement

Harman's presentation of the argument from moral disagreement can be found in Sections 1.2 and 1.3 of his book-length debate with Judith Jarvis Thomson.[273] This argument has a particularly important function in that work, since (as I understand it) Harman intends it to be the central argument he presents for relativism there.

Harman begins that section, whose title promises that it will be about "Explaining Moral Diversity", as follows:

> In this and the following section, I argue that the following claim is a reasonable inference from the most plausible explanation of moral diversity.
>
> > There is no single true morality. There are many different moral frameworks, none of which is more correct than the others.
>
> I begin by mentioning data to be explained: the nature and extent of moral diversity.[274]

Harman follows up admirably in *mentioning* the extent of moral diversity over the next two sections. However, he seems less keen to keep his earlier promise that he will go on to argue that the indented claim is a reasonable inference from the data he presents. He does mention, on two or three occasions, how a relativist might *account* for various

---

[273] Harman and Thomson, 1996
[274] *Ibid.*, p.8

features of moral language and discourse, but he never even claims outright, let alone argues, that these relativistic explanations would be in some way better than any others. It is true, no doubt, that a relativist might *account* for certain instances of moral discourse; but this does not provide good grounds for thinking that the relativistic account is the best available explanation or even a good one.

This, together with the powerful *prima facie* case against appraiser relativism outlined above, makes it very difficult indeed to see how – on any reasonable view – Harman can be seen as having met his argumentative burdens here.

Nevertheless, so as not to dismiss what may be an implicit argument too quickly, I will assume that Harman intends his argument to go roughly as follows:

P1. Moral diversity exists.

P2. The moral diversity that exists involves some intractable disagreements.

P3. Appraiser relativism is the best explanation of intractable disagreements.

C. Therefore, appraiser relativism is true.

Criticism of Harman's Argument from Moral Disagreement

In order for this argument to work, P3 must be correct. But is it? If it is, then intractable agreements about morality, in themselves, provide at least some reasonably good evidence for appraiser relativism. *How* good that evidence is will depend, among other

things, on the viability of alternative explanations. However, a comparison of this argument with parallel arguments about other sorts of disagreements makes it unclear why one should take this fact to provide reasonably good evidence for such a view.

To start with a simple case, let us suppose that Mr. and Mrs. Jang are out shopping. Mr. Jang feels confident that there is already an unopened container of orange juice at home, but Mrs. Jang feels sure that all the orange juice has been drunk. They stand in front of the orange juice at the market, attempting to assure one another that their views are correct; but neither can make any headway. Global relativists might explain this by saying that, since Mr. and Mrs. Jang sincerely hold opposing views on the matter, they must both be right. However, it does not follow from this that the Jangs' disagreement provides *reasonable evidence* for global relativism: it's just that global relativists would see things this way. Those who believe in the objective truth of some statements about orange juice would not have reason to doubt their commitment to this objectivity by hearing the Jangs argue. They can unproblematically maintain that either or both of the Jangs are simply in error about the facts.

It might be objected, however, that the case of Mr. and Mrs. Jang is not really parallel with an instance of the relevant kind of moral disagreement, since Mr. and Mrs. Jang can resolve the issue to their satisfaction once they return home and inspect their refrigerator. Such an objection, I feel, confuses a question about reality (i.e. about whether there really are objective moral facts, or whether there really is enough orange juice at home) with a question about our knowledge of that reality (i.e. whether a given pair of subjects are or

could be in a position to know the answers to these questions about reality). Nevertheless, I will entertain this objection by considering what might not be as simple a case: the case of John and Janice. Both John and Janice have a great deal of confidence in their beliefs on a number of issues. The two of them are now watching the second hand of a clock with great interest, waiting for it to reach the top of its circuit. When it does, they have agreed that they will use their intuitive and intellectual powers to determine whether the population of Vancouver Island a thousand years – to the second – before that instant was an odd or an even number. The second hand reaches the top, and Janice shouts out, "Even!" while John shouts out, "Odd!". Now, surely, either John is correct or Janice is (even in the very unlikely event that there happened to be no people at all on Vancouver Island at that second, we can count that as even). However, there seems to be no rational means whatever by which they can resolve their disagreement. Even if they render themselves fully informed by all available means, it does not seem that they will be any closer to discovering which of them is right: the relevant evidence simply does not exist any longer and cannot be reconstructed. Now, someone who is already committed to global relativism (say) might explain this disagreement in such a way that both Janice and John are correct, despite the fact that their beliefs seem to most of us to be incompatible with one another. However, the most natural judgment to make about such a situation is that either Janice is objectively right about this, or else that John is, but that neither we nor they have any way of knowing which it is, and that they both ought to be a little less confident in their assertions and beliefs. This natural, objectivist view does not seem to be threatened in any way by the fact that Janice and John cannot resolve their disagreement.

But the appraiser relativist might object, again, that Janice and John do not have an unresolvable disagreement in this case since there are some facts – albeit facts that are no longer accessible – that would settle the matter. If John and Janice could use a time machine, say, to go back to the second in question, and if they could then freeze time in such a way that they could run around and take a census while being assured that nobody else on the island would leave it, come to it, be born on it, or die on it, they would have the information they need to resolve the issue. Therefore, the objection might go, even though there is no way for them to recover this information in the actual world, the fact that such information-gathering is theoretically possible is enough to disqualify the Janet and John case as relevantly parallel. Further, the appraiser relativist might insist that the theoretical possibility of resolving the Janice and John dispute by obtaining this new information shows that Janice and John are not perfectly informed on all relevant issues, which is essential in order to show that some dispute is intractable.

It is not easy to see how to modify these attempted parallels to satisfy this requirement. There do seem to be cases where it is physically impossible for someone to know something despite being confident enough to assert that he or she does, such as when Jim claims to be certain that he knows the position and momentum of a certain subatomic particle. But if we are to be permitted to assume, counterfactually, that John and Janice can obtain normally inaccessible information about the precise population of ancient Vancouver Island, why not also assume, counterfactually, that Jim can violate Heisenberg's uncertainty principle?

To avoid all these problems, let us just hypothesize, for the benefit of the appraiser relativist, a very generic parallel: there exists a person X, and there exists a person Y, such that X and Y confidently hold apparently incompatible beliefs on some issue Z, such that Z is not an ethical issue and there is no way for X and Y to resolve their disagreement about Z. The problem, however, is that even this very generic description violates the criteria that X and Y need to be perfectly informed just as the Janice and John story does. For there are some relevant facts – namely, the correct answer to issue Z and all facts that follow from that fact in some way – that neither X or Y knows. For this reason, any parallel with objective facts seems impossible.

We can imagine, though, that appraiser relativists may be happy with this conclusion. They might hold that this is exactly their point: they might hold that questions of morality are questions of value, not of fact; that there is a strict dichotomy between facts and values; and that, while all factual matters are resolvable in principle by fully informed, rational people, matters of value are not.

If all this were true, then we would have good reason to doubt that there could be any adequate parallel to be drawn between a dispute over morality and a dispute over facts. But why should we accept that all this *is* true? It is no use pointing to a large range of cases where *apparently* reasonable and adequately-informed individuals have been unable to persuade one another on some ethical matter after a given amount of time. Such cases are always open to the responses that

a) despite the confidence of the appraiser relativist, the individuals in question might not be fully rational, and (in particular) they might not have drawn out all the relevant logical implications from their beliefs;
b) since the individuals under observation are presumably far from omniscient, they might after all not be adequately informed;
c) either or both of these individuals might be prey to psychological factors that make it difficult for them to accept that they are mistaken when they should know that they are; and
d) more time might be needed to resolve the dispute.

Furthermore, it is not difficult to find issues in science, philosophy, and elsewhere that appear to be as intractable as those in ethics. What reason do we have for thinking otherwise? The empirical evidence Harman presents is certainly inadequate to lead us to that conclusion, and it is not clear how he could have done better. However, there also seems to be no *a priori* reason for thinking that moral disagreements are intractable while those of science and philosophy, say, are not. Those who are not appraiser relativists seem to be under no compulsion, therefore, to accept such an assumption.

But let us extend the principle of charity and *assume*, for the moment, that there is good reason to believe that disagreements in morality, unlike other disagreements, are unresolvable in principle. Would appraiser relativism follow *then*? No. Even with that assumption, there is no reason for those inclined to believe in objective morality to accept the relativistic conclusion, since a plausible interpretation – and also, it seems, the most

natural – is that the moral facts in question are simply unknowable. The appropriate position to adopt, in that case, would be skepticism rather than relativism. One would only be constrained to accept the truth of appraiser relativism if, in addition to the above-mentioned assumption, one had good reason to adopt the further assumption that moral claims cannot be mind-independently true or false (perhaps because one has good reason to believe that morality is constructed). But no such reasons are offered by the argument from moral disagreement.

Still, let us extend the principle of the charity even further and assume *both* that moral disagreements are unresolvable in principle *and* that morality is mind-dependent. From both these assumptions together, there seems to be some good reason to believe that appraiser relativism is true. However, it is worth remembering, first of all, that these two additional assumptions are needed in order for appraiser relativism to follow from the argument from moral disagreement; and second, that those who oppose appraiser relativism have no good reason to accept either assumption. Also, it should be noted that these two assumptions, together, seem to imply appraiser relativism already, even without the other premises of the argument.

What follows from this is that the argument from moral disagreement is only effective if one already assumes (which one need not) that appraiser relativism is true. Since the argument is meant to *establish* the truth of appraiser relativism, it is clear that it commits the fallacy of begging the question.

