Moral Relativism and Moral Universalism: A Coherentist Approach

by

Damien Anthony Cole

A thesis
submitted to the Victoria University of Wellington
in fulfillment of the
requirements for the degree of
Doctor of Philosophy
in Philosophy

Victoria University of Wellington
2004

Abstract

In the following work I examine the long-standing dispute between the moral relativist and their opponent, who I call the moral universalist. I examine this dispute from a coherentist perspective. For both moral relativism and moral universalism I identify a key conviction underlying the position and I attempt to deliver a theory that accommodates both of these two convictions in a coherent fashion. This involves distinguishing between distinct types of moral principle and showing how the relativist's conviction is a reaction to one type while the universalist's conviction is a reaction to another type. These types of moral principle are not mutually exclusive, however, and a large part of my project will be to explain how each type of principle is able to be included as part of a coherent whole.

Acknowledgements

Table of Contents

Part One
Introduction

1.1 An important question.

For almost as long as there have been philosophical questions, one of the more contentious has been the following:

Q1) Is morality relative?

In some moods this can appear to be one of the most significant philosophical questions that may be asked about morality. Producing an answer to Q1 might be thought to be logically prior to the discussion of ordinary everyday ethical questions and being able to give a plausible answer can be seen as a precondition for being able to say what we ought and ought not to do. This is because if Q1 is answered in the affirmative then our everyday ethical discussions may have to proceed along rather different lines from those along which they proceed if Q1 is answered in the negative. Since a major goal of our everyday ethical discussions is to produce moral judgements telling us how we ought to conduct ourselves, the judgements we produce will to a greater or lesser degree depend on how we answer Q1.

The implications of the disagreement about the correct way to answer this question extend beyond the merely theoretical and affect our thinking about how to live our day to day lives. The importance in a pragmatic sense of finding a plausible answer to Q1 should not be overlooked. Perhaps this is one of the reasons why moral relativism is, even among the non-philosophical public, quite a widely known (although not necessarily widely understood) philosophical position.

My own experience of this comes from trying to run reasonably structured discussions of controversial moral issues whilst a tutor in applied ethics courses. Inevitably, at some point during the course a student would bring up, with varying degrees of sophistication, a view that could be called some variety of moral relativism. The student might suggest that the moral discussion that was taking place was a pointless debate that could never be resolved, because: "Each side is right from their own point of view", or a comment along similar lines to this. Sometimes this sort of comment was deliberately used as an attempt at an 'argument-stopper', particularly when disagreement was intense.

To agree that what the student had said rendered the debate pointless would imply that there was little to be gained from further participation in the discussion, not a useful result when participation in the discussion was the purpose of the tutorial. However, to disagree would mean the discussion would have to become more self-consciously metaethical, so that the merits of possible answers to a question like Q1 could be debated. Generally I would attempt to neutralise the objection so that the moral discussion could continue.

The possibility of moral relativism must almost always be dealt with at some point when running tutorials in an applied ethics course and there were a number of ploys I used to defuse this objection. Perhaps the student in question had expressed and argued for a certain moral view on one of the previous controversial issues in the course and was at that earlier time not prepared to accept that the views of those who disagreed might be the correct moral viewpoint for them. Pointing out the inconsistency here would often be enough to make the student forego their moral relativism, as they were more firmly committed to their moral stance on the previous issue. Occasionally the student would instead choose to relinquish their previously held moral stance in favour of a commitment to moral relativism, or perhaps the student had not previously taken any strong stance on a moral issue. At that stage I generally preferred to suggest that for the remainder of the course we would be assuming for the sake of argument that moral relativism was not the case.

Furthermore, in order to avoid derailing the topic of the tutorial, that was unfortunately all that could be said about the matter.

While a response such as this was sufficient to bring the discussion back onto the topic, it is not the most intellectually satisfying of responses and the neophyte moral relativist who is thus rebuffed may complain with some justice that their point has not been addressed. Of course, making an assumption in order to facilitate discussion is not something that is peculiar to moral discussion in applied ethics tutorials; there are many sorts of philosophical discussions for which fundamental assumptions, sometimes in themselves quite philosophically controversial assumptions, must be made before the discussion can even get off the ground. For instance, a paper in philosophy of mind may begin with the declaration that the author assumes that the so-called 'problem of other minds' has already been adequately addressed, so that a certain proposed solution to the problem will be taken as given.

Even making this observation, however, is to concede the point that a question like Q1 and the issues that it raises must be acknowledged and somehow dealt with on the way to a fully philosophically satisfactory excursion into questions in applied ethics. Some attempt at an adequate answer to Q1 must be on hand if we want to take our moral discussions with the seriousness they warrant.

We do want to take our moral discussions seriously because of their close connection to our actions. Morality advises us about what we ought to do, about the actions we ought to take, and so the theoretical foundations for such advice will ideally be as firm as possible, if our actions are not to be arbitrary and subject to capricious whim. Anyone, but especially a philosopher, who wishes to offer ethical instruction to others must be able to face and give a response to a question like Q1, and to give reasons why they answer it as they do, before they can have much credibility when venturing moral advice.

Yet despite the apparent desirability of coming up with a plausible answer to Q1, the question has remained recalcitrant. Many would answer the

question, even as it is posed above, with an emphatic 'No!' but there have always been those who answer 'Yes', and in the absence of any clear consensus here the question continues to be a live one. The unresolved nature of the question has the unfortunate result of lending an air of fatuity to our everyday ethical discussions – if we cannot settle Q1 prior to these discussions then it seems we must admit to some possible uncertainty about the status of our ordinary moral judgements and it may fairly be wondered what business we have making these moral judgements with any conviction. Do we have to put our moral lives on hold until some exceptional philosopher answers Q1 once and for all? I think not, and will try to show why below.

1.2 Clarifying the question

Stated in the form it is above, Q1 is vague and its significance is unclear. The question seems to be incomplete without some indication of what morality might be supposed to be relative to. Before a serious attempt to answer Q1 can be begun the question itself must be examined and made more precise. Part of the goal of this first chapter will be to do this, and also to raise some preliminary points that must be dealt with before any effective treatment of Q1 can begin. The question is further complicated by the way in which it is often entangled with other issues that are not directly relevant to it. One of the major themes of this study will be the necessity of preventing the investigation becoming sidetracked into other metaethical issues. In this chapter and also throughout the work I will continue to try to clear away those concerns that are peripheral to the main topic, and maintain the focus on matters that are central to the question.

As a first stab at the philosophical notions that are at work when the proposal is made that morality is relative, consider the ancient Greek practice of keeping human slaves. To those living in Western democracies in the early 21st century it is uncontroversial that the practice of human slavery is morally wrong, but our view about this seems on the face of it to be different to the view of those living in ancient Greece in the 5th and 4th centuries B.C.E. When

someone answers 'Yes' to Q1 they are saying something like the following: the ancient Greek view of the moral status of human slavery actually is the morally correct view for the ancient Greeks, but not for us, and our apparently incompatible view of the moral status of human slavery actually is the morally correct view for us, but not for the ancient Greeks. It is not just a case of 'When in Athens, do as the Athenians do', but that what the Athenians do in Athens is what is right for Athenians to do whilst they are in Athens.

In Plato's dialogue *Theaetetus* the philosopher Protagoras is reported to have said that: "whatever any community decides to be just and right, and establishes as such, actually is what is just and right for that community and for as long as it remains so established." (177d). This looks like a philosophical claim similar to the one made in the previous paragraph. The kind of moral relativism advocated here is that what is morally right is relative to each community, so that a claim about the morally just and right way in which to behave must contain a reference to the community in which this behavior is to take place.

Nowadays we find contemporary moral relativists saying much the same sort of thing as Protagoras may have said. Gilbert Harman's recent characterisation of his version of moral relativism is a good example of this:

> (M)oral 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. (1996, p 3)

In this quotation Harman's abstract idea of a framework has replaced the more prosaic notion of a community, but the sort of position he is offering seems not dissimilar from the one attributed to Protagoras. Harman's position will be examined in detail in Part Two, but for now his framework metaphor provides us with one possible way of answering Q1 in the affirmative, by stating what it is that morality is relative to:

10

A1) Morality is relative to a moral framework.

However, it seems to be inevitable that different moral relativists will have different views about the best way to answer Q1 in the affirmative. For instance, David Wong is a moral relativist who gives a different answer to Q1:

A1') Morality is relative to the extension of the phrase 'adequate moral system'.

The notions that Harman and Wong are relying on here, and the exact ways in which their versions of moral relativism differ, will be explored in depth in sections 2.5 and 2.6. Different versions of moral relativism offer up different ways of answering 'Yes' to Q1, but they all agree to the extent that an affirmative answer is the correct one to give. On the other side of the fence, there also appear to be a variety of ways that Q1 may be answered in the negative, although this is not as straightforward as it seems. The person who answers 'No' to Q1 wishes to deny that morality is relative to anything, but it proves to be somewhat tricky to delineate the exact nature of the claim that is being made here. The range of different positions available on either side of the question increases the complexity of the task of finding an answer to Q1. This is not helped by the vagueness of the two key terms in Q1, 'morality', and 'relative'. I propose to replace both of these words with terms that I believe make the question more precise and thus more easily answered. However, before doing this there are some preliminary matters to be dealt with.

1.3 Preliminaries

Firstly, according to the traditional distinction between metaethics and normative ethics, Q1 would be considered a metaethical question. Traditionally, metaethics is the philosophical discipline of reflecting on moral discussion, both disputes and agreements, while normative ethics involves

actually taking part in moral discussion. For example, it would be an exercise in normative ethics to argue that human slavery is morally wrong, while it would be an exercise in metaethics to argue that such a conclusion does not apply in societies in which human slavery is routinely practiced and accepted by all. While the distinction does at times become a little blurry, my study is an attempt to put forward a metaethical theory and not a normative view.

Secondly, when discussing moral relativism, the metaethical issue must be kept separate from what could be called the descriptive, or the empirical issue. One of the necessary features of any adequate metaethical discussion of Q1 will be the taking of some sort of stance towards apparent moral diversity. The moral relativist stresses the relativist implications for morality of the apparent diversity in customs and values across time and across cultures and across groups within cultures (and perhaps also across individuals). It seems to be a matter of empirical fact that there is, or at least appears to be, such diversity, and in what follows I will take it as given that (at least apparent) moral diversity of the sort described above has been established through empirical disciplines such as anthropology and others. Granting this does not, however, mean that there may not be some essential underlying similarities across what appear on the surface to be diverse moral systems. This very issue of what metaethical conclusions to draw from given empirical evidence is one of the hotly disputed aspects of the argument between the moral relativist and his or her opponent.

Thirdly, it is worth noting the distinction between what has been called discriminating relativism and indiscriminate relativism. (Snare, p 91) The latter position holds that in some important sense *everything*, all claims about anything whatsoever, are relative, and moral relativism is just one instance of this overall relativism. The former position holds that morality is relative in a way that is somehow due to the nature of morality itself. Indiscriminate moral relativism may be an interesting position but it is doubtful whether it is a form of *moral* relativism at all. The issues that it raises in metaphysics and epistemology do not seem to be distinctively metaethical issues. To

12

investigate them properly would take me too far from my intended area of inquiry, so in what follows I will focus exclusively on discriminating moral relativism.

Fourthly, in my examination of moral relativism I will restrict myself to explicit consideration of the so-called categories of deontic evaluation, that is, the moral rightness or wrongness (or optionality) of actions, rather than the so-called categories of value, that is, the goodness or badness (or value-neutrality) of people and things. (Timmons, p 7-10) What is said about the former categories will be capable of being rephrased so as to apply to the latter categories as well, although it would be unnecessarily cumbersome to continually keep doing so.

Fifthly, I will sometimes use the phrase: "the practice of morality". For my purposes, the practice of morality includes (at the least) both normative moral discourse as well as actions that are, so to speak, morally inspired, or taken with the purpose of doing what is morally right. In general moral action is preceded by moral discourse about the action. In what follows I will focus on moral discourse about actions rather than actions considered in isolation, as the former seems to be the ground on which the disagreement between the moral relativist and their opponent is most straightforwardly contested. When I mention actions I take it that these are actions that are subject to moral discussion for the purpose of making moral evaluations about them.

Finally, the issue of cognitivism versus noncognitivism is another subject of much metaethical dispute, although it does not have such a long history and is not so widely known about in non-philosophical circles. The cognitivist believes that moral claims may be literally true or false whereas the noncognitivist denies this, believing that moral claims are instead expressions of emotion that are not literally true or false, although they may still be meaningful. This is not an issue that separates moral relativism from its opposition, as the philosophers I examine, both those who are moral relativists and their opponents, are all cognitivists. This is not because I think noncognitivists have nothing important to contribute to the discussion, but

merely to avoid becoming entangled in issues surrounding the dispute between cognitivists and noncognitivists. Hence, I will overlook this dispute in what follows.

1.4 Moral principles and moral judgements

The presence of the less-than-exact term 'morality' in Q1 contributes to the vagueness of the question. I propose to sharpen the focus of my investigation by replacing this term. To do so I need to distinguish between what I will call 'moral principles' and what I will call 'moral judgements'. Moral principles occur in moral discussion as standards by reference to which actions are judged. They are quite general rules, for example: 'Breaking a promise is morally wrong', or 'Human slavery is morally wrong'. When moral principles are combined with facts about the circumstances of a particular situation they give rise to moral judgements. Moral judgements are specific instructions about what is the morally correct thing for a particular person to do in a particular situation. Q1 seems to be an attempt to frame a metaethical question about moral principles and not about moral judgements.

For instance, consider two moral agents, both of whom perform an action of type A. Let us assume that A is in itself a morally neutral action. If one of the agents has previously made a promise not to A and the other agent has previously made a promise to A, then moral judgements would likely be forthcoming that the first agent ought not to A, while the second agent ought to A. Supposing that there is a general moral principle that requires agents not to break promises they have made, we can say that this one moral principle, in conjunction with different circumstances on the part of different agents, that is, the different promises they have made, leads to different moral judgements about the actions of each agent. This is one way in which moral judgements may be said to be relative; they are relative to the circumstances of the agents whose actions they are about.