Hence, Harman's argument from moral disagreement fails. That leaves us with only one more argument to consider. This sole remaining argument only supports agent relativism, which accepts that it can be an objective matter whether some particular agent ought to perform some particular act. Hence, we can already rest assured that Harman has presented no argument that threatens this kind of objectivity. The only thing left in question at this point is whether there are moral imperatives that are binding on us whether or not we can be made to care about them or the moral patients they concern. I turn now to Harman's final argument, for this second type of relativism.

I. Agent Relativism and the Argument from Moral Reasons

I will begin this section by presenting what I take to be the most important problems with agent relativism. As in the previous section, this will make clear the hurdles that an argument for the sort of relativism in question will need to overcome. Next, I will present Harman's discussion of moral reasons, and see which sort of argument Harman can most charitably be construed as making. Along the way, I will argue that Harman, and indeed many other metaethicists who have dealt with these issues, have obscured the issue by not clarifying what they mean by a moral reason. I will use this discussion to help clarify what I take to be the strongest version of the sort of argument from moral reasons that Harman could be construed as making. I will then try to show that, even on this interpretation, the parallel arguments strategy shows that moral anti-relativists seem to have no reason to fear this argument.

Problems with Agent Relativism

Again, I follow Streiffer in setting out the *prima facie* case against agent relativism. Streiffer presents two arguments to this end.[275] The first of these, which he calls the *Argument from Extreme Heinousness*, he rejects as ineffective; but the second, the *Argument from Morally Relevant Differences*, he accepts as providing a strong case against agent relativism. I will discuss both of these arguments here, and take issue with Streiffer's reasons for rejecting the *Argument from Extreme Heinousness*.

The *Argument from Extreme Heinousness* rests on the single premise that there are some actions so heinous that they should never be performed. The existence of any such action logically implies that, for any agent X, it would be wrong for X to perform that action. But it is precisely this that agent relativism denies. Therefore, it follows that agent relativism is incorrect.

Streiffer's reason for rejecting the *Argument from Extreme Heinousness* is as follows:

> Some people think that for any action alleged to be sufficiently heinous, they can conceive of possible circumstances where an agent is forced to choose between performing that action on one person and performing the same action on some absurdly large number of people. In such circumstances, these people say, the consequences of not performing the action in question are themselves so tragic that morality permits, perhaps even requires, that the agent perform the heinous action on one person. If you think, as I do, that numbers can matter, then you are unlikely to be persuaded by the Argument from Extreme Heinousness.[276]

---

[275] Streiffer 2003, pp.31-32
[276] *Ibid.*, p.31

The reasoning here seems to be flawed. Surely, it is reasonable to assume that those considering whether an action is extremely heinous are capable of taking the circumstances of the action into consideration as their moral theories require. Let us consider some particular action that may be deemed extremely heinous by some: Murphy is watching a movie at the cinema, but he and many other filmgoers are distracted by a shrieking infant whose parents have, perhaps unwisely, brought it to the movies with them. After waiting politely for the parents to remove the infant, Murphy grabs it and hurls it into a padded chute in the wall that has been installed by the cinema for the purpose of swiftly transporting shrieking infants from the cinema to a room in the basement from which its parents can retrieve it at their convenience. Now, it seems that Streiffer feels that a Kantian, or someone else who denies 'that numbers can matter', can classify this action as unequivocally heinous, whereas an act utilitarian (say) cannot; and that this is a problem with the argument from extreme heinousness. But is that correct? It seems that, *if* it follows from act utilitarian principles that the throwing of the baby is morally justified or obligatory in this case, then on that same view the act will either a) not be extremely heinous after all (because it is morally permissible or obligatory), or else b) be morally permissible or obligatory *despite* its being extremely heinous. But neither of these seems to present a difficulty for the *Argument from Extreme Heinousness*. If a), then the objection to the argument misses the point because it has failed to identify an action that, *given the circumstances in question*, would be either extremely heinous or wrong for anyone to perform. There is a world of difference between an action being such that nobody ought to perform it in a given set of circumstances and an action's being such that no imaginable circumstances could justify

its being performed. If b), on the other hand, then the objection gives us reason to reject the *letter* of the argument, but not its *spirit*: the essential point of the argument is that there is something that should not be done by any agent in exactly the circumstances given. If it turns out that, on some analysis, there are some extremely heinous but morally permissible or obligatory actions, then the question of the heinousness of those actions will be irrelevant to the point at issue.

Still, the argument from heinousness faces an important difficulty: if there are agent relativists who genuinely believe that there are no actions such that it would be wrong for any agent to perform them, then they will have no reason to be persuaded by the *Argument from Extreme Heinousness*. Moreover, it seems that any agent relativist who has thought carefully about the ramifications of the view, and who nonetheless maintains agent relativism, will have that belief. Hence, the *Argument from Extreme Heinousness* seems to commit the fallacy of begging the question against any thoughtful and committed agent relativist.

The other main argument against agent relativism, which Streiffer calls *The Argument From Morally Relevant Differences*,[277] is (as Streiffer claims) somewhat stronger. The argument goes roughly as follows (I am modifying Streiffer's argument here in a way that, I hope, strengthens it): agent relativism implies that, for any action X, there are two agents Y and Z such that it is morally permissible for Y to perform action X and not morally permissible for Z to perform action X. But it seems there must be some morally relevant distinction between Y and Z if that is the case. Perhaps the distinction is that Y

promised to perform action X, while Z did not. But in that case, the two actions are not really the same: doing something that you promised to do is not morally interchangeable with doing something that you never promised to do. Perhaps the distinction lies not in the circumstances surrounding action X, but rather in the inner states of Y and Z: maybe Y is committed to a moral view according to which it is important to perform action X, while Z is not. However, in this case (granting temporarily, for the sake of argument, that this is a morally relevant sort of difference) the grounds for making this distinction apparently rest on a broader moral principle that applies to both Y and Z: 'You are morally obliged to do X if, and only if, you have a moral view that entails that it is important to perform action X'. Further, there is now a particular moral prohibition – the prohibition against neglecting to perform action X if your moral view entails that you ought to perform action X – that both Y and Z would be wrong not to follow, and this applies equally to both Y and Z even by the agent relativist's lights, and moreover it is a fuller statement of precisely the moral situation that is needed by the agent relativist to make the case for relativism. So it seems that agent relativism entails its own negation. Nor does it seem possible for an agent relativist to overcome this difficulty simply by insisting that it is merely a brute fact, dependent on no broader, underlying reason, that it is permissible for Y and not Z to perform action X: even if the distinction *is* merely a brute fact, the rightness of Y and not Z performing that action is covered by the general rule 'It is wrong for Z, but not Y, to perform action Z'.

The *Argument from Extreme Heinousness* presents a strong case against agent relativism for anyone who agrees with the general moral view that there are at least some particular

---

[277] *Ibid.*, p.32

actions, and some particular circumstances, such that it would be wrong for anyone to perform those actions under those circumstances. The *Argument from Morally Relevant Differences* presents an even stronger case against agent relativism, and one that seems effective even against those who deny that there are any such actions and sets of circumstances. Together, they make a very powerful initial case against agent relativism. Therefore, the burden on those who would defend agent relativism with the argument from moral reasons is a very heavy one. I will now consider how Harman attempts to deal with this burden.

Harman's Presentation of the Case for Agent Relativism

Harman begins his presentation of this case by attempting an analysis of the term 'ought' in morality:

> Strictly speaking, we can say, an "ought" statement has the form, 'Relative to C, P ought to do D." For example, "Relative to the fact that you have promised to do something, you ought to do it." An all-things-considered "ought" judgment would then have the form, "All things considered, P ought to do D" – in other words, "Relative to all things considered, P ought to do D." This amounts to the claim that the prima-facie "ought" is basic and that the all-things-considered "ought" is to be defined in terms of the prima-facie "ought." We might then define the basic, prima-facie "ought" like this: "Relative to C, P ought to do D" means the same as "C gives P a reason to do D."[278]

There is much here that is more puzzling than one might desire. To start with, to say that some agent ought to do something *relative* to something else is not a usual way of speaking, and needs to be translated into clearer language if it is to make any sense. But

---

[278] Harman 1977, p.121

how to translate it? From Harman's first example, "Relative to the fact that you have promised to do something, you ought to do it", it at first appears that 'relative to' is here serving the purpose that 'given' does in everyday language. But while it seems natural to say that I ought to perform an action *given* that I promised to do it, this interpretation of Harman's phrasing runs into difficulties when we come to his second example. If Terry has been told by Sam's rich uncle to tell Sam about the location of a precious diamond, and Terry has good reason to believe that Sam, once told about this, would immediately sell the diamond and gamble away the money, that (together with other considerations) might provide Terry with a good moral reason not to honour the last wishes of Sam's uncle by telling Sam this information at once. However, if we say "All things considered, Terry should not fulfill Sam's uncle's final wish", we are not merely saying "*Given* all things considered, Terry should not fulfill Sam's uncle's final wish". This latter statement would mean that, given all the things that someone (presumably the speaker, or Terry) has *considered*, Terry should not fulfill the wish. But 'all things considered' in this sort of context does not imply that the truth of the statement that follows is based solely on what the agent, or the speaker, has considered. Even if Terry had not considered, and had not even known about, Sam's gambling addiction, it would nonetheless be correct for the speaker to say that Terry ought not give the money to Sam, 'all things considered'. Further, there does not seem to be any difficulty in saying "All things considered, Terry ought to hang on to the money – unless I am mistaken". The reason why this statement makes sense is that 'all things considered' means nothing more than 'considering all the relevant facts'. It is because one can have incomplete information about a situation that one can be mistaken about what someone should do, all

things considered in that situation, even if one knows what one has considered and what follows from a consideration of that limited range of factors.