The moral relativist seems to be making a claim that is over and above this one, however. They are claiming that it is the moral principles themselves that

14

are relative, although what they are said to be relative to will vary depending on which moral relativist is making the claim. If the moral principle requiring agents not to break promises is, as Harman puts it, relative to a moral framework, then it is conceivable that both of the agents in the example above might have made a promise not to A, and yet have different moral frameworks in place, so that all other things being equal it is morally wrong for one of them to break a promise but it is not morally wrong for the other one. This would of course also lead to different moral judgements being made about each of the agents as before, but this time the reason for the difference is due to the relativity of moral principles rather than the difference in circumstances that we saw in the first example. The versions of moral relativism that I will be examining, and it seems to me any interesting version of moral relativism, will take the view that it is moral principles as well as moral judgements that are relative, and not only moral judgements by themselves.

There are of course numerous other metaethical questions that might be raised about moral principles and moral judgements and the relationship between them, but as I have mentioned above, I want to bypass these as not directly relevant to my discussion of moral relativism. The ontological status of moral principles and moral judgements and whether one, both, or neither 'exist' or are 'real' will be overlooked. Questions about the possible psychological processes going on while moving from moral principles to moral judgements, where exactly the line must be drawn as to the level of generality required of a moral principle, and whether 'morality' might contain more elements than these two, must all be skirted in order to focus on the matter at hand. I intend the distinction between moral principles and moral judgements to be a neutral one that leaves it open just what sort of method of moral discussion might be used to move from moral principles to moral judgements. It might involve some sort of utilitarian calculation, an empathetic identification with others, some kind of rational deduction, or any of a number of other possibilities. I do not want to rule anything in or out regarding matters such as these, as their impact on my topic is minimal.

15

I propose to replace Q1 with Q2, a replacement that I think makes the question more precise:

Q2) Are moral principles relative?

What I want to achieve with this replacement is to emphasise the difference between the commonplace and metaethically uncontroversial observation that moral judgements are relative to the circumstances of particular cases and the controversial metaethical claim that moral principles are relative. I take it that it is not an argument for the relativism of moral principles to point out that moral judgements are relative. It may be accepted that moral judgements can be relative without such acceptance in itself committing one to the view that moral principles are also relative. In what follows I will focus on the question of whether moral principles are relative.

1.5 Opposition to moral relativism

In this section I will argue that, despite the longevity of the debate, the traditional dimensions along which the dispute between the relativist and their opponent has been characterised are misleading and inadequate. Progress can be made towards clarifying the meaning of the word 'relative' as it occurs in Q2 by examining more closely the position occupied by those who would answer 'No' to this question. Two common labels are used for the person who answers 'No' to Q2. The labels are 'Absolutist' and 'Objectivist'. Neither of these are entirely suitable.

'Absolutist' is an inaccurate label for the opponent of the relativist because it implies that answering Q2 in the negative commits one to a belief in absolute moral principles that may not be overridden. To an absolutist there are no such things as defeasible moral principles or mitigating circumstances concerning the obligations that arise from moral principles – if one has a moral obligation to perform or refrain from an action then this obligation is unconditional,

16

indeed, it is absolute. According to the absolutist, morality is in an important sense an all or nothing affair. Moral dilemmas are only apparent, so that the notion of weighing moral obligations one against the other and deciding that a moral obligation may have been outweighed in a particular situation is not the correct approach to morality. Now, some who answer Q2 in the negative may indeed be absolutists, but it is not clear that all who do so must be committed to such an uncompromising view of morality. Someone may believe that moral principles are not absolute and yet this belief in itself need not make them a relativist.

Furthermore, there is no contradiction in the position of someone who answers 'Yes' to Q2 and at the same time takes an absolutist stance towards moral principles. One may believe that moral principles impose absolute requirements to perform or refrain from certain actions while also maintaining that these principles are relative. For instance, a relativist may hold that it is an absolute moral requirement that those living in 21st century democracies not keep human slaves, and that there are no possible mitigating circumstances that would make it morally permissible for such people to keep human slaves. This belief need not conflict with their belief that it was morally acceptable for the ancient Greeks to keep human slaves.

So it seems that answering 'No' to Q2 is neither a necessary nor a sufficient condition for being a moral absolutist. Not all those opposed to relativists need be absolutists, and some absolutists may be relativists. The question of whether moral principles are or are not absolute is separate from the question of whether moral principles are or are not relative.

'Objectivist' is also an inaccurate label for those who answer 'No' to Q2, for a similar reason. The metaethical dispute between the moral objectivist and their opponent the moral subjectivist, is distinct from the issue that the moral relativist is seeking to raise. The objectivist has the view that moral principles are in some way ontologically separate and independent of the cognitive states, in particular the beliefs, of those who are bound by them, while the

17

subjectivist has the view that moral principles are somehow dependent on these cognitive states. There seems to be nothing inconsistent in the claim that moral principles are objective and yet also relative. On this view it would just be an objective fact that the ancient Greeks were bound by a different moral principle regarding human slavery from the one that we are bound by. An analogy here would be with the legal rules of the road. It is an objective fact of the matter as to which side of the road a driver must drive on, but this fact is relative to the part of the world that the driver is driving in. It seems that even moral principles held by but a single person could be objective in this way; it may objectively be the case that someone has done something morally wrong according to their own personal standards.

Furthermore, the subjectivist's rejection of external moral principles need not necessarily commit them to moral relativism. It may be held that moral principles are dependent in some way on cognitive states but that the relevant cognitive states are ultimately the same in all moral agents so that the same moral principles are ultimately behind all of the apparent moral diversity that the moral relativist points out.

So, as with the absolutist above, answering 'No' to Q2 is neither a necessary nor a sufficient condition for being an objectivist. What I have said in the above paragraphs is not to deny that many or even most of those who oppose moral relativism may also happen to be absolutists or objectivists or both absolutists and objectivists. The point is that neither the moral absolutist nor the moral objectivist are necessarily in opposition to the moral relativist with regard to the central issue that the relativist raises.

I want to suggest that a more accurate label for the person who answers 'No' to Q2 is 'Universalist'. Given the above difficulties with the labels 'Absolutist' and 'Objectivist', in what follows I will use the term 'Universalist' to refer to the person who is opposed to moral relativism. I believe this term shows most clearly just what it is that opposition to moral relativism commits one to. The universalist wants to assert that moral

18

principles are universal – they apply to all moral agents across all times and cultures and across all groups within cultures. According to the universalist there is one and only one correct set of moral principles, so that, for instance, there is one and only one correct stance to take towards a moral question such as the question of the moral status of human slavery. If we differ with the ancient Greeks about this status then one (or possibly both) of us must be incorrect. To the universalist there is a logical connection between moral agency and a single set of moral principles.

As an aside to forestall possible confusion, the position of universalism is not especially connected with the formal principle of universalisability. As applied to moral judgements, formal universalisability means that if a moral judgement is made about a moral agent in a particular set of circumstances, then any other moral agent in that same particular set of circumstances would be subject to the same moral judgement as well. So if somebody makes a moral judgement that someone else has a certain moral obligation to them, then they must be prepared to admit that they themselves would be likewise morally obligated towards this other person if their positions were reversed. This formal principle is something that the moral relativist may accept as easily as the moral universalist. Formal universalisability is not an issue that separates the moral universalist and the moral relativist, in spite of the similar sound of two of the words. (See Singer, p 822-3 for an account of the difference between formal universalisability and substantive principles such as Kant's categorical imperative.)

I contend that the moral universalist is the true opponent of the moral relativist. To answer 'Yes' to Q2, as the moral relativist does, is to commit oneself to the possibility that there could be more than one set of moral principles - that no one set of moral principles is universal. Because of this, the moral relativist thinks that asserting a moral principle on its own without giving any context for the principle is not something that can legitimately be done in moral discussion. Some reference must always be supplied to indicate the group (or perhaps the individual) that any given moral principle is relative

19

to. In other words, to the moral relativist moral principles are irreducibly context-dependent. I think that these observations put us in a better position to see more clearly just what is being asserted or denied when an answer to Q2 is given. We can rephrase the question once more to reflect this insight. I suggest replacing Q2 with Q3:

Q3: Are moral principles context-dependent?

Answering 'Yes' to this question is to assert the view that a moral principle is analogous to a statement like 'It is ten to five'. The latter statement may be pragmatically adequate if the context is sufficiently obvious, but it is accepted that it is an abbreviation. A complete statement of the time must include a reference to the location in which the utterance is made and the moment at which it is made – something like: 'It is now ten to five in Wellington, New Zealand.' A statement about the time is relativised to place and moment and these context-dependant elements are an irreducible part of such statements, although they may often be omitted as unnecessary because of the obviousness of the context. When asked for the time it would seem overly pedantic to respond: 'Do you mean the time at this moment and in this location?', but that only underlines the irreducible context-dependency present here. The likely response to the latter question will be, 'Yes, of course!' and certainly not, 'No, I want to know the time in a universal sense, without relativising it to anything. I am aware that at different moments and locations people may give differing answers when they are asked for the time, but I am not interested in that. I want to know which of these answers (if any) gives the universally correct time.'

The situation is quite different if asked a question such as: 'What is the moral status of human slavery?' Certainly someone may ask this question in the same way in which people ask for the time. Perhaps a trader in human slaves may be trying to ascertain the business prospects in a new city. However, until we can give a plausible answer to Q3 it is not obvious how

appropriate it would be to declare: 'I am aware that at different places and times people may give differing answers when they are asked about the moral status of slavery, but I am not interested in that. I want to know which of these answers (if any) gives the universally correct moral status of human slavery.'

Notice that, even if the universalist (who does, of course, declare this) thinks that they are able to supply the answer to the question about human slavery, this answer in itself will not enable them to answer the question: 'Is 'Human slavery is morally wrong' an objective or a subjective moral principle?' Nor will it enable them to answer the question: 'Is 'Human slavery is morally wrong' an absolute moral principle or not?' This again emphasises the difference between the moral universalist and the moral objectivist and moral absolutist. Of course the universalist does need to take a position on the latter two questions, but their commitment to universalism need not force their hand as to what their position will be.

The above analogy between moral principles and statements about the time has its limitations. There is, at the level of normative ethics, a degree of controversy about moral principles that has no echo in statements about the time. The controversy is evident from the example, commonly raised by moral relativists, of the dispute over the moral status of abortion. This heated quarrel divides members of our own society in a way that they are never divided over statements about the time. Not surprisingly, the dispute over the moral status of abortion is perceived differently by the moral relativist and the moral universalist. To the relativist both the pro-abortion advocate and the anti-abortion advocate are making a metaethical error in believing that there is a universally correct moral principle to be fixed upon here, while to the universalist both sides are metaethically correct, although only one side (or perhaps neither side) can be said to be advocating the correct approach to abortion at the normative level. Needless to say, the moral relativist and the moral universalist may agree (or disagree) at the normative level about an issue like abortion without this affecting their metaethical disagreement. I will return to the examples of human slavery and abortion later in my discussion.

21

As I have already said, the answer one gives to Q3 has practical implications for the way one does ethics and it will impact on the moral judgements one makes as well as the attitude one has towards moral principles. Q3 has considerable ramifications for ordinary everyday normative ethics. Depending on how one answers this question, the practical business of making everyday moral judgements will have fundamentally different theoretical underpinnings that will effect the process of making these judgements. To the moral relativist, having some workable means of determining the relevant and appropriate context upon which moral principles are dependant will be something of crucial importance. Conversely, moral universalists must have some workable way of determining just what the one universal set of moral principles is, even if this way is nothing more sophisticated than an appeal to authority or to tradition. These two different approaches to the practical business of making moral judgements may not necessarily lead to moral relativists and moral universalists giving different moral judgements about a particular case, but even if their moral judgements concur it would be something like a happy accident that they have reached the same destination by taking different paths. Any differences due to their metaethical disagreement would be over and above any variation in moral judgements that may occur as a result of using different methods to move from moral principles (whatever their status is thought to be) to moral judgements.

1.6 Relativism, scepticism and nihilism

Now that I have clarified the metaethical position of the person who answers 'No' to Q3, I think it is necessary to make a brief but important point about the position of the person who answers 'Yes' to Q3. Although it is sometimes presented as such by the universalist, moral relativism in itself has nothing at all to do with the metaethical positions of moral scepticism or moral nihilism. Although there may be a number of varieties of moral scepticism, the essence of the position is to doubt either the legitimacy of moral principles or the

possibility of making moral judgements or both, while the moral nihilist is of the opinion that there are no legitimate moral principles or that it is impossible to make moral judgements or both.

In the absence of an argument demonstrating otherwise, there seems to be nothing inherent in the position of moral relativism itself that entails either moral scepticism or moral nihilism. To the moral relativist the context-dependency of moral principles by no means renders them illegitimate, and according to the moral relativist it is indeed possible to make moral judgements. It matters to the moral relativist what is right and wrong just as much as it matters to the moral universalist, and the relativist is willing to engage in moral discussion to try to discern the correct moral judgements to make in particular circumstances, possibly even more lengthy and complicated discussion than the moral universalist sees the need for. As we saw above, the moral relativist may be both an objectivist and an absolutist, and they may make their moral judgements as fervently and with as much commitment to morality as the moral universalist.

That moral relativism actually leads inevitably to moral scepticism or moral nihilism is something that would have to be established by argument. Some have tried to do this. In section 2.3 below I will look at an argument from the moral universalist Pojman, who tries to discredit moral relativism by arguing that it leads towards nihilism about morality. J. L. Mackie also argues that moral relativism leads towards nihilism about morality, but with the rather different goal of discrediting the institution of morality by showing that all moral claims are actually false. Mackie's argument will not be examined as he targets the objectivity that he thinks is built into moral language, rather than universality, which is my focus.