The suggestion at the end of this paragraph of Harman's is somewhat less perplexing. Harman tells us there that "'Relative to C, P ought to do D' means the same as 'C gives P a reason to do D'", which is intelligible if we substitute in 'Given C' for 'Relative to C'. However, this still leaves unexplained how any of this talk about given circumstances relates to agent relativism. Consider, again, Terry, who coincidentally had an eerily familiar encounter with Sim's dying uncle immediately after her encounter with Sam's dying uncle. The major difference between the two cases is that Sim, unlike Sam, is a philanthropist and not a gambling addict. Let us suppose that there is also no further consideration that gets in the way of Terry's *prima facie* duty to tell Sim about the diamond that is waiting for her. We might then – while wholly denying the truth of agent relativism – agree that, *given* that Terry made a promise to Sim's uncle, and *given* further that she has no morally relevant reason not to tell Sim about the diamond, Terry has a moral obligation to fulfill Sim's uncle's dying wish. In saying this we may, again, concede that anyone in a position relevantly identical with Terry's ought to do the same. Finally, if 'relative to' is to be translated as 'given that', it seems that a proper moral claim is always a claim made 'relative to' everything, in this sense of 'relative'. If there is some *incomplete* set of relevant factors R that I am considering in mulling over a moral judgment, and some further set of factors N that I am not considering even though they are relevant, then – as agent relativists would presumably agree – my moral judgment is simply a hasty one on the basis of inadequate evidence, and quite possibly incorrect. A

good moral judgment, according to agent relativism *and* anti-relativism, is one that takes all relevant factors into account (and this should not, of course, be confused with the issue of whether it is possible to be sure that one *has* taken everything relevant into account). So even if Harman intends us to read 'relative to' as 'given that' in these sentences, this passage does not seem to offer any case for agent relativism.

Harman attempts to further analyze 'ought' by clarifying what he takes to be an important difference between "the moral ought" and "the ought of rationality".[279] We sometimes say, he points out, that we say that – given what we know to be someone's intentions – that person *ought* to do some particular thing to satisfy such intentions, even if we do not personally endorse either the intentions or the action. For instance, we might say that a bank robber *ought* to use the back door, even if we do not endorse bank robbing and are only making a non-moral claim about how the robber might best satisfy his or her purposes. Harman uses this example to show the difference between these two uses of ought – the ought of rationality and the moral ought – respectively, as follows:

> When I say that a bank robber ought to use the rear door, I do not endorse his doing so; I am not indicating that I am in favor of his doing so. But, when I say that the bank robber ought to give up his trade, I do endorse his doing so; I am indicating that I am in favor of his giving up his trade. The moral "ought" is therefore the "ought" of rationality plus something else. When I use the moral "ought", I presuppose that the agent and my audience accept certain practical principles that I also accept, and I make my judgment relative to those principles.[280]

Harman seems to feel that this is a sufficient basis for a case for agent relativism:

---

[279] *Ibid.*, p.122

> Consider also the ways in which we react on learning that an agent does not have the goals we assume he had. I judge that a bank robber ought to use the rear door because I suppose that his goal is to rob the bank and get away unobserved. If I learn that he does not intend to rob this bank but is merely making a deposit, then I withdraw my judgment that he ought to use the rear door and I say that I was mistaken. On the other hand, suppose that I use the moral "ought" to say that the bank robber ought to give up his trade. In saying this, I am presupposing that certain principles that the bank robber accepts give him reasons to stop being a bank robber. *If I learn that the bank robber is, however, totally amoral and that, given his goals and plans, he has absolutely no reason not to continue being a bank robber, then although I will withdraw my judgment I will not do so by saying that I was mistaken.*[281]

One difficulty with all this, of course, is that it *assumes* rather than *establishes* the truth of agent relativism. But nothing leading up to this passage has established that agent relativism is true, either. This robs Harman's case for agent relativism here of all its force. However, he does also seek to establish agent relativism in his paper, 'Moral Relativism Defended'. I will turn to his discussion in that paper now.

Harman begins by admitting that the relativity of moral obligations to circumstances or considerations (as in my stories of Sam and Sim) does not imply a relativity of moral obligations to motivating attitudes, and it is "not as easy to exhibit" this second kind of relativity.[282] "Normally", Harman continues,

> …a speaker who makes this sort of moral 'ought' judgment intends the relevant motivating attitudes to be ones that the speaker shares with the agent and the audience,

---

[280] *Ibid.*, pp.122-123
[281] *Ibid.* (italics mine)
[282] Harman 2000a, p.9

> and normally it will be obvious what attitudes these are. But sometimes a speaker does invoke different attitudes by invoking a morality the speaker does not share. Someone may say, for example, 'As a Christian, you ought to turn the other cheek; I, however, propose to strike back.' A spy who has been found out by a friend might say, 'As a citizen, you ought to turn me in, but I hope that you will not.' In these and similar cases, a speaker makes a moral 'ought' judgment that is explicitly relative to motivating attitudes that the speaker does not share.[283]

This seems to be the full extent of Harman's argument. He turns from this point almost immediately to a consideration of moral bargaining under a section by that name, and speaks no further of the relativity of moral obligations to motivating attitudes.[284] It is ultimately on the strength of the two cases discussed in this paragraph, therefore, that Harman's explicit case for agent relativism depends. Let us, therefore, examine these two cases carefully.

First, we have the case of the non-pacifist non-Christian who is speaking to the pacifist Christian. The non-Christian proposes to 'strike back' against someone who presumably wronged both the characters in question. A charitable reading of this is that the non-Christian is implying, through his or her proposal to strike back, that striking back would be morally permissible or obligatory for the non-Christian. However, we are also given that the non-Christian is morally advising the Christian to turn the other cheek instead. Why, given that the non-Christian feels that the situation calls for his or her own retaliation, does he or she feel that it does not call for the Christian's retaliation (assuming, for the moment, that their circumstances are otherwise relevantly identical)?

---

[283] *Ibid.*, p.9

[284] The one exception is in Harman 2000e, which I will consider shortly.

Presumably, the difference is accounted for on the basis of a belief held by the non-Christian to the effect that, in cases relevantly like this one, one should follow the teachings of one's religion (if one has a religion), provided perhaps that various other conditions are fulfilled. But then we have the same difficulty that Streiffer has pointed out: the speaker then seems to imply that the rule that one must follow one's religious teachings in such situations applies to all agents in the scenario in question. But the scenario was set up precisely to establish agent relativism – in other words, to show that there are scenarios in which there are no rules that apply to all agents. Hence, this counterexample fails to establish agent relativism.

Second, we have the case of the spy who, after being discovered by a friend, says "As a citizen, you ought to turn me in, but I hope that you will not." But what, exactly, does this spy think in general about turning in spies? Suppose that the spy discovers that his friend is a double agent. Would he feel obliged to turn his friend in then? Perhaps not: the spy may feel that the bonds of friendship always trump impersonal issues of national security. In that case, his request that his friend not turn him in does not point to an instance in which a speaker advises someone else to follow a moral code of which the speaker does not approve. Perhaps, instead, the spy feels that anyone discovering that anyone else is a spy has a moral obligation to turn the spy in. In that case, there is no reason to think that his view implies an acceptance of agent relativism, and his hope that his friend not turn him in should be read as just what it seems to be – a hope, rather than a moral consideration. Finally, there is the possibility that the spy really believes that it is morally obligatory for him to turn in any spies *he* discovers, but immoral for him to be

turned in by anyone else. Such a view might be understandable, for instance, if the spy felt he had good reasons for thinking that the cause he is working toward is good enough to justify spying, but that the other cause isn't. But this again seems to be just an instance of the spy following the moral rule that spying is ethically prohibited unless one is working for a good cause. Since Harman does not provide these important details, there is no way to assess the case in his favour. However, we have no reason to think that he has any way to overcome the *prima facie* case against agent relativism as outlined above.

Hence, Harman's explicit case for agent relativism is a non-starter. His remarks on the subject can only be interpreted favorably, therefore, if he is relying on some *implicit* argument for the position he is assuming, and if he is read as assuming also that his audience will be aware of this implicit argument. Following the principle of charity one more time, I will make this interpretation of Harman; and with this in mind I will try now to set out the best *implicit* argument that can consistently be read into his work.

A Stronger Argument for Agent Relativism

Streiffer sets out, on behalf of agent relativists, the following version of what I am calling the 'argument from moral reasons':

> **P1 (The Practicality Requirement)**: "For any action, if an agent is morally required to perform that action, then there is a reason for that agent to perform that action."[285]

[285] Streiffer 2003, p.30

P2 (Reasons Relativism): "For any action, there [could be] some agent such that there is no reason for that agent to perform that action."[286]

Conclusion (Agent Relativism): "For any action, there [could be] some agent who is not morally required to perform that action."[287]

Streiffer goes on to discuss, for the following twenty-two pages, the various moves that can be made for and against this argument. His verdict is that the Practicality Requirement ought to be accepted, but that Reasons Relativism ought not. While I grant the force of Streiffer's arguments in this connection, it seems to me that there is something odd involved in the Practicality Requirement, which makes it is as strange to argue for it as it is to reject it. Once this is shown, I will suggest, it begins to seem that there is something wrong with the argument from moral reasons as a whole. I turn to this discussion presently.

What Is Meant by a 'Reason'?

It is not my intention here to give a definite answer to this question, but rather to point to some difficulties that seem to have arisen owing to an apparent lack of clarity in the usage of the term 'reason'. Let us first consider Streiffer's Practicality Requrement. It tells us that, if X is morally obliged to do Y, then there is a *reason* for X to do Y. Now, what is meant by 'reason' here?

---

[286] *Ibid.* Streiffer has 'is' rather than 'could be'. I have modified his statement of the premise in order to present agent relativism in the fairest possible light.

There are, it seems, two main meanings of that term in this sort of context. First, there is what I will call the **non-moral** meaning: to say that X has a (non-moral) reason to do Y is to say that X has some set of interests, desires, purposes, motivations, etc. that would be served in some way by his or her doing Y. I take non-moral reasons to include reasons one might have for doing the aesthetically-best thing (if one cares overwhelmingly about aesthetics in some situation), the most polite thing (if one cares overwhelmingly about etiquette in some situation), etc.

Second, there is what I will call the **moral** meaning: to say that X has a (moral) reason to do Y is to say that morality in some way demands that X perform Y. Both of these definitions are familiar to all, I trust, from everyday usage; and they correspond quite closely with Harman's 'ought of rationality' and 'moral ought'; so I take it no further examples are needed here. Having combed the entry under 'reason' in the *Oxford English Dictionary* and found no other relevant definitions of the term, I take myself to be warranted in concluding that these two meanings exhaust, for all practical purposes, the chief meanings of the term in contexts like "X has a reason to do Y" or "There is a reason for X to do Y" (which may or may not mean just the same thing. In order not to prejudice the case against those who hold non-moral reasons as contingently or necessarily identical with moral reasons in all cases, I am tentatively adopting an agnostic position on that issue; but certainly, the claim that the two are identical is not self-evident and would, I suggest, need to be argued for).

---

[287] *Ibid.*

Now, I take it as clear that, if one is to use a term in a work of philosophy, one must either use the term according to its conventional meaning, or else signal one's unusual usage of it by providing one's own definition of the term. Since Streiffer does not provide such a definition of 'reason', I take it that he must be using that term according to one of its conventional meanings.

With that, and the two conventional uses of 'reason' in similar contexts, in mind, let us consider the two things that could be meant by Streiffer's Practicality Requirement.

Practicality Requirement 1 (the *non-moral* reading): If X is morally obliged to do Y, then X has some set of interests, desires, purposes, motivations, etc. that would in some way be served by his or her doing Y.

Practicality Requirement 2 (the *moral* reading): If X is morally obliged to do Y, then morality in some way demands that X do Y.

At this point, I suspect, it should be clear that the game is up. For let us suppose, first, that the non-moral understanding of 'reasons' is intended. In that case, P1 should be read as Practicality Requirement 1. On that reading of 'reasons', also, P2 (Reasons Relativism) will be so obvious as hardly to be worth stating: since there are, and more broadly, *could be*, genuine sociopaths (such as Harman's Hitler), and since there also seem to be, and at least *could be*, other people who have literally no concern for their own well-being and cannot be motivated to care about it, then it is true that, for any

action, there could be some agent whose motivations, etc. would not be served by his or her performing that action (which is what P2 would say on this reading of 'reasons'). But since this version of P2 seems uncontroversially to be true on this reading (even anti-relativists like Streiffer do not deny it), the argument from moral reasons is merely an inference from a single premise – Practicality Requirement 1 – to agent relativism. But those opposed to agent relativism have, and are given, no good reason to accept Practicality Requirement 1. To establish Practicality Requirement 1, agent relativism – or something very much like it – would presumably need to be established first. Therefore, on the non-moral reading of 'reasons', this argument seems to beg the question.

The difficulty that non-relativists have no good reason to accept P1 can be solved if we read it, instead, as Practicality Requirement 2. However, Practicality Requirement 2 is trivially true. Therefore, all the heavy lifting in this second reading of the argument must be done by the sole remaining premise, P2 (Reasons Relativism). But in order to avoid the fallacy of equivocation, the term 'reasons' in the statement of reasons relativism must, following the understanding of that same term in what is now P1, be understood under the *moral* reading. Hence, while P1 will be trivially true, P2 must now be read as saying that for any action, there could be some agent such that morality will not demand that that agent perform that action. But P2 is clearly nothing more than a restatement of agent relativism, which is just what is being argued for. So the argument begs the question on this reading also.

To summarize: if 'reasons' is understood on the non-moral interpretation, then P1 of the argument from moral reasons is question-begging. If 'reasons' is understood on the moral interpretation, then P2 is question-begging. These seem to be the only two plausible readings of 'reasons' in this context, and no other reading is suggested. Therefore, we are warranted in dismissing this version of the argument from moral reasons as begging the question.

Interestingly, similar difficulties can be found in every other presentation of the argument from moral reasons I have been able to find. In none of these cases does the writer either a) define the term 'reason' in a way that avoids the dilemma I have indicated, or b) dismiss the argument at once as begging the question. Instead, each has engaged in a lengthy discussion of the merits (or otherwise) of the argument. If my analysis here has been correct, however, any discussion of the argument that does not amount to a mere dismissal of it for committing the fallacy of begging the question must first be justified by a good definition of 'reasons' that differs relevantly from both the non-moral and the moral sense. Lest the reader think me too cavalier in dismissing this argument, here are two further sample presentations of it by prominent metaethicists, both of which go the same way.

The first is from David Brink's *Moral Realism and the Foundations of Ethics*:

1. To be under a moral obligation to do x, one must have reason to do x.
2. One has a reason to do x just in case x would contribute to the satisfaction of one's desires.

3. Hence one can have a moral obligation to do x only if doing x would contribute to the satisfaction of one's desires.
4. No everyone has the same desires.
5. Hence, there is no single set of moral requirements that applies to everyone; there will be different moral requirements that apply to different people in virtue of their different motivational sets.[288]

Brink goes on to discuss this argument at considerable length, ultimately declaring it unsound. He grants from the outset, however, that "[t]he first two premises are, of course, crucial."[289] It seems from the analysis I have just presented, however, that at most one of these premises can be crucial; the other is simply true by definition. If 'reason' here is to be understood in the non-moral sense, then P2 is trivially true; if it is to be understood in the moral sense, then P1 is trivially true.

My second example comes from Russ Shafer-Landau's excellent book, *Moral Realism: A Defense*. In Part IV of that work, Shafer-Landau carefully considers what he calls the *Desire-Dependence Argument*:

1. Necessarily, if S is morally obliged to Φ at *t*, then S has a good reason to Φ at *t*. (*Moral Rationalism*)
2. Necessarily, if S has a good reason to Φ at *t*, then S can be motivated to Φ at *t*. (*Reasons Internalism*)
3. Necessarily, if S can be motivated to Φ at *t*, then S must, at t, either desire to Φ, or desire to ψ, and believe that by Φ–ing S will ψ. (*Motivational Humeanism*)
4. Therefore, necessarily, if S is morally obligated to Φ at *t*, then S must, at t, either desire to Φ, or desire to ψ, and believe that by Φ–ing S will ψ.[290]

---

[288] Brink 1989, p.52
[289] *Ibid.*, p.53
[290] Shafer-Landau 2003b, p.170

Again, depending on the meaning given to 'has a good reason', either Moral Rationalism or Reasons Internalism will be a coherent and interesting view. But unless some new usage of the term is intended – and, as usual, no new definition is given, and it is not clear what it could plausibly be – it is not possible for *both* of these views to be coherent and interesting without equivocation on 'reason'. Nonetheless, Shafer-Landau devotes a full chapter each to a discussion of the merits of these two theses.

Since Brink's version of the argument from moral reasons really amounts to a fuller explication of the one considered by Streiffer, it seems to suffer from the same difficulties. Shafer-Landau's argument is not a version of the argument from moral reasons: its conclusion is distinct from agent relativism, and its third premise brings in the question of moral motivation. Nonetheless, either P1 or P2 of Shafer-Landau's argument seems to be trivially true and hence redundant, but we cannot tell which it is until we have learned how Shafer-Landau intends the term 'reason' to be understood.

With all this in mind, we can turn to what I take to be Harman's most sophisticated discussion of relativism: his 'Is there a Single True Morality?' There, he presents his most careful account of the argument from moral reasons. He begins with a general discussion of reasons:

> Consider what it is for someone to have a sufficient reason to do something. Naturalism requires that this should be explained in terms congenial to science. We cannot simply

> treat this as irreducibly normative, saying, for example, that someone has a sufficient reason to do something if and only if he or she ought to do it.[291]

Depending on whether Harman is using 'reason' in what I have called the non-moral sense or the moral sense, the final sentence of this passage may be read in two different ways. First, suppose that he is using the term in the non-moral sense. In that case, what the sentence means is that we cannot simply say that someone's interests, desires, purposes, etc. will be served by performing some action if and only if he or she ought to perform it; and this is no doubt correct. It is a little strange to read the sentence in the other way, using the moral sense of 'reason'; but this is still intelligible. In that case, Harman is saying (in essence) that it is unsatisfactory to assert that morality obliges someone to do something if and only if that person is morally obliged to do that thing. It is somewhat strange to say this, because the assertion he discusses would then be trivially true; but (for all we have seen up to this point) perhaps Harman is merely making the point that one cannot satisfactorily answer a 'why' question by saying 'Because that's the way it is'.

The next few sentences add some clarity to this matter:

> Now, presumably, someone has a sufficient reason to do something if and only if there is warranted reasoning that person could do that would lead him or her to decide to do that thing. A naturalist will suppose that a person with a sufficient reason to do something might fail to reason in this way to such a decision only because of some sort of empirically discoverable failure. Such a failure might be inattention, lack of time, failure to consider or appreciate certain arguments, ignorance of certain available evidence, an error in reasoning, some sort of irrationality or unreasonableness, or weakness of will. If

[291] Harman 2000e, p.86.