1.7 An irresolvable dispute?
Although the moral relativist themselves may not necessarily be sceptical or nihilistic about morality, the long-standing opposition between the moral relativist and the moral universalist casts a shadow over the operation of

normative ethics. It is as if these foes inhabit two sets of ancient fortifications set up opposite one another and between which the business of making moral judgements must be conducted. Taking into account their theoretical polar opposition to one another, not to mention the practical ramifications of their dispute, it seems that in order to maintain their own credibility there is a pressing need for the moral relativist and the moral universalist to each refute the other. Unless one side or the other is able to emerge triumphant, moral judgements made by the supporters of either faction may continually be sniped at from the other side. However, each side has become entrenched in their respective positions, and each of them are able to make what seem to them to be reasonable enough criticisms of the other side. This has made an outright victory on the part of one side or the other elusive.

For the universalist, a fundamental flaw in moral relativism is that it appears to lead to an unacceptable situation in which 'anything goes' as far as moral principles are concerned. The moral universalist places great weight on what for them is a elementary consideration. The consideration may be expressed by saying that there are some kinds of acts, such as for instance murder or breaking a promise, that must always be regarded as being of moral concern. The very notion of what constitutes morally acceptable behaviour means that such acts must come to our notice and that we must have something to say about them when we are making moral judgements.

As I have argued above, the universalist is not necessarily an absolutist, so they may believe that there could be situations in which, once all things have been considered, we may make a moral judgement that an act of murder or of breaking a promise is required in order to do what was morally right. But such a situation would be one in which the moral agent was forced (hopefully with regret) to commit a murder or break a promise. The universalist will here call attention to the distinction between giving an excuse and giving a justification concerning such an act. An excuse provides mitigating circumstances to show why an agent has been forced to do something that still remains morally wrong. They may have had to choose the lesser of two evils. A justification shows that

the sort of action the agent has taken is not morally wrong at all. The universalist view is that a particular murder or breaking of a promise might be able to be excused, but not justified. An act of murder or of breaking a promise is itself still immoral and an object of concern to us when we are making moral judgements. (See section 3.3.2 for more on this)

Taking as an example the idea that moral principles are dependent on the context of a moral framework, consider how we would react towards someone who agrees that an action they are contemplating is unquestionably an act of murder, but who says that they use a moral framework according to which murder is not something that is of moral concern, and so the fact that an action is a murder is morally irrelevant. Our reaction could be that such a person seems to be missing something important about morality. The moral universalist accuses moral relativism of being unable to accommodate this reaction. To the moral relativist just about any act, including murder, could be regarded with moral indifference, or even approval, provided that a suitable moral framework were in place. 'Murder is of moral concern in your moral framework, but not in ours', seems to be an intelligible statement under moral relativism. The universalist may think that this makes a mockery of the very notion of morality. As we will see in section 2.3 below, Pojman attempts to develop this concern into the argument that moral relativism does inevitably lead to moral nihilism.

Furthermore, moral relativism also seems to open the possibility that apparently morally irrelevant aspects of life could be treated as being of moral concern: 'Whistling on a Tuesday is of moral concern in our moral framework, even if it is not of moral concern in yours.' A discussion in which the accusation of murder earns no more than a shrug of the shoulders, but the accusation of whistling on a Tuesday elicits condemnation in the strongest terms would seem to the universalist to be a discussion that has little to do with morality. Whatever sort of judgements are being made here, they are not moral judgements. This line of thought will be further developed in section 3.3.2, where I propose a mechanism to accommodate this universalist concern.

The relativist, on the other hand, may be guided by the consideration that universalism refuses to adequately take into account the implications of moral diversity. To the relativist, universal moral principles are something that nobody can ever really be in a position to claim that they have access to. Articulating such principles would seem to require the ability to step outside of one's own personal and cultural context and pronounce upon morality with an impossibly impartial eye. To the relativist there appears to be more than a little high-handedness and wishful thinking in the notion that the universalist considers it possible make such pronouncements. This is especially so when universalism means that large segments of our society, and indeed entire cultural groups, must be somehow incorrect in following the moral principles that they do, or have done. There must be some more convincing way to account for this than just to say that a lot of people are simply mistaken. The apparent possibility of such widespread error about moral principles severely damages the credibility of moral universalism, the relativist may claim.

Each of the considerations put forward is likely to be dismissed by the other side as instances of question-begging. The moral relativist may say that they do not have to accept that there are any moral principles that are above rejection by some group or groups, while the moral universalist may say that they do not have to accept that there is any difficulty with some groups being incorrect about the set of universally correct moral principles. What I propose to do is to develop a metaethical theory that will give weight to both of these considerations and show how they may co-exist within the same theory.

So, there are two fundamental convictions that need to be adequately accounted for. I will call them Conviction U and Conviction R:

I) The moral universalist's fundamental conviction (Conviction U): At least one moral principle is universal, that is, it counts whenever moral judgements are made, and it counts in the same way.

II) The moral relativist's fundamental conviction (Conviction R): At least some moral diversity is due to different groups being committed to different moral

principles. Not every instance of moral diversity can be explained away as due to mistake or misinformation.

These apparently opposing convictions lead the universalist and the relativist to give differing answers to Q3. In my thesis I propose an approach to Q3 that attempts to find a compromise between universalism and relativism and that will leave room for both Conviction U and Conviction R.

1.8 A coherentist approach

In what follows I will offer a compromise between the two apparently incompatible positions of moral universalism and moral relativism. I aim to do justice to the legitimate concerns of both the moral universalist and the moral relativist while at the same time avoiding the problems that I will raise with the various theories I examine during the course of my study.

My approach to the dispute between the moral universalist and the moral relativist could be characterised as a coherentist approach (or coherentism). The long history of the dispute and the current impasse between the protagonists strongly suggests that we do not have available a metaethical decision procedure that directs the complete ruling out of either moral relativism or moral universalism. Neither Conviction U nor Conviction R is able to emerge completely triumphant over the other. This implies that it may be more fruitful to seek a means of allowing these two fundamental convictions to co-exist with each other. The nature of the compromise that I offer will be a metaethical theory that has room for both Conviction U and Conviction R.

In doing this I am advancing a metaethical version of a method that is more often seen at the level of normative ethics. In order to motivate my brand of coherentism I will first look briefly at how the coherentist approach has been employed by others in the normative arena. Norman Daniels is a prominent advocate of a coherentist approach towards normative questions. He calls this "wide reflective equilibrium." According to Daniels this involves, "attempting to produce coherence in an ordered triple of sets of

beliefs held by a particular person, namely, (a) a set of considered moral judgements, (b) a set of moral principles, and (c) a set of relevant background theories" (p 22). He distinguishes this from "narrow reflective equilibrium," an approach that only recognises the first two sets of beliefs. Daniels' idea is that any of the above three sets of beliefs may be adjusted in the light of any of the others, in order to find a stable point at which further adjustment would only reduce coherence. The overall goal is to, "render theory acceptance in ethics a more tractable problem" (p 21), by providing a method that recognises the role played by nonmoral background theories in normative decision making. There is no decision procedure available that allows us to favour any one of these sets of beliefs over the others, and so all three must be taken into account.

The coherentist approach at the normative level has been taken up and further developed by David DeGrazia, who places emphasis on the theoretical virtues to be sought from the method. DeGrazia explains that coherentism aspires to more than just logical consistency, but also what he calls "argumentative support" and "global illumination" (p 14-5). The first is the demand that something that counts as a reason or justification in one area of a coherentist theory must be recognised as counting as a reason or justification across all parts of the theory. The second is the demand that, "there must be a tolerable amount of illumination over the whole" of a coherentist theory. The theory must be able to explain how its various elements fit together, particularly if they seem initially to be at odds.

An example of how DeGrazia employs the coherentist methodology is his investigation into the question of the nature and extent of the mental lives of animals. It is worth looking at how he does this, to see whether there is any parallel between his project and mine. He seeks to fit together evidence about mental life from four different sources: human phenomenology, the study of animal behaviour, functional-evolutionary arguments, and physiological evidence (p 77). He says: "The idea is to treat with respect each kind of evidence – without prioritizing any of them – and work toward theses about

animal mentation that do justice to all of the available evidence. The four (sub)methods approach mental phenomena in different ways, but they are compatible" (p78).

DeGrazia's first source of evidence is human phenomenology, that is, how things seem and feel to humans. This is useful both as a means for setting the agenda about what mental life in animals could be like, and also as a starting point for understanding the qualitative features of mentation. The second source of evidence, observations of animal behaviour, is valuable because of the close connections between behaviour and mental life. This includes observations of how animals behave both in laboratory studies and also in their natural habitats. The third source of evidence, functional-evolutionary arguments, involves investigation of the possibility that possession of mental states confers a selective advantage on the possessor. The fourth source of evidence, animal physiology, also includes neurophysiology, the study of the brain and central nervous system, and reflects the importance of understanding the physical systems most likely to be associated with mental life (p 78-84).

DeGrazia says that, although to him it seems obvious that all four of these spheres of evidence have relevance in looking at the question of the mental life of animals, many philosophers and researchers tend to prioritise one source of evidence to the exclusion of the others. For instance, DeGrazia mentions Descartes and other dualists as prioritising phenomenological evidence, behaviorists and functionalists as prioritising behavioral evidence, and so-called 'tough-minded' scientists as prioritising physiological evidence (p 88-92). Trying to make results from all of these disparate sources of evidence fit together into a coherent whole requires the jettisoning of prejudices about one source or another taking precedence over the rest.

DeGrazia's actual investigation of the mental life of animals, during which he draws upon all four of these sources of evidence and works their results into a coherent whole, is not easily summarisable. Fortunately such summary is unnecessary for my purposes. The point of bringing up this example is to

highlight the coherentist methodology DeGrazia pursues. Four sources of evidence normally taken to be divergent can be shown to fit together when they are all taken on an equal footing with none being given priority over the others. The parallel between DeGrazia's project and mine is that in both cases apparently conflicting considerations must be made to cohere.

I can see no reason why coherentism must be restricted to the normative level. A coherentist approach is attractive in any situation where we do not have an acknowledged decision procedure that allows us to settle on one of a number of conflicting views, and I will try to show this to be the case with the dispute between the moral relativist and the moral universalist. In relocating coherentism to the metaethical level I will take the core methodological recommendations articulated by Daniels and DeGrazia and apply them to the problem of the dispute between the moral universalist and the moral relativist. The coherentist method essentially requires the fitting together of diverse ingredients into a consistent theory and so I will try to achieve a metaethical account that shows how Conviction U and Conviction R can fit together. Rather than attempting to appeal to some elusive master principle or over-riding decision procedure that rules out one or the other of these convictions once and for all, I will not deny or privilege either of them. My goal is a maximally coherent account that has room for both convictions. The most coherent account in this case could be seen as a form of compatibilism, as the two convictions to be fitted together seem to oppose each other and so must be shown to be compatible.

DeGrazia notes two criticisms that coherentism faces. Both of these criticisms focus on the initial elements that the coherentist tries to fit together, and could even be said to be different versions of the same criticism. They are, firstly, that coherentism is only fruitful if there is some sort of credibility attaching to the initial elements (p 19-22), and secondly, that the initial elements may be biased towards or against some normative position (p 22-31). However, the opposition between Conviction U and Conviction R protects the metaethical coherentism I am putting forward from these criticisms. Firstly, the

credibility of both Conviction U and Conviction R has already been thoroughly tested during the course of the long-running dispute between the moral relativist and the moral universalist. Although neither side may believe that their opponent's fundamental conviction is very credible, there seems to be no argument that clearly discredits either of them. This is, of course, the main reason why the dispute has proved so intractable. Secondly, it would be stating the obvious to say that Conviction U is biased towards moral universalism and Conviction R is biased towards moral relativism, but this is not a problem as these opposing biases balance each other out. Making room for both of these convictions and not privileging either of them prevents my approach from being biased towards either of them.

I do not try to say that either the moral relativist or the moral absolutist is completely correct while their opponents are completely incorrect. To anticipate, my answer to Q3 ('Are moral principles context-dependent?') is, roughly, 'There are two types of moral principles, those that are context-dependent and those that are not', although this answer requires a number of clarifications and qualifications to be made in due course. With this answer I try to undermine the supposition that all moral principles are either context-dependent or not context-dependent, a supposition that appears attractive to both sides of the dispute. For this reason, the theory I advance may possibly not be palatable to hard-line advocates of either moral relativism or moral universalism. However, it seems to me that the hard-liners have fought each other to a standstill on this question, and that the time is ripe for a compromise. I hope to do sufficient justice to both sides so that the compromise I offer is seen as a viable way forward.

One of the primary merits of the coherentist approach to the metaethical dispute between the moral relativist and the moral universalist is the implications it has for the practical business of making everyday moral judgements. In showing how Conviction U and Conviction R can coherently fit together I hope to go some way towards easing the metaethical anxieties fueled by this dispute. For instance, undermining the force of the challenge

31

raised by the student in the applied ethics tutorial who maintains that moral relativism makes any moral discussion pointless. Much metaethical work remains to be done even if my approach is acceptable, but I believe something significant will have been achieved in showing that the two convictions that motivate the moral universalist and the moral relativist need not necessarily be assumed to be incompatible with one other.

1.9 Thesis plan

Pressing the case for a coherentist accord between Conviction U and Conviction R requires that I make an in-depth examination of the current state of the dispute. In Part Two I will take a close look at moral relativism and the details of the theories of some prominent moral relativists. First I survey some universalist criticisms of moral relativism along with relativist replies to these criticisms before focusing on the theories offered by three prominent contemporary moral relativists. These relativists are Harman (who has already been mentioned), David Wong, and Robert Arrington.