> the person does not intend to do something and that is not because he or she has failed in some such empirically discoverable way to reason to a decision to do that thing, then according to the naturalist the person cannot have a sufficient reason to do that thing. This is the first premise in a naturalistic argument in support of the relativist.[292]

Now, what sense of 'reason' does Harman intend in the phrase "sufficient reason" here? If he means moral reasons, then this premise is certainly very controversial: it asserts precisely what anti-relativists deny. Not only that, but the conclusion he seeks to establish with this argument is "that there are people who do not have sufficient reasons, and therefore do not have sufficient moral reasons, to adhere to the general prohibition against harming or injuring others"[293]; and, more broadly, that agent relativism is true. Hence, on the 'moral reasons' reading of the above passage, Harman would be arguing in a circle. Therefore, the principle of charity demands that we read him as intending 'sufficient reason' to mean sufficient *non-moral* reason. But if this is what he means, then this first premise tells us nothing about what morality demands of such people; and the second (and only other) premise Harman offers is that "there are people, such as certain criminals, who do not act in accordance with the alleged requirement not to harm or injure others, where this is not due to any of these failings. The argument for this is simply that there clearly are people who do not adhere to the requirement in question and who do not seem to have failed in any of the ways mentioned."[294] But this premise doesn't tell us anything about the actual requirements of morality, either. On this second reading, therefore, the premises of Harman's argument in its most fully-stated form are

---

[292] *Ibid.*
[293] *Ibid.*, p.87
[294] *Ibid.*

simply irrelevant to its conclusion. Therefore, his argument fails to establish agent relativism.

Some Final Worries, and a Parallel Argument

There are, however, two points of slight weakness in the overall line of argument I have used to resist the argument from moral reasons. First, while I have offered what I take to be strong grounds for thinking that 'reasons' must mean either non-moral or moral reasons, or else some illicit equivocation between these two meanings, it is *possible* (surprising though it would be) that Harman, Streiffer, Shafer-Landau, Brink and several other metaethicists have a clear third definition of 'reasons' in mind – a definition none of them take the time to mention despite the central role the term plays in their discussions, and despite also the fact that it does not appear in the Oxford English Dictionary – and that this third definition also happens to solve the problems I have raised. Second, while I have made the case that the argument from moral reasons, in all the forms I have considered, either begs the question or commits some other fallacy, it may also be countered by agent relativists that there is no non-question-begging argument *against* agent relativism. Even if this could be shown, admittedly, it would only suffice to leave the inquirer in a state of *ataraxia* between agent relativism and its negation, by destroying the arguments on both sides. However, though even this would be sufficient for the critical task I have taken on, one might wonder whether there is not something even stronger against agent relativism.

To deal with the first worry, let us suppose that there is, in fact, a third definition (we can call it 'Definition 3') of the term 'reasons' that is a relevant and viable alternative to both the non-moral and the moral definition of 'reasons'. Let us also suppose, in the spirit of charity, that these Definition 3 reasons – D3 reasons, for short – are such that they really do make it intelligible to wonder both a) whether there can be moral reasons for X to do Y without there being D3 reasons for X to do Y (Streiffer's P1), and b) whether, for any action, there could be some agent such that there is no D3 reason for that agent to perform that action (Streiffer's P2), with the idea that a positive answer to both a) and b) would imply that agent relativism is true. If all these suppositions were correct, and if agent relativism is to have any interesting consequences (as it would not if, say, it asserted merely that different people might have very different preferences), it seems that D3 reasons would need to stand in the same sort of relationship to moral reasons that anti-relativists hold moral reasons to stand in relative to non-moral reasons. To anti-relativists, if it would be personally inconvenient to Melanie to hang up on the friend with whom she is having her fifth informal cell-phone call of the day in order to fulfil her moral duty of calling an ambulance for a severely injured person she sees, then so much the worse for her non-moral reasons: moral reasons trump them. Similarly, if the argument from moral reasons is to work, then D3 reasons must be able to trump moral reasons. If they could not, then the argument from moral reasons would not establish agent relativism, since it would only show that not everyone has D3 reasons to do what is *morally* obligatory, and not why the absence of D3 reasons to do what is obligatory is an adequate overall reason for saying that one need not do it anyway.

If such a definition exists, it would at least exonerate the argument from moral reasons on the charge of begging the question (though many other charges, not mentioned here, have also been raised against it) by making clear how one could accept both P1 and P2 of Streiffer's, Brink's, Shafer-Landau's, and Harman's arguments without equivocation and without reducing either premise to a triviality. It would also serve to illuminate many otherwise puzzling passages in the metaethical literature, such as this one from Bernard Williams:

> While we are sometimes guided by the notion that it would be the best of worlds in which morality were universally respected and all men were of a disposition to affirm it, we have in fact deep and persistent reasons to be grateful that this is not the world we have.[295]

Though this strategy might, in some way, remove the difficulties I have been discussing, I doubt very much that any definition could serve the role of D3 here in a way that makes this position plausible. There is, I think, something fundamentally wrong with the idea that one can have an overall moral obligation to do something, which is not defeated by any other moral considerations whatsoever, but which one nonetheless is not at all obligated to do and which there is no reason for one to do. Let us consider, for instance, what Williams is saying. Streiffer puts this point quite well:

> What I find puzzling about William's position is that it seems obvious to me that to the extent that these non-moral reasons do strike us as weighty enough to outweigh a moral requirement, they will also strike us as relieving us of the moral requirement with which they allegedly conflicted. Morality can, in unfortunate circumstances, require great personal sacrifice, but typically, morality itself makes room for other non-moral pursuits

[295] Williams 1981, p. 23

> and values. And when it doesn't, I for one am inclined to think that what one ought to do is abide by one's moral requirements. Morality not only makes room for other non-moral pursuits and values, it makes the right amount of room for them. So, for this reason, I think it is troubling to think that morality can require someone to do something even though there is no reason for that person to do that thing.[296]

However, I think that Streiffer does not go far enough. It is, as far as I understand the term 'moral obligation', either unintelligible or the nearest thing to unintelligible to say that "morality can require someone to do something even though there is no reason for that person to do that thing". To say that morality requires X to do Y already implies that this is true with all relevant information taken into consideration. It of course may be that someone who says that morality requires X to do Y has overlooked some important information in such a way that that judgment is not true; but this just means that *morality does not, after all, require X to do Y*. If morality *does* require X to do Y, and if there is nonetheless some reason against X's being required to do Y, then (as I and I think most people understand the term 'overall moral obligation'), X is required to do Y, regardless. As for Williams' argument, I can only make sense of it by assuming that he means by 'morality' some moral view that he does not fully accept. If someone does not accept *Christian* morality, say, I can understand the claim that we should be grateful that this sort of morality is not universally followed. However, to say that it would be much better if morality on some view were not followed is, it seems to me, tantamount to a rejection of that moral view, at least in part. Since this seems to be a fundamental fact about the proper use of the word 'morality', and hence of the phrase 'moral reasons', I fail to see how it could be possible that there could be any definition D3 of 'reasons' such that these

---

[296] Streiffer 2003, p.36

reasons trump moral reasons. At the very least, the burden of offering such a definition of 'reasons' must lie with those who claim that such a definition is not only intelligible but plausible.

As for the second final worry – that both agent relativists and anti-relativists may be condemned to beg the question against one another in their arguments that there can or cannot be moral obligations without corresponding desires, motivations, concerns, etc. in the morally obliged subject – a response is not strictly needed for my purposes, but I will offer two. The first is simply the invocation of what appears to be a good principle of philosophical argumentation: that if someone presents an claim of the form "You can't have an X without a Y", the onus is on that person to give an argument as to why one can't, rather than on the audience to demonstrate that such a thing is possible. And the second is the *prima facie* case against agent relativism I presented back at the start of this section.

While I think that all this is more than adequate as a response to Harman's argument from moral reasons, I think there is a perhaps even stronger reply: a parallel argument.

The parallel argument this time is between moral reasons and epistemic reasons. We begin with the argument from moral reasons. Here, again, is Streiffer's version:

> **P1**: For any action, if an agent is morally required to perform that action, then there is a reason for that agent to perform that action.

P2: For any action, there [could be] some agent such that there is no reason for that agent to perform that action.

Conclusion: For any action, there [could be] some agent who is not morally required to perform that action.

And now, here is the parallel 'argument from epistemic reasons':

P1: For any belief, if a subject is epistemically required to reject that belief, then there is a reason for that subject to reject that belief.

P2: For any belief, there [could be] some subject such that there is no reason for that subject to reject that belief.

Conclusion: For any belief, there [could be] some subject who is not epistemically required to reject that belief.

Let us take the belief in question to be that there are more people living in the Queen Charlotte Islands than in Vancouver, Victoria, and all the rest of British Columbia combined. By P1 of the argument from epistemic reasons, if a subject is epistemically required to reject that belief about the Queen Charlotte Islands, then there is a reason for that subject to reject it. Now, my old roommate Rob was a strong and vocal proponent of just this belief, and nothing in his belief system and principles of reasoning seemed sufficient for him to abandon that belief even when presented with what would normally

be considered excellent evidence.[297] Also, nothing non-trivial follows from the argument from epistemic reasons unless we interpret 'reasons' along the lines of what I called the non-moral definition of reasons in the discussion of the argument from moral reasons. Putting this together, it follows by *modus tollens* from P1 and this fact about Rob that Rob was not epistemically required to reject his belief about the Queen Charlotte Islands. Given the ridiculous beliefs that some people (at least *one* person, and at the very least some *possible* person) maintain in the face of extremely powerful reasoning to the contrary, even when they are apparently fully aware of that contrary reasoning, the move from P1 to the conclusion is quite straightforward. Hence, once we have accepted P1, we seem committed to the conclusion that all it takes to be justified in any belief is to rid your belief system and set of reasoning practices of anything that might be used against your maintaining that belief. But surely, this is absurd. Rob was not warranted in maintaining his belief about the Queen Charlotte Islands, no matter how much he believed that he was. But if we have good grounds for rejecting the argument from epistemic reasons, it seems that by the same logic we need to reject the argument from moral reasons. Russ Shafer-Landau makes this point very nicely, and is worth quoting at length on this issue:

> [I]f reasons exist regardless of one's beliefs, desires or interests, then where do they come from? (And how can we know them?) What, other than an agent's own perspective, could serve as a source of normative authority? To insist that a set of facts could contain within themselves normative authority for agents, regardless of their outlooks on life, seems obscurantist, and gives the appearance of prematurely cutting off any helpful explanation of normativity.

---

[297] This is not, incidentally, something absurd I have made up for the sake of argument: I really did have such a roommate, and he really did steadfastly maintain this belief in the face of what I would have taken to

If this is obscurantist, I think we have no choice but to embrace the mysteries. I think that intrinsic normativity is ineliminable. To see this, consider the parallels between conditions of epistemic and moral assessment. We say that agents, if they have reason to believe anything at all, have reasons to believe the truth, and to conform their reasoning to truth-preserving schemas, even if believing the truth is not conducive to the goals they set themselves. Suppose someone, on the basis of excellent evidence, accepts the truth of a conditional and its antecedent, but denies that she has any reason to accept the consequent. It's not just that she may have (possibly overriding) reasons which oppose making such an inference. Someone might correctly believe that all passengers aboard a drowned airliner have been killed, while knowing that her brother was among them, and yet resist drawing the terrible conclusion. Practical considerations, such as sustaining emotional stability, may militate against believing the truth, and may, for anything said thus far, be so strong as to outweigh it in given instances. But the sister who holds out hope against all evidence, and contrary to the logical implications of her own beliefs, is in some real sense acting against reason. Indeed, she is, in one sense, behaving irrationally, though also in a way that is fully understandable. She is acting contrary to sufficiently good reasons – reasons are there to tell her, and anyone in her epistemic situation, that she ought to believe something that she cannot bring herself to believe.

To say such a thing commits one to the existence of what I shall call *non-perspectival* or *intrinsic* reasons. I believe that there is intrinsic reason to think that two and two are four – the fact itself provides one with reason to believe it. One needn't show that such belief is somehow related to one's adopted goals in order to justify believing such a thing. If, unusually, success at basic mathematics was entirely unrelated to one's preferred activities, one would still have good reason to think that two and two were four, not five or three.[298]

---

be extremely persuasive evidence to the contrary.

[298] Shafer-Landau 2003b, pp.205-206. In the very recent past, a number of other metaethicists have begun to draw similar parallels between moral and epistemic normativity. See, for instance, Matthew S. Bedke, 'Might All Normativity Be Queer?' (Bedke, forthcoming) and Sharon Street, 'Evolution and the Normativity of Epistemic Reasons' (Street, forthcoming).

Further, let us assume for a moment that we are faced with some diehard agent relativists (in morality) who, since they concede that they cannot produce any good reasons for both accepting the argument from moral reasons *and* rejecting the argument from epistemic reasons, opt to bite the bullet and accept both arguments, rather than reject them. Such diehards will, therefore, endorse the conclusion of the argument from epistemic reasons, which is simply epistemic relativism: 'For any belief, there is some subject who is not epistemically obliged to reject that belief.' The subject referred to here is, presumably, any subject whose personal set of beliefs and reasoning practices involves the endorsement of (or at least ambivalence about) the belief in question. Now, since this is true of any belief, it is true of belief in the conjunction of moral anti-relativism (of both the appraiser and agent varieties) and epistemic anti-relativism. And those who are anti-relativists about morality and epistemology seem to qualify as the subjects in question. So, by the diehard relativist's own lights, anti-relativists are entirely warranted in believing that agent relativism is false (or at least being ambivalent about it) and, further (if desired), that all agent relativists are (or could easily be) objectively mistaken, and that the matter is closed. This is, of course, the familiar 'recoil' argument against epistemic relativism, introduced via the parallel arguments approach into the agent relativism discussion.

What more could be wanted than this? Why not, on the strength of this case, reject the agnosticism of the parallel arguments approach entirely, and join Shafer-Landau, Bedke, Street, and company in an endorsement of moral realism? I will address this final, but important, point in my Conclusion.

Conclusion: Moral Relativism and the 'Profound' Philosopher

If my reasoning so far has been sound, persuasive, and complete, I have by now established for those in the humour of Hume's practical philosopher that Harman's case for relativism is not successful, even with the benefit of liberal applications of the principle of charity. Among other problems with Harman's case, we have seen that accepting any of his arguments for moral relativism would entail our also accepting, by the same logic, arguments for skepticism[299] about science (Chapter Two) or for epistemic relativism (Chapter Three). But the conclusions of these latter arguments would be absurd for us to accept. The central commitments of skepticism about science, in the form that the parallel arguments to Harman's would require, would certainly be liable to something like Hume's reply to general skepticism:

> These principles may flourish and triumph in the schools; where it is, indeed, difficult, if not impossible, to refute them. But as soon as they leave the shade, and by the presence of the real objects, which actuate our passions and sentiments, are put in opposition to the more powerful principles of our nature, they vanish like smoke, and leave the most determined sceptic in the same condition as other mortals.[300]

---

[299] It will be remembered that Harman's famous argument from the distinction between seeing a proton and seeing the wrongness of burning a cat is an argument for moral *skepticism*, and not moral relativism. Harman only moves from this argument to the conclusion that moral relativism is correct by suggesting that an acceptance of moral relativism is the most reasonable response to the initial case for moral skepticism (see Chapter One for more on this). Hence, again, the parallel arguments discussed in Chapter Two are for scientific and epistemic skepticism, not relativism.

[300] Hume 1993, p.116

Radical epistemic relativism is a position the practical philosopher will be even less worried about: to hold that anyone's standards of epistemic justification are as good as anyone else's, is to deny that one is following what are objectively the epistemic practices in concluding that epistemic relativism is true. But this seems to imply that one has no objectively adequate reason for accepting epistemic relativism, and hence that one is doing nothing objectively wrong in rejecting that view. And, since the strongest arguments Harman offers for moral relativism are no better than this, the parallel arguments can also be rejected by the practical philosopher, who need not accept anything so crass as Jamesian pragmatism[301] to reach this conclusion.

But a profound philosopher will rightly be left unsatisfied by this. Even if one can never be objectively justified in *holding* that radical epistemic skepticism, say, is true, it might nonetheless be the case that it *is* true. Barry Stroud makes a similar point in discussing the 'repugnancy' Berkeley finds in supposing that one can think of an extended body that is not coloured:

> Another feat you cannot perform is to say that you are talking without talking. It is not possible for you to say that you are not speaking and thereby say something true. There is in that sense a "repugnancy" involved in saying truly that you are not speaking: it simply cannot be done. But there is no "repugnancy" or impossibility involved in what you are saying about yourself when you say that you are not speaking. It is not

[301] "It is as if James is *objectifying* belief. He treats it as one might treat an ornament, for which the only questions would be: does this suit me; is it a good thing to wear to the social party? … A related charge is that James is *privatizing* belief, concentrating not upon the social trust that is at the forefront of Clifford's discussion, but upon the private satisfactions that follow upon settling a matter in one's own mind." (Blackburn 2005, pp.8-9)

> impossible for it to be true that you are not speaking. It is true of each of us much of the time.[302]

The parallel arguments approach I adopt forces a choice between accepting relativism (or skepticism) about both moral *and* epistemic standards, etc., or the rejection of both. For the practical philosopher, the answer is obvious: global relativism or skepticism must be rejected. For the profound philosopher, however, the parallel arguments approach is not so decisive.

So how do we proceed with the profound philosopher? The position I find it best to adopt in that context is one of limited agnosticism. I will spend the rest of this Conclusion spelling out what this involves and why it seems best.

A Limited Agnosticism

There is a substantial philosophical trend in early modern philosophy – from at least as far back as the publication of Locke's *Essay* in 1690 – that combines a skeptical position on all or most metaphysical (or, more broadly, philosophical) issues with a non-skeptical position on everyday questions. Some philosophers see evidence for this sophisticated version of skepticism in classical antiquity, and have therefore called it 'Pyrrhonian scepticism'.[303] A similar move has recently been made along relativistic lines: Steven D. Hales, in his *Relativism and the Foundations of Philosophy*,[304] argues (reluctantly)[305] that there is no good basis for rejecting epistemic relativism about philosophical claims, so

---

[302] Stroud 2002, p.67
[303] See Sinnott-Armstrong 2004 for an excellent anthology on this subject.
[304] Hales, 2006
[305] *Ibid.*, p. 184

long as that relativism is carefully distinguished from relativism about everyday matters. I know of no good argument that cleanly disposes of either of these worrisome views, which makes me hesitant in espousing a strong philosophical position against either.

Epistemic difficulties in establishing anything secure in philosophy of the profound sort are only part of the problem, however: there are also *linguistic* difficulties. James Harold has recently argued[306] that our folk-psychological term 'belief' is too crude to allow us to hope for a decisive result in the debate between ethical cognitivism (the view that our ethical commitments are best understood as beliefs) and non-cognitivism (the view that they are best understood as something else). According to Harold, the level of analytical detail we have already achieved in our effort to resolve this problem has already become much more precise than anything our somewhat unclear term 'belief' will bear. Russ Shafer-Landau, while arguing for non-naturalism (the view that moral properties are real properties that are neither identical with nor reducible to the properties studied by natural or social scientists) and rejecting non-reductive naturalism (the view that moral properties are somehow a part of, or continuous with, the properties studied by natural or social scientists, while not being reducible to any other natural properties) has the following to say:

> ...we need to ask just how much of the debate has become a matter of taxonomic bookkeeping, and how much of philosophical significance is really at stake. What matters is not what we call a view, but whether the view, whatever its denomination, can solve the central philosophical problems that have generated the division between naturalists and non-naturalists.