In turning to examine moral universalism I am presented with an immediate difficulty. Rather than three prominent theorists, it is fair to say that there are literally scores of moral philosophers who could be seen as presenting theories that are universalist to some degree. Added to this is the problem that these theories seem to be only partially or indirectly universalist and defend moral universalism in conjunction with a defense of moral absolutism and/or moral objectivism. This is connected with the confusion discussed in section 1.5 over just which metaethical position is opposed to moral relativism. When I come to discussing universalism in Part Three I will therefore take a different tack from that of the discussion of relativism in part Two. I will focus on universalism from a more general perspective, attempting to separate out the issues that confront moral universalism as a position, rather than pursuing close criticism of individual theorists.

In my discussions of these positions I will be looking to decide whether there is a reason to privilege the fundamental conviction that lies behind each

of them, Conviction U in the case of moral universalism and Conviction R in the case of moral relativism. The most obvious way to privilege one of these convictions would be to rule out the other. To anticipate, I find that there is not a reason to privilege Conviction R, as the discussion of Part Two will show that the three most prominent contemporary versions of moral relativism all contain serious flaws. But nor is there a reason to privilege Conviction U, as it is plausible that there are more types of moral principle than only universal moral principles. I put forward a conception of what such universal moral principles are like in Part Three, and then in Part Four I present a theory about the nature of the context-dependent moral principles that are also a legitimate part of the practice of morality.

Part Two
Moral Relativism

2.1 Introduction

Moral relativist Gilbert Harman has an interesting analogy between morality
and motion that he uses to help clarify what his view of moral relativism is.
He reminds us that it is only possible to judge whether something is in motion
or at rest relative to some spatio-temporal framework. For example, a
passenger sitting in the dining car of a train will judge that a cup of tea in front
of him is not moving, and relative to his spatio-temporal framework it is not; a
person standing beside the tracks watching the train go by will judge that the
cup of tea is moving, and relative to *his* spatio-temporal framework it is.
Morality is relative in a similar way to the way in which motion is relative,
suggests Harman. (1996, p 3-4) It is only possible to judge whether something
is morally right or morally wrong relative to some moral framework. For
example, an anti-abortion advocate judges that abortion is morally wrong, and
relative to her moral framework it is; a pro-abortion advocate judges that
abortion is not morally wrong, and relative to *her* moral framework is it not.
Furthermore, just as there is no spatio-temporal Archimedean point from
which we can judge whether something is 'really' moving or at rest, so there
is also no moral Archimedean point from which we can judge whether
something is 'really' morally right or morally wrong. Everything depends on
the framework within which we make our judgements, either about motion or
about morality.

Harman does not present this analogy between morality and motion as an
argument for moral relativism. The analogy is an elucidatory device designed
to bring out the importance for him of acknowledging the differing moral

frameworks within which people's moral judgements are made and their moral principles operate. While other moral relativists may not agree with Harman that moral frameworks are what morality is relative to, his analogy between morality and motion still manages to convey a nice general image of what the moral relativist is claiming about morality.

2.2 Some initial universalist objections

As well as the general universalist worry surrounding the perceived relativist threat to Conviction U, there are a number of other more specific universalist criticisms that may initially come to mind when moral relativism is advanced. In this section I will review some of the opening objections that might be made against the view that moral principles are context dependant and some possible replies that moral relativists could make to these objections.

Criticism I) Moral relativism is inconsistent.

The universalist might accuse the relativist of inconsistency (See Arrington, p 201). If one answers 'Yes' to Q3, then they are asserting that all moral principles have a context-dependent element (or elements). But this relativist assertion is presumably meant to apply to all moral principles in any situation or cultural context, and there is no mention of such a relativist assertion itself having any context-dependent elements. In other words, the relativist is making a universalist assertion in the very act of stating their thesis. The relativist will be able to meet this objection by noting that the position they are advancing is a metaethical one, the assertion they are making when stating it is not itself a moral principle and so no problem of inconsistency arises here.

Once this version of the inconsistency objection is cleared out of the way, there is another sort of inconsistency that the moral relativist may be guilty of. Moral relativism may be accompanied by the claim that tolerance should be exercised towards those who operate in contexts other than ours (See Wong, p 175-190). This would include both those in other cultures as well as those within our own culture who disagree with us on controversial moral issues.

It would be inconsistent of the moral relativist to claim that tolerance followed from moral relativism, as the normative claim that we should be tolerant of others is itself an expression of a moral principle that will, if the relativist is correct, be context-dependent. To be consistent, the moral relativist would have to say that whether or not we should be tolerant of others depends on the context. Of course, a metaethical commitment to relativism does not prevent the moral relativist from arguing for tolerance (tolerance within a certain context) at the normative level. Wong is a moral relativist who takes this avenue (see section 2.6.6). Arrington, on the other hand, is a moral relativist who argues against tolerance (see section 2.9).

Criticism II) Moral Relativism is incoherent.

Moral Relativism is incoherent because it entails that moral principles are both true and false. Take the claim 'Abortion is morally wrong.' If moral relativism is the case then this is true when spoken by an anti-abortion advocate and false when spoken by a pro-abortion advocate. Hence, the principle is both true and false, an insupportable situation (See Arrington, p 201). This objection has effectively been blocked by the replacement of Q2 with Q3. The sort of moral relativism endorsed by someone who answers 'Yes' to Q3 is one that sees a moral principle like 'Abortion is morally wrong' as being irreducibly context-dependent, so that the truth value of this moral principle is always linked to the circumstances in which it is uttered. If the moral relativist is correct and moral principles are context-dependent then this means that 'Abortion is morally wrong' may be true in those circumstances in which it is uttered by an anti-abortion advocate and false in those circumstances in which it is uttered by a pro-abortion advocate, but this is no more problematic for moral principles than for any other context-dependent statements in which the truth-value of the statement is tied in some way to the circumstances of utterance. 'It is now ten to five' may be true when spoken in Wellington and false when spoken in Ohio, but this variation in truth value is a harmless feature of the context-dependency of statements about time.

Criticism III) Moral disagreements become pseudo-disagreements.

This follows on from Criticism II. If the correct answer to Q3 is 'Yes', and moral principles do actually contain context-dependent elements, then this means that the so-called moral disagreements that people appear to have are in reality merely pseudo-disagreements (See Arrington, p 201-2). The anti-abortion advocate and the pro-abortion advocate, for instance, despite the often bitter and intransigent nature of their dispute, are actually only appearing to disagree with one another, or they are quite simply mistaken about whether there is a matter for disagreement here. Surely this is implausible, the universalist may say; there really is a genuine dispute about whether or not abortion is morally acceptable. We have to make moral judgements about abortion in particular real-life cases, and some are willing to kill and die for the sake of their side of the dispute.

However, while there is no doubt that arguments about what should be done in particular cases does occur and that there must be a clash of some kind going on when the pro-abortion advocate and the anti-abortion advocate argue, this in itself does not guarantee that the dispute is a genuine one in the sense that the universalist intends. The mere fact that there is a difference of opinion about something does not necessarily mean that one side must be correct and the other side incorrect. Certainly, if the moral status of abortion *is* the sort of issue for which it is possible to apply a non-context-dependant moral principle, one either permitting or forbidding it, then the debate over abortion is indeed a genuine one, and one side or the other has the correct answer while their opponents are incorrect. But this is not the sort of thing that the universalist can afford to simply assume if they want to avoid begging the question. First of all an investigation must be conducted into the issue of what the correct or most plausible answer to Q3 is, before the universalist can be in any position to make a judgement about the genuineness of the disagreement between the pro-abortion advocate and the anti-abortion advocate. But even if the universalist concedes this, they may insist that the burden of proof is now

more heavily on the moral relativist at this point. The relativist will need to give some sort of account as to how such a situation could arise and why both sides could be so mistaken as to the nature of the disagreement that they so fervently fight over.

Now the relativist may not be overly distressed at having to do this. Conviction R is itself partly a response to the observation of moral disagreement and diversity and the relativist claims that the wide variety of incompatible moral principles at large in the world is a feature of moral discourse that favours moral relativism rather than moral universalism. The relativist may suggest that passionate disagreement over a question in normative ethics such as the moral status of abortion really lends strong support to the relativist position on the metaethical question of whether moral principles are context-dependent, in spite of what those who actually argue about abortion may think. It could be that those who are embroiled in the normative dispute over abortion are unable to see metaethical matters very clearly in their desire to condemn the moral stance of those who disagree with them about the normative question. The moral relativist may give an account of the origin and causal development of moral discourse that tries to explain how and why moral principles have come to be context-dependent. Below I will examine the theories of two moral relativists, Harman and Wong, each of whom give their own such account in support of their own theory.

Criticism IV) The Genetic Fallacy.

At this point the universalist may introduce a fresh criticism - that the relativist is now in danger of committing the error is reasoning that is known as the genetic fallacy. (Snare, p 116-7) The genetic fallacy is a species of the naturalistic fallacy. The naturalistic fallacy originates with G. E. Moore's objection to the notion that it is possible to define a moral term like 'good' in terms of a natural property such as pleasure (Moore, p 13-4). Moore argued that a definition of this kind would always be vulnerable to his open-question argument. There will be properties about which it would be possible to

intelligibly ask: 'Such-and-such is pleasant, but is it good?' As the answer to a question like this could still be open, that is, the answer would not be self-evident just from knowing the meanings of the words involved, then the question is a legitimate question. This means that the purported definition of 'good' must have failed to fully capture the meaning of the word. However, the naturalistic fallacy is not only restricted to attempts to strictly define moral terms or phrases, but also works against attempts to analyse moral concepts in such a way that the analysis is vulnerable to an open-question argument.

The genetic fallacy is an instance of this. It applies to attempts to fully analyse moral concepts by giving a causal explanation of the genesis of the concept. For instance, a theorist might claim that for an action to be morally wrong is just for it to be disapproved of by those in authority, and support this claim by providing a causal explanation giving empirical evidence about how the concept of an action being morally wrong arose as a means for authority figures to control certain sections of society. This claim would be vulnerable to an open-question argument. It could always intelligibly be asked: 'Such-and-such is disapproved of by those in authority, but is such-and-such wrong?' The critic who proffers the open-question argument here would be supposing that the answer to the question was not self-evident merely from understanding the concept of moral wrongness. This means the question would be a legitimate one and the purported analysis would fail to adequately capture the moral concept. The critic is not saying that the causal explanation that was given was false, but even if it were true or likely to be true it could not do all the metaethical work that was required of it and so it could not constitute a complete analysis of the moral concept under investigation.

The genetic fallacy (and the naturalistic fallacy) rest upon the familiar fact/value distinction. This distinction is a recognition of the logical gap between saying that something is or is not the case, and saying that something should or should not be the case. A moral relativist would be committing the genetic fallacy if they were attempting to infer the metaethical conclusion that moral relativism is the case solely from empirical premises about the causal

genesis of the diversity of moral principles. A causal explanation giving empirical evidence as to why various different groups may adhere to the differing moral principles that they do is not sufficient to establish the metaethical case that moral principles are context-dependant. Although it is well known and accepted that the genetic fallacy is a error in reasoning, it never the less appears in the work of a number of ethical theorists (and not only moral relativists) in various different disguises, as I will try to show below.

Moral relativism is not automatically vulnerable to the genetic fallacy. It may be possible to explain why moral principles are context-dependent just by making use of other moral notions and without any explicit reference to casual origins. I will investigate this possibility below. In any case, while the moral relativist who is able to come up with a causal explanation for moral diversity may accept that this does not by itself suffice as a complete metaethical argument for moral relativism, they may well allege that the burden of proof has now been passed back to the universalist. It is now up to the universalist to give their account as to why there is only one correct set of moral principles, so that for any normative controversy such as that surrounding abortion, one (or perhaps both) of the sides to the dispute must be mistaken at the normative level.

Clearly more could be said by both the moral relativist and the moral universalist about the criticisms raised above, and the above argument sketches may be only the opening moves in a long and complicated metaethical campaign. What has emerged from the above brief discussion is that anyone who seeks to adjudicate between the moral relativist and the moral universalist and present a view as to how best to answer Q3, must be able to give an account of the nature of moral diversity and disagreement. Such an account must take seriously the points made above by both the moral relativist and the moral universalist. That is one of the things that I will try to do in the metaethical theory that I present over the course of this study.

Continuing the review of moral relativism, I will now look at some more criticisms of moral relativism that have been offered by two universalists, Louis Pojman and James Rachels. I find that their criticisms are directed against certain specific versions of moral relativism, and while their objections may be successful against these versions of relativism, they will only work against a moral relativist who makes some quite specific commitments.

2.3 Pojman and individualistic moral relativism

One criticism that might be advanced by a universalist is directed towards a certain interpretation of moral relativism. Moral relativism might seem to offer the unwelcome possibility that it could be up to an individual to decide for themselves just what moral principles it was appropriate for them to follow. This interpretation of moral relativism (I will call it individualistic moral relativism) is the target of criticism from Pojman, whose argument I will now examine.

According to the view offered by individualistic moral relativism each individual is the sole judge as to what are the morally correct principles for them. On such a view an individual's present acceptance of a moral principle is both necessary and sufficient for that moral principle to apply to them. Such acceptance may or may not be the result of some form of reasoning that leads to the principle. It may seem that even to speak of moral principles in conjunction with individualistic moral relativism would be going too far – it seems that it would be possible for the individual to make a fresh decision in each case about what it is morally right for him or her to do, without regard to previous decisions made in similar situations.

This is the very difficulty that Pojman finds with individualistic moral relativism (he calls it "subjectivism", although it is a different view from the view I called subjectivism in section 1.5). To Pojman's subjectivist, the morally right thing to do at any given moment is what one feels like doing. This view, says Pojman, entails that: "little or no interpersonal criticism or

judgement is logically possible." (Pojman, p 509) The ramifications of such a stance will be that individuals are morally incorrigible, or unable to be corrected on moral matters, and their moral judgements are not open to any real discussion, except perhaps about the accuracy of their reports as to what they feel like doing. An individual may feel like committing a murder on Monday and so would be doing nothing morally wrong in committing murder on that day, but come Tuesday they may feel differently, and so whether it is wrong for them to commit murder on that day will likewise fluctuate. This may reveal hypocrisy on their part, but so long as they feel like being a hypocrite, their apparently unprincipled behavior may not be questioned on moral grounds. Pojman: "For [the individualistic moral relativist], both hypocrisy and nonhypocrisy are morally permissible." As Pojman insists, the resulting incorrigibility of moral agents robs moral discourse of its normative character, and "makes morality a useless concept." (p 509) Pojman is trying to show that individualistic moral relativism leads inevitably to moral nihilism.