---

[306] Harold, 2007

> Historically, and on some contemporary views, this division is worth fighting for and signifies a vital theoretical boundary. However, when we identify the stakes at hand, we can see, pretty quickly, that there is very largely a pattern of shared strength or vulnerability, and often little else worth marking by this time-honoured division.[307]

I suggested in the last chapter that the term 'having a reason' might similarly lack the precision that is needed to make it clear what it would take for someone to have a reason to do something. Unless this can be sorted out clearly, I hold that it would be premature to commit oneself to a position on the reasons externalism issue, say (that is, the issue of whether one can have a reason to do something that one cannot be motivated to do – and of course, what counts as being such a person and what counts as being in such a situation that one *cannot* be motivated to do something, also need fleshing out). Concern over these and other ambiguities in our terminology, and recognition (as in the above quote from Shafer-Landau) that these ambiguities might be serious enough that "time honoured divisions" marked out by even the most precise definitions of these terms that we have been able to produce might in the end come to little or nothing, are not uncommon in the contemporary metaethical literature.

It may be argued that these linguistic issues might in the end be resolved (or even that they have been!), and similarly for the epistemological issues I mentioned earlier. Very well: I do not deny the possibility. But let us consider what it would take to clear up these issues to the point where a powerful argument for relativism – on the *profound* view of philosophy -- could successfully be made. First, all the relevant metaethical

terminology would need to be defined clearly and unambiguously, rather than metaphorically: no more would it be acceptable for J. L. Mackie to ask whether values are "part of the fabric of the world"[308]. Second, any ambiguities in the components of the definitions of metaethical terms would also need to be resolved. Even if we accept, with Geoffrey Sayre-McCord, that the ambiguities between "the many moral realisms" can be dealt with if we limit the definition of 'realism' to "just two theses: (1) the claims in question, when literally construed, are literally true or false (cognitivism), and (2) some are literally true",[309] other ambiguities remain. What, exactly, is meant by 'literally construed'? And what does it take for something to be 'literally true'? These terms will have to be sorted out unambiguously before one can make real progress along this path, if one is proceeding in the humour of the profound philosopher. Third, it will be necessary to make equally clear how all relevantly similar metaethical terms differ in their definitions. What, exactly, is the precise relationship between saying that X has a reason to do Y, that there is a reason for X to do Y, that X ought to do Y, that it ought to be the case that X does Y, that it is demanded by morality that X do Y, and so on? While some of these and other relations have been spelled out by various philosophers, these spellings-out have not always (or perhaps ever) been as precise or complete as the project of profound philosophy requires. Fourth, it will need to be made clear just what will count as an adequate demonstration that moral relativism (however defined) is the correct view, or that it has been refuted. It will, as I have suggested, probably be considered an inadequate refutation of moral relativism by the standards of profound philosophy to

---

[307] Shafer-Landau 2003b, p.62. The Stanford Encyclopedia of Philosophy entry on moral non-naturalism begins with a similar point: see Ridge, 2008.
[308] Mackie 1991, pp. 15, 21, 22, 23 and 24
[309] Sayre-McCord, 1988b, p.5

show that we have no better reasons for accepting it than we do epistemic relativism; but then, what *should* count as successful in that arena? And fifth, it will need to be shown (against the more sophisticated skeptics and relativists I discussed earlier, who limit their reservations about objective knowledge to philosophical matters) that we have adequate grounds for accepting that (say) moral relativism is objectively correct.

In setting all this out, I should make clear, I am not committing myself to the position that these tasks cannot or will not be performed. I am merely drawing attention to what the presentation of any adequate case for a metaethical position would be like on a committed 'profound' approach. As for Harman's arguments for relativism, I admit the possibility that someone might modify them in such a way that they could be successfully deployed along the lines set out above. This, as I see it, would be a two-stage process. In the first stage, one of Harman's arguments would need to be rebuilt in such a way that the non-parallel-arguments objections I have presented in the preceding chapters would no longer be effective against them. And in the second stage, all the terminology used in that argument, and the implicit criteria for that argument's success, would have to be clarified and defended following the five points in the previous paragraph.

If such a thing were achieved, then it would be most relevant to ask which form of relativism is established by that successful argument. If the form of relativism established is closely akin to what I termed 'extremely boring moral relativism' in the introduction[310], then I see no inconsistency in accepting that form of moral relativism on the profound approach and yet rejecting moral relativism on the practical approach: the

forms of relativism a practical philosopher will discuss, or consider worthy of the name, will be at least somewhat interesting, after all. It may also be, however (though I suspect this is much less likely) that a defender of Harman is able to present a sound argument for an *interesting* form of moral relativism after the manner of the profound philosopher. In that case, my overall position of limited agnosticism on this issue, as set out below, would have to be abandoned. But I take it that this possibility is remote.

Still, it might be thought that the extremely difficult conditions the profound approach to philosophy entails cuts both ways. Isn't it also a distant possibility that anyone could offer an adequate argument *against* relativism by these profound standards? Can I really think myself to have presented such an argument in these pages? And if I haven't, then how can I be justified (on the profound approach) in maintaining that moral relativism is false?

My answer is that I do not maintain the falsity of moral relativism on the profound approach. I maintain, as I argued in the three main chapters of this work, that we have excellent *pragmatic* reasons for rejecting moral relativism, and as far as I can see we have no good pragmatic reasons for accepting it. And this is true even if one rejects what I take to be the crass version of pragmatism espoused by James. However, I feel that the best approach to take on any issue in profound philosophy is a limited agnosticism (in this case, limited only by what I take to be slight intuitive grounds for thinking that moral relativism is false) combined with a reminder to the side making the claim that it bears the burden of proof. Too often, I feel, anti-relativists have fared poorly against relativists

---

[310] p.15

because they feel that an anti-relativist position is not complete unless it is accompanied by a positive position on such issues as I have described above (cognitivism vs. noncognitivism, naturalism vs. nonnaturalism, reasons internalism vs. reasons externalism, etc.). If there is (as I suspect) something to the worry that we have not yet laid down an adequate meta-metaethical foundation for stable and enduring arguments and positions in profound metaethics, this temptation to include a positive component in any response to relativism is misguided, and the difficulties pointed out by the interlocutor for the particular non-relativist position espoused and the arguments for it will lead to a general impression that relativism is faring far better than it actually is. So I avoid taking on any such positive commitments.

The suggestion that progress in some area of ethics needs to wait (perhaps forever!) for the construction of an adequate foundation is not new: early and mid- 20th century metaethicists were fond of arguing that work on applied ethics should be stopped until the theoretical foundations of morality were established.[311] While the practical need to resolve problems in applied ethics makes this suggestion impossible for all intents and purposes, the same does not seem to hold for metaethics, so long as the metaethical issues in question are limited to the fully profound. There is plenty of time for us to wait for those who wish to take on the burden of arguing for relativism in the profound but interesting sense to make their case adequately.

One does not need to wait, however, to see that Harman's case as presented is not successful. If my analysis has been correct, I have now shown that all of his arguments,

as offered, run into general problems of various kinds; that, so far as the practical philosopher is concerned, all of his main arguments fail because they commit us to absurd parallel arguments; and that, so far as the *profound* philosopher is concerned, none of his arguments (even when charitably reinterpreted) comes close to meeting the high standards of success in making a case for moral relativism. It is by showing this that I hope to have answered Harman's relativism.

---

[311] See Peter Singer's introduction to Singer (ed.) 1986 for more on this.

Bibliography

Alexander, Larry, Moore, Michael, "Deontological Ethics", *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/ethics-deontological/>.

Ayer, A. J., 1952. *Language, Truth and Logic*. Dover

Beardsley, Monroe C., 1965. 'Intrinsic Value'. *Philosophy and Phenomenological Research,* Vol. 26, No. 1 (Sep., 1965) pp. 1-17

Bedke, Matthew S., forthcoming. 'Might All Normativity Be Queer?'

Blackburn, Simon, 1984. *Spreading the Word*. Oxford University Press

Blackburn, Simon 1993. *Essays in Quasi-Realism*. Oxford

Blackburn, Simon, 2005. *Truth: A Guide*. Oxford University Press

Brink, David O., 1989. *Moral Realism and the Foundations of Ethics*. Cambridge University Press

Chalmers, David, 1997. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press

Cioffi, Frank, 1999. *Freud and the Question of Pseudoscience*. Open Court

Clifford, William; James, William; and Burger, A. J. 2001. *The Ethics of Belief*. Dry Bones Press

Crews, Frederick C. (ed.), 1998. *Unauthorized Freud: Doubters Confront a Legend*. Viking Penguin

Darwall, Stephen; Gibbard, Allan; and Railton, Peter, 1992. 'Toward Fin de siecle Ethics: Some Trends'. *The Philosophical Review*, Vol. 101, No.1, Philosophy in Review: Essays on Contemporary Philosophy (Jan., 1992), pp. 115-189

Dawkins, Richard, 1996. *The Blind Watchmaker*. W. W. Norton

Denker, John, 2005. 'The Laws of Thermodynamics'. URL = <http://www.av8n.com/physics/thermo-laws.htm#heat-hotness>

Double, Richard, 1996. *Metaphilosophy and Free Will*. Oxford University Press

Driver, Julia, 2007. *Ethics: The Fundamentals*. Blackwell Publishing

Elliott, Anthony (ed.), 1998. *Freud 2000*. Routledge

Erwin, Edward, 1995. *A Final Accounting: Philosophical and Empirical Issues in Freudian Psychology*. The MIT Press

Feinberg, Joel (ed.) and Shafer-Landau, Russ (ed.) 2007. *Reason and Responsibility: Readings in Some Basic Problems in Philosophy*. Wadsworth

Ferraiolo, William, 2002. 'A Dilemma for Robust Alethic Relativism'. Sorites, Issue 14, October 2002, pp.36-41

Frankena, William. K. 1939. 'The Naturalistic Fallacy'. *Mind*, 48: pp. 464-477.