However, while there is clearly something faulty about the individualistic moral relativism that Pojman describes, it is worthwhile looking at just why it is flawed. While individualistic moral relativism may not be criticised at the normative level on the grounds of hypocrisy, it may be criticised at the metaethical level on the grounds of the incorrigibility it leads to. Presumably it is this deficiency at the metaethical level that Pojman rightly finds so objectionable. Whatever we may think, normatively speaking, of the hypocrisy of the individualistic moral relativist described in the previous paragraph, it is their outright incorrigibility at the metaethical level that makes them appear so unreasonable. If we take the stance Pojman describes to be the paradigm of individualistic moral relativism then the incorrigibility of this view may indeed be criticised, as Pojman himself does. The very point of the concept of morality would, as he says, be under threat if such incorrigibility were permitted at the level of metaethical examination. The individualistic moral relativist comes across as really being an eliminativist about morality. Proposing to follow a system in which corrigibility has no place is proposing

to follow a system that is not a moral system. Incorrigibility leads straight to nihilism.

But it seems possible, given that incorrigibility is the stumbling block for Pojman's subjectivist, for there to be a version of individualistic moral relativism that avoided this weakness. Such a view could be called principled individualistic moral relativism. Here the individual, while following a set of moral principles that is unique to themselves, tries to follow their own singular moral principles as consistently as they possibly can, and is furthermore prepared to change their intended actions in the light of criticism from others that they are not consistently adhering to their own principles.

A position such as principled individualistic moral relativism would be possible to hold without violating metaethical scruples about incorrigibility and so not threatening the very concept of morality. Given this, it seems that the size of the group or groups that morality might be relative to does not matter all that much. It is the incorrigibility of the unprincipled individualistic moral relativist that is their undoing, not the fact that they follow a singular moral code. Pojman's criticism of this interpretation of moral relativism does make an important point – that any position that does not allow for the normative nature of moral claims cannot be a form of *moral* relativism (although it might be some other kind of relativism).

So it seems that whenever moral relativism is interpreted in an unprincipled individualistic manner it will be vulnerable to Pojman's criticism. However, there are other (more plausible) varieties of moral relativism that are neither unprincipled nor individualistic. Some of them I will examine below. Pojman's criticism does not offer any threat to them, as it is so specifically directed against unprincipled individualistic moral relativism only. In what follows I will proceed under the assumption that the theorist attempting to articulate the most reasonable possible version of moral relativism will avoid the version that Pojman identifies and criticises.

2.4 Rachels and cultural relativism

James Rachels is a universalist who gives a number of criticisms of a position that he calls cultural relativism. It emerges that cultural relativist positions form a subset of relativist positions. It is worthwhile looking at the arguments that Rachels has to offer against cultural relativism both in order to bring out the nature of the difference between cultural relativism and moral relativism, and also to serve as a first glimpse at what seems to be a ubiquitous theme in universalist arguments against relativists.

According to Rachels there are six principles of cultural relativism:

> 1. Different societies have different moral codes.
> 2. There is no objective standard that can be used to judge one societal code better than another.
> 3. The moral code of our society has no special status; it is merely one among many.
> 4. There is no "universal truth" in ethics – that is, there are no moral truths that hold for all peoples at all times.
> 5. The moral code of a society determines what is right within that society; that is, if the moral code of a society says that a certain action is right, then that action is right, at least within that society.
> 6. It is mere arrogance for us to try to judge the conduct of other peoples. We should adopt an attitude of tolerance toward the practices of other cultures. (Rachels, p 489)

From the point of view of the dispute I am concerned with, between the moral relativist and the moral universalist, principle 1 is an empirical claim that both can accept, and principle 2 is an objectivist claim that need not be a point of difference between them (see section 1.5), unless principle 2 is understood as: "There is no single objective standard that can be used to judge one societal code better than another", in which case the universalist must reject it. Principles 3, 4, and 5 seem to have the most bearing on the dispute, particularly principle 4, which I have characterised as the chief matter of contention between the moral relativist and the moral universalist. Principle 6 is an explicitly normative principle, and hence different in kind from principles 1-5 (See section 2.2). It may be that Rachels conceives cultural relativism as a sort of hybrid theory that contains both metaethical and normative commitments. In order to focus on metaethical argument, I will

ignore principle 6 in what follows. Rachels gives some criticisms of the cultural relativist who adopts principles 1-5. I will evaluate his criticisms.

Rachels condemns a certain type of argument that might be given for cultural relativism, one he calls the cultural differences argument. He characterises this argument as follows: "The strategy used by cultural relativists is to argue from facts about the differences between cultural outlooks to a conclusion about the status of morality." (p 489) He rejects this argument because it commits the genetic fallacy – the failure to distinguish facts about the causal origin of moral beliefs from arguments that establish the metaethical status of these beliefs (see section 2.2). This seems to be quite a respectable line of argument, as a bare appeal to facts about diverse moral beliefs will not in itself establish moral relativism.

However, as I mentioned in section 2.2, a universalist will have to say something themselves about diverse moral beliefs. Of the many varied moral beliefs that have been and are currently subscribed to, some convincing account needs to be given as to what makes only one particular set of these moral beliefs correct and the rest of them incorrect, and also an account of how it is that some people could come to be in error about which are the correct moral principles. Rachels shoulders the burden of proof here by arguing for "a general theoretical point ... that there are some moral rules that all societies will have in common, because those rules are necessary for society to exist." (p 493) The problem for Rachels is that this argument itself commits the genetic fallacy. I will look at this further below.

Rachels also presents a modus tollens argument designed to show that cultural relativism is not the case. He puts forward three implications he thinks follow from cultural relativism, and then tries to show that as these implications cannot be accepted, neither can the view that leads to them. Rachels says that if cultural relativism is the case then the following claims are true:

1) We could no longer say that the customs of other societies are morally inferior to our own.

2) We could decide whether actions are right or wrong just by consulting the standards of our society.

3) The idea of moral progress is called into doubt. (p 490-1)

But Rachels believes that that 1) – 3) are not true; we can say that the customs of other societies are inferior to our own, we cannot decide whether actions are right or wrong just by consulting the standards of our society, and moral progress has undoubtedly occurred. Hence, according to this argument cultural relativism cannot be the case.

Underlying Rachels modus tollens argument against the cultural relativist is an appeal to Conviction U; the notion that there are some kinds of acts that are morally intolerable and the resulting fear that moral relativism may open the door for such acts to somehow become morally justified in certain circumstances. This is a legitimate concern and I propose a way of addressing it when I present my own view in 3.3.2 below. The point I want to make here is that Rachels' modus tollens argument does not succeed in harnessing this intuition in a way that poses much of a challenge to the moral relativist.

The first thing to notice about this argument is that the statements 1) and 3) are really different ways of saying the same thing. To claim that moral progress has occurred just is to claim that the customs of societies in the past are morally inferior to our own. But when he makes his assertions that 1) and 3) are false, it is difficult to see how Rachels is doing anything other than begging the question. The claim that we can say that the customs of our own society are morally superior to the customs of other societies, whether these are societies that existed in the past or those that are contemporaneous with our own, is one of the very things that is at issue between the moral universalist and the relativist (whether a cultural relativist or a moral relativist).

Rachels' claim that 1) and 3) are false certainly has the potential to be an effective rhetorical ploy to use against the cultural relativist, and Rachels does make full use of the rhetorical possibilities here, invoking slavery, anti-Semitism, and gender inequality as instances of customs of other societies that are morally inferior to our own, but this ploy does not, of course, constitute an argument against cultural relativism. Rachels is not likely to have many of his readers disagree with him when he says, for instance, "slavery and anti-Semitism seem wrong *wherever* they occur." (p 491, italics in original) But this is to be expected given that his probable audience will be made up of members of his own society. But so long as 'seem wrong' means no more than 'judged by members of our society to be wrong', this would show only that applying one group's moral standards to the customs of another group will be likely to lead the first group to judge that the customs of the second group are morally inferior.

The cultural relativist (and the moral relativist) can accept this, as it does not damage cultural relativism (or moral relativism) for members of our society to judge that the customs of other societies are morally inferior. Undoubtedly members of other societies would judge that the customs of *our* society are morally inferior. After all, this is just what we would expect given that moral diversity occurs and that different societies have different moral codes.

However, it is here that an extra dimension to cultural relativism (according to Rachels' interpretation of the position), and the factor that seperates it as a distinctive subgroup within moral relativism, emerges. Rachels says that, according to the cultural relativist, believing in moral progress, or judging another societies customs to be morally inferior to one's own, "make[s] no sense." (p 491) This is a quite singular claim, and can be regarded as a seventh principle augmenting the list of principles above that delineate the view Rachels is calling cultural relativism. This principle 7 entails that for the cultural relativist it is not just an error to judge that moral

47

progress has occurred or that the customs of other societies are inferior to one's own, but it is literally senseless.

This is the point at which cultural relativism (as Rachels interprets it) is revealed as a somewhat extremist sect within moral relativism. The moral relativist can accept that it makes sense to moral universalists to use what are, according to the relativist, context-dependent moral principles as if they were non-context-dependent by making moral evaluations of the actions that occur in other cultures. Given that it does seem to be at least intelligible to make such evaluations about the customs of other societies, so long as it is understood that any such judgement is relative to one's own moral standards, the moral relativist does not have to deny that it is possible to understand the sentences that people utter when they make such evaluations. For the cultural relativist, as Rachels characterises the position, such evaluations would be on a par with judgements about the colours of numbers, or some other literally senseless utterances.

There are two major insights to be gained from the discussion of Rachels' arguments against cultural relativism. Firstly, cultural relativism (as Rachels interprets it) is a species of moral relativism that requires a commitment to the semantic claim principle 7 above (as well as the normative claim principle 6). Not all versions of moral relativism need to be committed to these principles. Secondly, an appeal to Conviction U must be carefully handled by the universalist. Asking a crowd of Athenians whether it is Athenians or Spartans who do what is right is not the way to argue for moral universalism. Some more legitimate way of accommodating Conviction U must be found.

Pojman's unprincipled individualistic moral relativist and Rachels' cultural relativist both contain glaring weaknesses that make them easy to criticise. It would perhaps not be entirely fair to call the unprincipled individualistic moral relativist and the cultural relativist straw men, as Pojman and Rachels may indeed have come across the sort of opinions that these two positions represent. However, there are far more rigorous and interesting versions of moral relativism available and any moral universalist who ignores

48

these theories would be doing themselves and their position a disservice. In the remainder of this chapter I will critically examine three serious versions of moral relativism, presented by Gilbert Harman, David Wong, and Robert Arrington.

2.5 Harman's moral relativism

2.5.1 Moral diversity

Harman is perhaps the classic example of a moral relativist who is impressed with Conviction R. His basic position is that moral relativism is: "a reasonable inference from the most plausible explanation of moral diversity". (1996, p 8) He gives many examples of moral diversity: across different cultures we find different moral attitudes towards cannibalism, marriage, the treatment of women, slavery, caste systems, theft and infanticide. Within contemporary Western society we find different moral attitudes towards the treatment of animals, abortion, euthanasia, the value of artifacts as compared with human life, how much help one ought to give another, whether killing is worse than letting die and the relative importance of liberty versus equality. (1996, p 8-11)

In Harman's view the most plausible explanation for all this heterogeneity is the operation of a wide variety of different moral frameworks. From this he believes we can reasonably infer that moral relativism is the case. While he is aware that moral diversity in itself does not entail moral relativism, Harman tries to show that it is more reasonable to infer that moral relativism is the case than that anything else is the case.

2.5.2 Harman's criticism of universalism.

Harman brings up and dismisses one possible universalist explanation of moral diversity. This is the notion that some people or cultures are rather poorly placed epistemologically. According to this universalist account, such individuals or groups may have the same relevant experiences as those who are following the universally correct moral system but their false beliefs about

49

nonmoral matters lead them to misinterpret this experience and so they end up adopting moral arrangements that are not correct. (1996, p 12)

Harman does not think that this is a plausible explanation and believes that this kind of account of moral diversity is the result of simple bias in favour of one's own moral framework: "The system of moral coordinates that is determined by a person's own values can be so salient that it can seem to that person to have a special status" (1996, p 13). A commitment to moral universalism in itself provides no guarantee that the moral universalist's own moral framework is actually the universally correct moral framework. The claim that there is but one correct set of moral principles leaves open the possibility that the universalist themselves is mistaken about what this set contains, or even the possibility that nobody at all is in a position to say what this set contains. Needless to say, moral universalism will not guarantee access to the truth about relevant nonmoral matters either. Imputing false beliefs about nonmoral matters to others as an explanation of moral diversity cuts both ways; these others may be just as inclined to pay the erstwhile imputer back in her own coin. If the nonmoral matters in question are themselves controversial ones, with no readily apparent procedure available for deciding them, such as for example, the question of whether the fetus is or is not conscious, then little will have been achieved if each side declares that the other side has false beliefs about this. Harman's point is that at the very least the moral universalist must do more than merely impute false beliefs to others in order to give a convincing explanation of moral diversity.

2.5.3 Inner judgements.

There are two main parts to Harman's attempt to give what he considers the most plausible account of moral relativism. 1) There is the logical thesis that moral relativism applies only to those moral judgements he calls 'inner judgements'. 2) There is the claim that agents are motivated to act morally as a result of an agreement in intentions that arise due to what he calls 'implicit bargaining'. I will look at each of these in turn.