Galileo, Galilee 1991 (1638). *Dialogues Concerning Two New Sciences*. Prometheus

Garner, Richard 1967. 'Beardsley, Firth and the Ideal Observer Theory'. *Philosophy and Phenomenological Research*, Vol. 27, No. 4 (Jun., 1967), pp.618-623

Gibbard, Alan 2003. *Thinking How To Live*. Harvard

Gibbard, Alan 2008. *Reconciling Our Aims: In Search of Bases for Ethics (Berkeley Tanner Lectures)* Oxford

Gowans, Chris, "Moral Relativism", *The Stanford Encyclopedia of Philosophy (Winter 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2008/entries/moral-relativism/>.

Graham, Gordon 1996. Review (untitled) of *Moral Relativism and Moral Objectivity* by Gilbert Harman and Judith Jarvis Thomson. *Philosophy*, Vol. 71, No. 278 (Oct., 1996), pp. 622-624

Hales, Steven D. 1997. 'A Consistent Relativism'. *Mind* 106, pp.33-52

Hales, Steven D, 2006. *Relativism and the Foundations of Philosophy*. MIT Press

Hardin, C. L. 1988. *Color for Philosophers: Unweaving the Rainbow*. Hackett

Harman, Gilbert 1975. 'Moral Relativism Defended'. Philosophical Review 84, pp.3-22

Harman, Gilbert 1977. *The Nature of Morality*. Oxford University Press

Harman, Gilbert 1986. 'Moral Explanations of Natural Facts: Can Moral Claims be Tested Against Moral Reality?' *The Southern Journal of Philosophy* (1986) Vol. XXIV, Supplement, pp.57-67

Harman, Gilbert 1999. 'Rationality'. Published in *Reasoning, Meaning and Mind*, Oxford University Press

Harman, Gilbert 2000. *Explaining Value and Other Essays in Moral Philosophy.* Oxford

Harman, Gilbert 2000a. 'Moral Relativism Defended'. Published in Harman 2000.

Harman, Gilbert 2000b. 'What is Moral Relativism?' Published in Harman 2000.

Harman, Gilbert 2000c. 'Relativistic Ethics: Morality as Politics'. Published in Harman 2000.

Harman, Gilbert 2000d. 'Justice and Moral Bargaining'. Published in Harman 2000.

Harman, Gilbert 2000e. 'Is There a Single True Morality?'. Published in Harman 2000.

Harman, Gilbert 2000f. 'Moral Philosophy and Linguistics'. Published in Harman 2000.

Harman, Gilbert 2008a. 'Naturalism in Moral Philosophy'. Draft of talk delivered at Kansas State University, April 2008; unpublished.

Harman, Gilbert 2008b. 'Using a Linguistic Analogy to Study Morality.' Published in Sinnott-Armstrong 2008, pp.345-351.

Harman, Gilbert and Thomson, Judith Jarvis 1996. *Moral Relativism and Moral Objectivity.* Blackwell Publishing

Harold, James, 2007. 'Meta-metaethics: Moderate Skepticism about Some Concepts of Metaethical Inquiry. *Pacific Division Meeting of the American Philosophical Association.* San Francisco, CA. April 2007.

Hobbes, Thomas, 1994 (1651). *Leviathan.* Trans. Edwin Curley. Hackett.

Hume, David, 1993 (1748). *An Enquiry Concerning Human Understanding.* Hackett.

Hume, David, 2000 (1739-40). *A Treatise of Human Nature.* Oxford Philosophical Texts

James, William, 1981 (1907). *Pragmatism: A New Name for Some Old Ways of Thinking.* Hackett

Johnson, Kristin, 2006. *Rescuers and Others.* URL = <http://www.utahcrossroads.org/DonnerParty/Rescuers.htm>

Kant, Immanuel, 1998. *Groundwork of the Metaphysics of Morals.* Cambridge University Press

Kemp Smith, Norman 2005 (1941). *The Philosophy of David Hume*. Palgrave Macmillan

Kim, Jaegwon, 1984. 'Epiphenomenal and Supervenient Causation'. *Midwest Studies in Philosophy*, Vol. 9, 1984, pp.257-70

Kolbel, Max, 2007. 'An Argument for Relativism'. *Think*, no. 14

Koons, Jeremy Randal, 2000. 'Do Normative Facts Need to Explain?' From *Pacific Philosophical Quarterly* 81:3 (September 2000), pp.246-272. Blackwell Publishers Ltd.

Kuklick, Bruce, 1981. Editor's introduction to James, William 1981.

Locke, John 2000 (1689). *An Essay Concerning Human Understanding*. Adamant Media Corporation

Locke, John 1993 (1689). *Two Treatises of Government*. Everyman

Mackie, J. L., 1977. *Ethics: Inventing Right and Wrong*. Pelican Books

Macmillan, Malcolm, 1997. *Freud Evaluated*. The MIT Press.

Matlin, Margaret, 1998. *Psychology*, 3rd Edition. Wadsworth.

McLaughlin, Brian and Bennett, Karen, 2005. 'Supervenience', *Stanford Encyclopedia of Philosophy* (http://plato.stanford.edu/entries/supervenience)

Midgley, Mary 1993. *Can't We Make Moral Judgments?* Palgrave Macmillan

Miller, Alexander 2003. *An Introduction to Contemporary Metaethics*. Polity Press

Nobis, Nathan 2005. *Truth in Ethics and Epistemology: A Defense of Normative Realism*. URL = <http://nathannobis.com/papers/dissertation/intro.html>

Plato, 1990. *The Theaetetus of Plato* (Ed. Burnyeat, Myles; Trans. Levett, M. J.). Hackett

Popper, Karl 1965. *Conjectures and Refutations*. Harper & Row

Pust, Joel 2000. *Intuitions as Evidence (Studies in Philosophy)*. Routledge

Quinn, Warren S., 1986. 'Truth and Explanation in Ethics'. *Ethics*, Vol. 96, No. 3 (April 1986), pp. 524-544. University of Chicago Press.

Rachels, James 2006. *The Elements of Moral Philosophy*. McGraw-Hill

Rawls, John 1971. *A Theory of Justice*. The Belknap Press

Ridge, Michael, "Moral Non-Naturalism", *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/moral-non-naturalism/>.

Rorty, Richard, 1981. *Philosophy and the Mirror of Nature*. Princeton University Press

Sayre-McCord, Geoffrey (ed.) 1988. *Essays on Moral Realism*. Cornell University Press

Sayre-McCord, Geoffrey, 1988a. 'Moral Theory and Explanatory Impotence'. Published in Sayre-McCord 1988

Schroeder, Mark 2007. *Slaves of the Passions*. Oxford

Schroeder, Mark 2008. *Being For: Evaluating the Semantic Program of Expressivism.* Oxford

Schroeder, Mark 2009. *Noncognitivism in Ethics*. Routledge/Taylor and Francis

Shafer-Landau, Russ 2003a. *Whatever Happened to Good and Evil?* Oxford University Press

Shafer-Landau, Russ 2003b. *Moral Realism: A Defence.* Oxford University Press

Sher, George, 1980. 'Moral Relativism Defended?' *Mind*, New Series, Vol. 89, No. 356 (Oct., 1980), pp. 589-594

Sinnott-Armstrong, Walter (ed.) 2004. *Pyrrhonian Skepticism*. Oxford

Sinnott-Armstrong, Walter (ed.) 2008. *Moral Psychology, Volume 1, The Evolution of Morality.* MIT Press, 2008

Skyrms, Brian 2004. *The Stag Hunt and the Evolution of Social Structure*. Cambridge

Smith, Michael 1994. *The Moral Problem.* Wiley-Blackwell

Street, Sharon, forthcoming. 'Evolution and the Normativity of Epistemic Reasons'

Streiffer, Robert, 2003. *Moral Relativism and Reasons for Action.* Routledge

Stroud, Barry, 2002. *The Quest for Reality: Subjectivism and the Metaphysics of Color.* Oxford

Sturgeon, Nicholas L., 1980. Review (untitled) of *Hume's Moral Epistemology* by Jonathan Harrison. *The Philosophical Review*, Vol. 89, No. 1 (Jan., 1980), pp. 124-129

Sturgeon, Nicholas L., 1985. 'Moral Explanations'. Published in Sayre-McCord 1988.

Sturgeon, Nicholas L., 1986a. 'Harman on Moral Explanations of Natural Facts'. *Southern Journal of Philosophy* 24, suppl., 69-78

Sturgeon, Nicholas 1986b. 'What Difference Does it Make Whether Moral Realism is True?' *Southern Journal of Philosophy* 24, suppl., 115-142

Sturgeon, Nicholas 1994. 'Moral Disagreement and Moral Relativism'. *Social Philosophy and Policy* (Winter 1994), 11(1):80-115

Williams, Bernard 1981a. 'Internal and External Reasons' in *Moral Luck*. Cambridge University Press, pp. 101-13.

Williams, Bernard 1981b. 'Moral Luck' in *Moral Luck*. Cambridge

Wollheim, Richard (ed.) 1974. *Freud: A Collection of Critical Essays*. Anchor Books