50

For Harman, an inner judgement is a judgement we can make about a person, "only if we suppose that he is capable of being motivated by the relevant moral considerations." (1982, p 190) Examples of inner judgements include a judgement that someone ought or ought not to have done something and a judgement that it was right or wrong of someone to have done something.

Harman draws a logical distinction between four different senses of the word "ought". These are what he calls the expectational sense, the rational sense, the normative sense and the moral sense (1982, p 192). For example, the sentence: "She ought not to have an abortion," may be interpreted in at least four different ways, employing each of the above four senses respectively. The sentence may mean that one would not expect her to have an abortion, or that it would not be in her own best interests to have an abortion, or that it would be a bad thing for her to have an abortion, or that she would be morally wrong to have an abortion. According to Harman, it is only the very last of these interpretations of this sentence that makes it an inner judgement, and it is this kind of judgement that Harman's theory of moral relativism solely applies to.

The distinction between the third and the fourth sense of "ought" is a fine one, but important for Harman. He explains it thus:

> The normative "ought to be" is used to assess a situation; the moral "ought to do" is used to describe a relation between an agent and a type of act that he might perform or has performed.

Harman also thinks that this is the case with the word "wrong". For Harman it is important to distinguish:

> using the word "wrong" to say that a particular situation or action is wrong from using the word to say that it is wrong *of someone* to do something. (1982, p 192, italics in original)

The point of the distinction seems to be to allay the universalist's fear that moral relativism can allow anything to be morally acceptable by giving Harman the conceptual room to say that, for instance, the situation that obtained in ancient Greece that allowed human slavery to be practiced was a morally wrong situation, but that despite this it was not morally wrong for individual ancient Greeks in that situation to enslave humans.

Harman singles out inner judgements as being susceptible to relativism because they imply that the agent who is the subject of the inner judgement has a reason (or reasons) to do or to refrain from an action and that the one making the judgement approves of this reason (and that the intended audience for the judgement also approves). So inner judgements may only legitimately be made relative to a set of shared reasons for acting. I may not make a genuine inner judgement (using the moral sense of 'ought') that someone ought or ought not to do something unless I (and also my intended audience) share the relevant motivational attitudes with the one about whom the judgement is made. This is the logical foundation of Harman's moral relativism. He analyses 'ought', when used in the moral sense, as a four-place predicate, in the following manner: " "Ought (A, D, C, M)," which relates an agent A, a type of act D, considerations C, and motivating attitudes M." (1982, p 194) The inclusion of M ensures that reference to shared motivating reasons for acting is built into the meaning of the moral sense of 'ought.' Although in this analysis Harman does not mention the notion of the framework that takes center stage in his later writings on relativism, a framework may perhaps be supposed to be constituted by a set of shared motivating reasons for acting.

Harman acknowledges that his analysis requires the following assumption: "I assume that the possession of rationality is not sufficient to provide a source for relevant reasons [for action], that certain desires, goals, or intentions are also necessary" (1982, p 194). However, even those who did not accept this assumption about the role of rationality in motivating action would

not need to reject Harman's analysis of inner judgements entirely, but would only need to insist that the inclusion of a reference to motivating attitudes is a redundant element, as speaker and audience will always necessarily share motivating attitudes if they are all rational. Someone with this view might want to insist that, strictly speaking, the word 'ought' when used in a moral sense is a three-place predicate and not a four-place predicate as Harman believes. Making this modification would allow Harman's analysis of moral judgements to be detached from his moral relativism and they could each stand alone as independent contributions to metaethics.

In any case, the metaethical issue concerning the role of rationality in moral motivation is distinct from the metaethical issue between the universalist and the relativist, and this is something I will discuss in more depth in part 3. For the sake of specifically evaluating Harman's version of moral relativism, I propose to accept the analysis of inner judgements given above.

2.5.4 Tacit bargaining.

Harman offers a hypothesis about the origins of the shared motivating attitudes that allow us to make legitimate inner judgements. He says that they arise as a result of implicit bargaining. (1982, p 196-8; 1996, p 21-4) In fact he says it is "essential" (1982, p 197) to his version of moral relativism that they be supposed to have originated in this way. A deal that is struck as the result of a bargain, for instance an agreement that a house will be sold at a certain price, is not considered to reflect any sort of universally correct reality about the value of the house. It is a pragmatic arrangement that allows a self-interested buyer and seller to do business and achieve their goals by reaching a compromise that both of them can live with. The buyer and seller of a house may use various argumentative strategies and appeal to various facts when bargaining over the price of the house, but an appeal to the price of the house will not be one of these, as the price of the house emerges only once the bargain has been struck. Harman stresses that the most effective bargaining

tool in such situations will often be a threat to withdraw from the bargaining process; the seller will refuse to sell if the price is too low, the buyer will refuse to buy if the price is too high.

According to Harman a similar sort of bargaining process is also how moral frameworks arise. There is no universally correct fact of the matter concerning what moral framework to follow prior to bargaining, and each of the self-interested parties to the bargain can, if the bargain is not to their liking, "threaten to withhold full participation in a moral framework" (1996, p 22). The resulting bargains form a sort of social contract that it is in each participant's self-interest to follow. Everyone agrees that there must be restrictions on some kinds of behaviour, such as murder, because the benefits to any self-interested individual in being able to commit murder at will are outweighed by the costs of the constant possibility that one could be murdered at will by anyone else.

Harman sees morality as "a compromise based on implicit bargaining." (1982, p 197) He believes that this is the best way to explain why "most of us assign greater [moral] weight to the duty not to harm others than to the duty to help others." (1982, p 196) His explanation is that because the implicit bargaining underlying morality is between people of differing degrees of power then, while all could agree to a moral principle forbidding harm to others because all would benefit from such a principle, only the poor and the weak would support a moral principle requiring everyone to help others, because only the poor and the weak would benefit from it. The rich and the strong would not support such a principle, and so implicit bargaining would result in a compromise whereby there was a robust moral principle requiring agents not to harm others, and a considerably more fragile moral principle requiring agents to help others.

This reflects Harman's conviction that morality has "a source in actual bargains, reached among actual people who have different powers and resources" (1996, p 63). This conviction suggests that Harman considers himself to be giving a causal explanation of why we regard it as more wrong

54

to harm others than not to help them. It is because our ancestors once struck a bargain to that effect some time before the beginning of recorded history, and since then social pressures along with other casual factors have entrenched this aspect of our morality. Furthermore, Harman seems to envisage this as an ongoing process. He says, "most people's values reflect conventions that are maintained by continual tacit bargaining and adjustment" (1996, p 22). He also believes that we may sometimes observe the process in action: "Recent changes in conventional sexual morality seem clearly to have been affected by moral bargaining" (1996, p 30).

Harman gives an example of what he thinks tacit bargaining is. The example is Hume's one of two rowers adjusting their rates of rowing to achieve a compromise rate (1996, p 22). However, while this may be a good example of *unspoken* bargaining, it is not clear that it is *tacit* bargaining. The physical circumstances that the rowers are in does not require speech in order for one rower to understand that the other rower is not matching their speed and to react accordingly, but there is some kind of nonverbal communication going on here. Other instances we might imagine that could be cases of unspoken bargaining also seem to have this feature. So, for instance, when people at a bus stop form a queue to get on the bus without having to speak a word it could perhaps be said that they are bargaining with each other to determine their place in the queue, but again this only seems to be possible because body language and other non-verbal means take the place of speech. Just because a situation is such that bargaining can take place without speech does not necessarily make the bargaining tacit. Two people could bargain without speaking to each other by writing messages on a blackboard, but this would not be tacit bargaining. Some indication is needed of just how implicit the bargaining would have to be before it would count as tacit bargaining.

Harman recognises that there are some clarifying questions faced by the proponent of a theory like his and he himself lists five of them:

1) If we have agreed, when did we do it?

2) Does anyone really remember having agreed?

3) How did we indicate our agreement?

4) How do those who do not want to agree indicate that they do not agree?

5) What are the consequences of their not agreeing? (1982, p 201)

Harman responds to these questions by stipulating that the sort of agreement he has in mind is not supposed to be some sort of datable ritual but rather an agreement in intentions: "there is an agreement in the relevant sense when each of a number of people has an intention on the assumption that others have the same intention." (1982, p 201). Harman says he is using the word "intention" here in an extended sense to indicate a certain disposition or habit, rather than a conscious intention. (1982, p 196). This takes care of questions 1-3. As it is an agreement in intentions it is not the kind of agreement for which we can give a date for the commencement of the agreement, or remember agreeing to, or indicate our agreement in any formal way. We just habitually act in the way we think is morally right (as far as we can) and assume that others will do likewise. This also answers the question raised in the previous paragraph; no communication of any kind occurs between the parties to the agreement. Participants in the bargaining process will not be aware that they are undertaking such bargaining, but will simply think that they are debating about what is morally right.

As for questions 4 and 5, Harman says: "Someone refuses to agree to the extent that he or she does not share these intentions. Those who do not agree are outside the agreement; in extreme cases they are outlaws or enemies." (1982, p 201). The way such outlaws are dealt with will depend on the details of the moral framework, although some sort of sanction will presumably be likely.

However, it is significant that Harman continually talks of agreement when he is answering these questions, and not bargaining, even though it is the notion of tacit bargaining that his theory is built upon. There is an

important difference between agreement and bargaining when it comes to answering questions 4 and 5. Those who are unsatisfied with an agreement are simply regarded as being outside it. Harman says: "in extreme cases they are outlaws or enemies." (1982, p 201). In contrast, those who are unsatisfied with a bargain, Harman says, "can threaten to withhold full participation in a moral framework unless their disadvantage is lessened or removed." (1996, p 22). With an agreement one is either a part of it, or outside of it, but in a bargaining situation things are less straightforward as there is a third possibility. One may either be a part of the bargain, outside of the bargain, or be threatening to withdraw from the bargain unless certain demands are met.

A difficulty with Harman's theory is that it is crucial for participants to be able to make a threat to withdraw from the bargaining process, but it is a mystery as to how such a threat could be delivered tacitly. It seems that explicit communication between the parties is a necessary element of any bargaining situation. Recalling Harman's earlier example, it would seem to be impossible for a buyer and seller to tacitly bargain over the price of a house. Part of the nature of bargaining involves making offers and counteroffers and it is difficult to see how anyone could do this tacitly, or tacitly threaten to withdraw from the negotiations if they are not getting what they want.

To see the difficulty, consider the following situation: Suppose I am a member of a club. I intend to pay my club fees so long as I believe that the others in the club also intend to pay their fees. Nothing to this effect has ever explicitly been stated, and so the club members could be thought of as being participants in a tacit agreement. So far, so good. Now suppose I become aware that myself and a few other members of the club pay far higher club fees than the rest of the members while receiving exactly the same benefits of membership as they do. I am dissatisfied with this arrangement and I (along with the minority of other club members who also pay higher fees) would like the agreement to be modified so that the amount we pay is in line with what the rest of the club members pay. Presumably bargaining is called for, and if necessary I and the others in the minority will threaten to not pay our fees at

all or even to leave the club if our fees are not reduced. But it is difficult to see how such a threat can be made in a tacit manner. Perhaps I should simply stop paying my fees or pay a lesser amount than I am currently paying and expect that the rest of the club members will tacitly understand what my grievance is. An appeal to intentions does not clarify matters. It seems that my intention has changed, or at least become more refined: I now intend to pay my club fees so long as I believe that the others in the club are charged fees comparable to mine and intend to pay these fees.

This problem for Harman also effects the other examples of rowing and queuing for a bus. Suppose one of the rowers decides she wants to rest for a while and wishes to strike a bargain with the other rower about this. The sensible thing to do would be to speak up and explicitly discuss a new agreement. Someone who pushes their way from the back to the front of a bus queue may have rejected the bargaining process, or they may only be tacitly threatening to reject it because they think it was unfair that they were at the back. It seems impossible to say which is the case without communicating with them.

Harman's notion of tacit bargaining as the basis for morality faces the difficulty that there does not seem to be any plausible way for bargaining to proceed tacitly. If we do not communicate explicitly with others about the bargain then we have no means of determining what aspects of morality they intend to bargain about nor can we determine when or why they may be threatening to withdraw from the bargain. Claiming that tacit bargaining is the basis for morality leaves it unclear whether, for instance, a particular murderer is an outlaw operating outside the bargaining process or whether they are someone who is provisionally withdrawing their commitment to morality as a part of the bargaining process. (In a recent paper Barbara Fried has raised related concerns about how the practical difficulties of exiting a social contract type arrangement tend to be suppressed when such theories are presented.)

It seems to me that the notion of tacit bargaining is problematic in so far as Harman does not explain how it is possible to conceive of bargaining occurring without explicit communication between the parties to the bargain. Yet Harman would not want to rely on the notion of explicit bargaining as the basis for morality because no explicit bargaining about morality takes place, at least so far as I can determine. For instance, the recent changes in sexual morality that Harman mentions do not seem to have occurred because people who were once thought to be immoral due to their sexual orientation threatened to start lying and murdering if they were not more fairly treated. Bargaining, either explicit or implicit, does not seem to play the role in morality that Harman believes it does. This makes his version of moral relativism, which is built on the notion of bargaining, less plausible.

2.5.5 Wong's criticisms of Harman.

Before moving on to examine the version of moral relativism offered by David Wong, I will look briefly at the objections Wong makes to Harman's brand of moral relativism. He has two lines of objection. Firstly, he proposes that in a situation in which no bargain about morality was in force, we would still have some moral obligations. Wong says: "Suppose that in the next Great Depression, the fabric of society unravels into a Hobbesian war of all against all. Many of us think we would still have the elementary duty not to kill others for amusement, even if we know that others had no intention of reciprocating." (Wong, p 24). According to Wong, Harman's theory is not able to explain where this duty would come from.

Harman might reply that it is doubtful whether the kind of extreme situation that Wong envisages, wherein every individual is set against every other individual, could ever really arise. Even the most divisive of civil wars will be between groups of people and not individuals, and any time people are interacting with one another in groups there must be some sort of an implicit agreement in force within the group, and that is all Harman needs for his

59

moral relativism to get off the ground (assuming he is able to find some way to account for the difficulties I raised in the previous section).

Wong could perhaps maintain that the situation he envisages is not supposed to be one that could ever arise in practice but rather that it is a logical possibility that society could completely fragment into warring individuals and yet that even in such an extremity the moral duty not to kill others for amusement would hold. This would perhaps be to imagine that some sort of individualistic relativism (see section 2.3) has taken over. Wong might think that it is this logical possibility that Harman's theory cannot account for. Perhaps it could not, although it is unclear just how much of a victory this would be for Wong to score. In setting up this extreme situation as something that a theory of moral relativism must be able to explain in order to be credible, Wong is in danger of defeat at the hands of the monster he created, as I argue below that his own theory is not able to explain this situation either (see section 2.6.2). I think it is fair to say that this situation is something that many moral theories would have difficulty explaining but that this has a lot to do with the incredible nature of the scenario that Wong creates. After giving the details of Wong's theory below I will ask whether it fares any better than Harman's in explaining this scenario.

Wong's second objection is that Harman's view does not take into account our propensity to criticise other groups for the agreements they make. This is reminiscent of Rachels objection that moral relativism implies that we could no longer say that the customs of other societies are morally inferior to our own (see section 2.4). In fact, Harman takes some pains to distance his moral relativism from such an implication: "I am not denying (nor am I asserting) that some moralities are "objectively" better than others or that there are objective standards for assessing moralities." (Harman 1982, p 190) In putting forward this objection Wong seems not to recognise that Harman's moral relativism is supposed to apply only to inner judgements. Harman's distinction between the normative sense and the moral sense of words such as "ought" (see section 2.5.3) is supposed to forestall the very objection Wong

60

raises. Perhaps Wong could have tried to cast doubt on the plausibility of this distinction of Harman's, but he does not appear to be aware that it has been made. In summary, it seems to me that neither of the criticisms that Wong puts forward against Harman's version of moral relativism are especially successful. Now I will turn to the type of moral relativism that Wong advocates.

2.6 Wong's moral relativism

2.6.1 Six claims.

David Wong, in his book *Moral Relativity*, advances a theory that he says "gives us the best explanation of moral experience." (p 1) It is "a theory built around the claim that there is no single true morality." (p 1) Like many others, Wong characterises the metaethical dispute about this claim in terms of conflicting views about whether morality is "objective" or "subjective". For my discussion of Wong I will recast what he says in the terminology I have already introduced in section 1.5, for the reasons that I give there. When I do use Wong's terminology I will place it in quotation marks to indicate that it is his.

Wong lists six claims on the first page of his book. How many of these claims one is willing to accept indicates the degree to which one is an "objectivist" or a "subjectivist" about morality. The claims are:

1. Moral statements have truth values;
2. There are good and bad arguments for the moral positions people take;
3. Nonmoral facts (states of affairs that obtain in the world and that can be described without use of moral terms such as 'ought,' 'good,' and 'right') are relevant to the assessment of the truth value of moral statements;
4. There are moral facts (that may or may not be claimed to be reducible in some way to nonmoral facts);
5. When two moral statements conflict as recommendations to action, only one statement can be true;
6. There is a single true morality. (p 1)

Wong rejects Claim 6, correctly identifying it as being the real issue that separates moral universalists from moral relativists. He recognises that a moral relativist may accept at least claims 1-4, while it would be possible for a moral universalist to deny them. Claim 5 seems to be the odd one out in this list as whether or not it was accepted or denied would seem to depend greatly on just how the two moral statements are supposed to be conflicting as recommendations for action. Wong appears to think that the two moral statements must contradict one another. The example he gives of two conflicting moral statements is: "You ought not to act coldly" and "You ought to act coldly" (p 2-3). Less implacably opposed pairs of conflicting moral statements seem to be possible however, such as the sort of situation that might be described as a moral dilemma, perhaps a case in which someone has made a promise that requires them to tell a lie, in which moral recommendations to action may conflict without strictly speaking contradicting each other. The acceptance or rejection of Claim 5 depends on the details, and will not necessarily separate the moral universalist from the moral relativist.

Wong commits himself to accepting Claims 1-4: that moral statements have truth values, that there are good and bad arguments for the moral positions people take, that nonmoral facts are relevant to assessing the truth values of moral statements, and that there are moral facts. But he rejects Claims 5-6: he believes that two moral statements recommending conflicting courses of action can both be true and that there is no single true morality. (p 7) Wong categorises himself as a moral relativist, but at the same time he wants to take on board some of what he thinks are the more plausible of the claims that moral universalists might make. Wong wishes to achieve a compromise between moral relativism and moral universalism by conceding some claims that are often central to an universalist account of morality while at the same time insisting that the nature of morality is ultimately relativistic. The strategy is to give the moral universalist almost all they could reasonably ask for but still in the end maintain a theory that would deserve to be called a

form of moral relativism. This seems to me a sound strategy and my attempt in this thesis to present a coherentist approach that makes room for both Conviction U and Conviction R shares some similarities with what Wong is trying to do, although the details of our two approaches differ markedly.

Wong, like Harman, is impressed with Conviction R and like Harman he critically examines some common universalist strategies for explaining moral diversity and disagreement. The universalist might say that in a moral disagreement one side or the other is mistaken about some fact or facts, that one side or the other has made errors in reasoning, or that the disagreement is beyond the powers of human reason to solve (Wong, p 151-3). Wong thinks that the universalist is not able to satisfactorily say just what the facts are that one side or the other is mistaken about. Nor does he believe that universalists are able to say just what errors in reasoning have taken place, or why exactly the disagreement is beyond human power to solve. "The general failing of the [universalist] methods is not meeting the principle of the best explanation. We are left with gaping holes in our theory of the speakers of moral language." (Wong, p 153) Wong here exploits Conviction R in very much the same way as Harman, by alleging that universalists are not able to give an acceptable explanation of why moral diversity occurs.

2.6.2 Wong's analysis of morality.

A central part of Wong's reconciliation project between the "objectivist" and the "subjectivist" is to make use of the developments in philosophy of language pioneered by Tarski, Quine and Davidson (p 5-6). Once Wong reveals just how he intends to do this it becomes obvious that his acceptance of the notion that moral statements have truth values was not merely a concession made to the "objectivist" in the spirit of reconciliation, but is in fact essential to Wong's own project. He wants to "present a truth conditions analysis of moral statements" (p 17) that will "show how 'ought' and 'good' contribute to the truth conditions of [moral] statements." (p 19). The moral statements he concentrates on are statements of the form: "A ought to do X,"

63

and also: "X is a good Y". This sort of approach is of course only feasible if we assume from the outset that moral claims are the sorts of things that can have truth conditions, and this will only be the case if they have truth values. This makes Wong a good example of an explicitly cognitivist moral relativist.

To support his theory Wong produces a causal account of the origin and nature of moral statements. I will briefly set out what he says about this and offer some criticism. Wong gives a schematic explanation of how morality arose:

> My thesis is that rules [without explicitly moral terms] were the first means by which people began to formulate and to recommend to each other actions or policies of action. At first, they may have used rules of the simplest form "A is to do X." As they became more sophisticated and reflective about the use of their rules, they noted conditions under which this or that policy of action was to be carried out, and they used rules of the form "If C, A is to do X." They went on to develop general rules of the form "If C, everyone is to do X," from which more specific rules about particular agents could be derived. It was useful to develop a language for talking about which rule applied to agents given certain conditions, or which rule applied to agents given certain rules applying to all agents. An "A ought to do X" statement may be interpreted as telling us (in part) that by not doing X, A would be breaking a rule under certain conditions that obtain. (p 37-8)

Wong goes on to say that as they mature groups or societies will develop different sets of rules that have different functions. Morality arises from "the set [of rules] for resolving internal conflicts of requirements (stemming from an individual's different needs, desires and goals) that affect others and for resolving interpersonal conflicts of interest in general." (Wong, p 38) It is disappointing that Wong does not include a few examples here to explain the transition from normative rules with no moral content to moral systems. This transition is discussed at a level of generality that seems to eschew providing concrete instances. Wong does not give us any particular reason to share his confidence that: "An "A ought to do X" statement may be interpreted as telling us (in part) that under certain conditions that obtain, A would be

breaking a rule belonging to that set of rules for resolving internal and interpersonal conflicts." (Wong, p 38)

Wong continues by considering how morality is to be distinguished from other nonmoral sets of rules, such as legal statutes, that are also at least in part devoted to resolving conflict. Out of all the possible systems of rules that are concerned with conflict resolution, Wong tries to narrow the field to those that could be said to be moral by specifying that: "true moral 'A ought to do X' statements are founded on moral systems that are adequate with respect to some ideal of morality." (p 39) He then introduces what will be for him a fundamental notion, that of an adequate moral system. Its adequacy resides in its meeting the standards for moral systems. Wong recognises that this reliance on standards for moral systems may possibly be viciously circular but he escapes this by stipulating that the standards for adequate moral systems do not contain the word 'ought' and are merely descriptions of the characteristics of adequate moral systems. "A complete list of such standards specify the ideal of morality" (Wong, p 40).

With this structure in place, Wong can present his analysis of moral 'A ought to do X' claims:

By not doing X under actual conditions C, A will be breaking a rule of an adequate moral system applying to him or her. (p 40)

Wong's analysis of 'X is a good Y' statements is:

Under actual conditions C, X satisfies those standards for Y's contained in adequate moral systems applying to X. (p 69)

These formulae allow Wong to present the claim that is the cornerstone of his book: "The extension of 'adequate moral system' could vary, as the term is used in different groups and societies." (p 44) Because different societies may use different standards to measure the adequacy of moral systems there will be

in these different societies varying extensions of the phrase 'adequate moral system'. For Wong this is the source of moral relativism; the extension of the phrase 'adequate moral system' will be relative to the society in which the phrase is being used.

Wong discusses some contemporary theories of reference in order to decide which of them would be most appropriate to determining the reference of the phrase 'adequate moral system' and also account for how the extension of this phrase could vary across different groups. He elects to follow a descriptive-causal theory of reference. Wong summarises the theory in this way: "certain descriptions associated with 'adequate moral system' fix the referent of the term, and the referent, once fixed, plays an appropriate role in a causal explanation of speaker's beliefs about it. In the end, we have a theory with causal and descriptive elements in it." (p 61) The descriptions that fix reference are weighted, in that some count more in the reference-fixing process than others, because they are more firmly associated with the phrase (p 60). Wong does not think that the process of reference-fixing takes place only at an initial stage, but rather he: "allow[s] for the reference to change at any time ... when the moral beliefs that embody reference-fixing descriptions (standards for the adequacy of moral systems) change." (p 63)

I propose not to question the plausibility of the descriptive-causal theory as a theory of reference. It would take me too far from the topic of my thesis to delve into the issues in philosophy of language that would arise in a complete discussion of what Wong says in support of his preferred theory of reference. However, there is one issue that I think is significant enough to be mentioned. There is an apparently unresolved tension between two different approaches to philosophy of language that Wong attempts to combine under the banner of his moral relativism. I will briefly indicate where I think the tension lies. Wong explicitly says he is following the Tarski-Quine-Davidson approach to language (Wong p 17-22) in looking to give a truth conditions analysis of moral statements, and yet this approach seems to conflict with his use of the descriptive-causal theory of reference to make decisions about the

extension of a single phrase. One of the central features of the truth conditional approach of Quine and Davidson is that the smallest semantic unit is the sentence, as it is sentences that are true or false and thus have truth conditions, and in their view it is fruitless to attempt to do the sort of semantic work with sub-sentential phrases that Wong is trying to do. (See for instance Quine (p 26-79) and Davidson (p 215-225, p 227-241.))

It is characteristic of the Tarski-Quine-Davidson approach to philosophy of language that a candidate for the extension of a sub-sentential phrase will emerge only gradually as a theory of interpretation is built up for a whole language by looking to assign possible truth conditions to utterances of sentences of that language, in the process Davidson calls radical interpretation. Even when this has been done, on such an approach there is no such thing as *the* extension of a word or phrase, because assigned extensions only represent the best guesses of the interpreter in an attempt to present the overall best theory of interpretation for the language as a whole. (Davidson, p 125-139) On this view we cannot in principle give the extension of a sub-sentential phrase. This is especially the case with an extremely abstract phrase such as 'adequate moral system', because the truth conditions of sentences in which it appears will not be apparent from observing the circumstances in which they are uttered. Quine and especially Davidson are eliminativists about the semantic relation of reference and it is difficult to see why Wong thinks he can graft the very non-eliminativist descriptive-causal theory of reference onto their truth-conditional approach to arrive at a theory of moral relativism. It seems that if Wong really was a follower of Quine and Davidson then he would see little point in the claim that the extension of the phrase 'adequate moral system' is relative to different societies because, strictly speaking, there is nothing there to be relative.

It might be replied that this is a misconception of the results of the truth conditional approach. Because on this approach the extension assigned to any given phrase in a language will always be relative to a theory of interpretation for that language and because such theories of interpretation may quite

evidently be relative to societies or groups, there is no difficulty with Wong's claim that the extension of the phrase 'adequate moral system' is relative to societies or groups. This may be true, but it will do nothing to assist the case that Wong is trying to make for moral relativism. If any given phrase in a language is relative to a theory of interpretation, and not just the phrase 'adequate moral system', then there is nothing special about this phrase from the point of view of the truth conditional approach to philosophy of language. But for Wong there is something special about this phrase. The extension of 'adequate moral system' is supposed to be relative to societies or groups in a way that the extensions of other phrases are not. That is what makes his theory of moral relativism discriminating relativism and not indiscriminate relativism (see section 1.3). Perhaps it would be possible for Wong to find a way to defuse this apparent tension between the different elements within his theory, but as things stand the apparent tension between Wong's desire to employ a certain kind of truth-conditional analysis of moral statements and his desire to make claims about the extension of a particular phrase is problematic.

Now that the details of Wong's theory are present it may be asked whether it can explain the situation that he thought was essential for Harman's theory to be able to explain. This was the scenario Wong created of the utterly fragmented population of warring individuals that never the less still have the moral duty not to kill one another for amusement (see section 2.5.5). It seems that Wong's theory faces difficulty explaining this situation. A population of warring individuals would apparently have no need for rules to resolve interpersonal conflict. As every individual is pitted against every other individual in a struggle for survival, interpersonal conflict will be the order of the day and resolution of it will be through use of force. Maybe there is some internal conflict that may only be resolved by a moral duty not to kill other individuals for amusement, but this would need to be spelled out by Wong.

Another option for Wong would be to say that the moral duty in question derives from the rules of conflict resolution that were present in the society

prior to its fragmentation into warring individuals. So the explanation of this duty would involve tracing its origin to an earlier pre-fragmentation epoch in which it was essential. A sort of intellectual inertia has allowed it to survive beyond the conditions for which it was appropriate. Of course, if Wong gives this sort of explanation then it would be open to Harman to do so as well. Harman could claim that a situation in which warring individuals somehow retain the moral duty not to kill one another for amusement, may be explained by tracing this duty back to an implicit bargain that was in force in an earlier epoch.

It may be wondered how good either of these explanations are but Wong's theory does not seem to be on any stronger ground than Harman's with respect to explaining this particular situation. In any case, the very extremity of the situation does not make it a very credible candidate for being a sort of test case that a moral theory must be able to explain. There is a strong inclination to say that in a "Hobbesian war of all against all" morality would have collapsed and the warring individuals would owe no moral duties to one another. Wong himself makes a remark that seems to indicate that he would also agree with this assessment: "A society in which each person is a moral community unto him- or herself does not have a viable morality" (p 64). Perhaps the reason why the scenario Wong invents is so difficult to explain is that it is quite far-fetched, even to its creator.

2.6.3 Moral relativism and tolerance.

It is worth considering in a bit more detail a claim Wong makes towards the end of his book. This is a claim that he recognises is a controversial (even "disreputable" (p 177)) one for a moral relativist to espouse. He says that he wants to advocate a version of moral relativism that is combined with a principle he calls the "justification principle", and that this version of moral relativism leads to tolerance of other groups that operate with differing extensions of 'adequate moral system' from one's own. The justification principle is a moral statement that Wong derives from Kant and it says: "one

should not interfere with the ends of others unless one can justify the interference to be acceptable to them were they fully rational and informed of all relevant circumstances." (p 181)

A critic might ask under what circumstances the justification that gives this principle its name could ever occur. After all, Wong has earlier criticised the universalist view that serious moral disagreement with other groups could be due to ignorance of facts or error in reasoning. (See above, section 2.6) It might seem that according to Wong's moral relativism, a serious moral disagreement that occurs because different groups subscribe to different extensions of the phrase 'adequate moral system', is not a matter of ignorance or error on the part of one or the other of them, and that such disagreement may proceed even when both groups are reasoning correctly and fully aware of all the facts. Neither side will be able to justify to the other any interference with their ends and so the justification principle is misleading as Wong has stated it. It would be more accurate for the principle to simply read: One should not interfere with the ends of others. (Perhaps with an explanatory note to the effect that 'others' means those in other groups that subscribe to a different extension of the phrase 'adequate moral system', because of course the justification principle is not concerned with possible interference with the ends of other individuals within one's own group.)

However, Wong might be able to reply to this line of criticism by appealing to his belief that morality arose as a set of rules for resolving internal and interpersonal conflicts. One may be able to justify interference with the ends of those in another group if one can show that there are some facts about conflict resolution that this other group is unaware of and that under the extension of the phrase 'adequate moral system' that this group cleaves to, the ends they are pursuing will increase internal and interpersonal conflict rather than resolving it. The other possibility is that one could show that there is some failure of reasoning on the part of this group that leaves them unable to recognise that the ends they are pursuing do not resolve conflict but rather increase it.

Wong admits that there is nothing about moral relativism or the justification principle in themselves that is likely to make a moral relativist particularly inclined to accept the justification principle. He has two major lines of argument to try to convince a moral relativist to accept it. One is that sympathy and compassion, in conjunction with the desire to be a part of a group in which there is mutual refraining from interference with the ends of others, may lead the relativist to adopt the justification principle. (p 185) This reply ignores the fact that the justification principle is supposed to regulate conduct towards *other* groups, of which one is not a part. As these other groups adhere to moral codes with which one disagrees, it is not at all clear what feelings may be held towards members of these other groups, and whether they would be sufficient to motivate accepting the justification principle rather than some other principle for dealing with other groups.

Wong's other reply is that it is no contradiction for a moral relativist to subscribe to the justification principle: "the justification principle is interpretable in such a way that the content of what it prescribes is independent of any proposition entailing the falsity of moral relativism" (p 185). It may be true that it is logically possible to combine moral relativism with the justification principle but a form of moral relativism that is merely logically possible seems to be a weaker goal than that which Wong is aiming for. There are many kinds of moral relativism that are logically possible and Wong himself thinks that something stronger is required in a plausible version of moral relativism.

This emerges in his response to the universalist's Conviction U, the worry that moral relativism may imply that any moral system, no matter how seemingly evil or repugnant, is as acceptable as any other. Wong's reply to this is to say, "let me point out that it is only in the logical sense of possibility that a group can conduct so arbitrary a determination of adequate moral systems" (p 74). He appears to be claiming here that a version of moral relativism that is merely logically possible may be overlooked or disregarded as not being a serious version, and yet when he discusses the justification

principle he thinks it is a point in his favour that it is logically possible to combine his version of moral relativism with the justification principle. Perhaps he is only making the modest observation that his version of moral relativism is able to be combined with the justification principle without contradiction, and that this is more than some versions of moral relativism might manage. However, Wong does look to be guilty of applying double standards with respect to the credibility of holding positions that are merely logically possible.

2.7 Arrington's criticisms of Harman and Wong

Although I have identified what I think are problems with some of the details of the theories of moral relativism advanced by Harman and Wong, it might be thought that suitably revised versions of their theories could still stand as credible opposition against universalism. In this section and the next I will examine two general lines of criticism against the approach they take. The first comes from Robert Arrington, who during the course of developing his own theory of moral relativism (a theory that I will examine more fully in section 2.9), criticises the versions presented by Harman and Wong.

Arrington argues that the positions of Harman and Wong slide into either what he calls, "objectivism" or "subjectivism". First I will look at what he says regarding objectivism and then subjectivism. Arrington divides objectivism into two distinct varieties. These are metaphysical objectivism, the view that: "true moral judgements describe a subject matter that is independent of the thoughts and feelings of finite sentient beings," and epistemological objectivism, the view that: "at least some moral judgements can be supported by evidence that any knowledgeable and rational individual would have to accept." He goes on to say that moral relativists "usually deny one or both of these objectivist theses." (Arrington, p 193)

The difficulty he sees for the theories of Harman and Wong is that: "both theories offer explanations of how the rules of a social group come into being, and these explanations make the rules susceptible of an objectivist

interpretation" (p 248). This is because for both Harman and Wong the relativity of morality is seen as being the result of different social groups allegedly following different moral principles due to their different social, geographical and other circumstances. Harman and Wong see moral principles as having non-moral goals. In their view the metaethical justification for a moral principle is that it fulfils a non-moral role, in Harman's case creating a stable society through tacit bargaining, in Wong's case resolution of conflict. This is thought by them to lead to moral relativism, because different groups in different situations will require different sets of moral principles to fulfil the same non-moral role. I will have more to say about this approach in section 2.9 below, when I give my own criticism of it. Arrington has his own avenue of criticism.

Harman's claim that moral principles arise through tacit bargaining and Wong's claim that moral principles arise in order to resolve conflict leads to the problem, as Arrington sees it, that: "these rule-generating activities are constrained by standards either of truth or of rationality" (p 248). Under the conception offered by Harman, there will be some moral principles that it is more or less rational to bargain towards, while under Wong's there will be some moral principles that undeniably do or do not effectively resolve conflict. To reject some of the moral principles arrived at through tacit bargaining or determined to be effective conflict resolution devices is to make manifest some sort of irrationality or error. Both Harman's and Wong's theories thus appear to be unwittingly supporting epistemological objectivism. Arrington believes that this is problematic for Harman and Wong.

For example, suppose that the members of a social group have arrived, through a process of tacit bargaining such as Harman envisages (assuming for the sake of argument that such bargaining is possible), at a moral rule stating that it is wrong to wear green. Given that this moral rule is the result of such bargaining, it seems to Arrington to be irrational for any moral agent within that social group to reject this rule. After all, any moral agent who wished to wear green without thereby doing anything morally wrong could have tacitly

73

bargained against such a rule. So, as tacit bargaining is supposed to provide both an explanation *and* a metaethical justification of the moral rule stating that it is wrong to wear green, Harman seems to Arrington to be endorsing epistemological objectivism. A moral agent may not reject the moral rule against wearing green without this rejection entailing that they are making an error or manifesting some sort of irrationality. Arrington says: "if moral bargaining produces a result that can be evaded only by being irrational, do we not have a basis for morality that is independent of custom and opinion – moreover a basis that is likely to yield the same set of principles, or largely overlapping sets, for most people?" (p 213)

But then, if Harman is an epistemological objectivist, it seems to Arrington to be implausible for Harman to simultaneously claim to be a moral relativist. He thinks that Harman is in effect trying to maintain that there is no one set of true moral judgements while also maintaining that it is irrational to reject a certain set of moral judgements. Analogous reasoning leads Arrington to the same critical result regarding Wong's theory. But now the strain of maintaining a metaethical theory that is supposedly morally relativist but also simultaneously epistemologically objectivist seems to Arrington to be too much for Harman's and Wong's positions to bear. He thinks it is the morally relative part of their theories that must give way. He says: "Both Harman and Wong then, picture rules in such a fashion as to convert what purport to be relativist theories into at least quasi-absolutist ones. If this is relativism, absolutists have little to fear" (p 249).

Possibly Arrington has tripped himself up with his own complex terminology in stating this conclusion (his terminology I explain more fully in section 2.9). The thrust of his criticism was advertised as being directed against what he perceives as Harman's and Wong's epistemological *objectivism*, rather than their "quasi-absolutism", a term that he does not explain. Furthermore, as detailed in section 2.9, Arrington himself explicitly accepts what he calls moral absolutism (p 195) while rejecting what he calls epistemological absolutism. If moral absolutists have little to fear from the

theories of Harman and Wong then it seems they have little to fear from Arrington's theory either.

However, putting this aside, Arrington's criticism of Harman and Wong is not as damaging as he thinks it is. As explained in section 1.5 above, the opposition between objectivists and subjectivists is quite distinct from the opposition between relativists and universalists. So a moral relativist may, for example, quite consistently be an epistemological objectivist. Moral relativists may usually deny objectivist theses as Arrington claims, or they may not, but they are by no means compelled on pain of inconsistency to reject either metaphysical objectivism or epistemological objectivism. Moral judgements may be independent of the thoughts and feelings of finite sentient beings and supported by evidence that any knowledgeable and rational individual would have to accept and yet for all that also be context-dependent. The friction that Arrington believes he has identified between the objectivism and the moral relativism in Harman's and Wong's theories seems to be illusory.

The other side of Arrington's criticism of Harman and Wong is that if they do not take an epistemological objectivist line then they must find themselves committed to some sort of subjectivism. If standards of truth or of rationality do not constrain moral agents in their tacit bargaining or their attempts to resolve conflict, then it seems that Harman's and Wong's versions of moral relativism will collapse straight into subjectivism. Tacit bargaining could lead to practically anything being adopted as a moral rule if the parties involved are free to strike irrational bargains, while conflict resolution will be

force them to abandon the idea of moral principles that are (metaphysically and epistemologically) objective and also relative.

2.8 The naturalist strategy in metaethics

2.8.1 The genetic fallacy

At this point I would like to present what seems to me to be a more credible general objection that attempts to cast doubt on the overall metaethical strategy followed by both Harman and Wong. The objection is that they are committing the genetic fallacy (see section 2.2). If this objection is correctly handled it is one of the most effective weapons of argument that the universalist has available, although it can be something of a double-edged sword. This is because the genetic fallacy is not confined to relativists alone. In this section I will examine how the genetic fallacy and its parent the naturalistic fallacy impact on the debate between relativists and universalists. I return briefly to the universalism of Rachels to show how a universalist may commit the genetic fallacy, before turning back to Harman and Wong to evaluate how damaging this line of criticism is to their theories.

Recall that in section 2.4 Rachels argues: "that there are some moral rules that all societies will have in common, because those rules are necessary for society to exist." (Rachels, p 493) He gives three examples of these rules, and argues that each of them are required in any sort of complex human society. The examples are: a general rule requiring that infants be cared for, a general rule prohibiting lying, and a general rule prohibiting murder. In each case he argues that society would fall apart or would not be possible if these rules were not generally adhered to. Rachels' arguments are, firstly, that a society that did not care for infants would not last, as the young would not survive to replace the elder members of the society as they died. Secondly, in a society in which people constantly lied to each other, "[c]ommunication would be extremely difficult, if not impossible,"(p 493) making any complex society out of the question. Thirdly, if murder were not prohibited then, "[e]veryone

would have to be constantly on guard. People who wanted to survive would have to avoid other people as much as possible," with the result that, "[s]ociety on any large scale would collapse."(p 493)

Now even if it were to be conceded that these three rules are, for the reasons that Rachels gives, instances of rules that all societies must follow if they are to continue to exist, establishing this would not support the conclusion that Rachels wishes to argue for. This is because his argument attempts to leap the fact/value gap in what seems to be an illegitimate manner. Rachels would not have demonstrated the metaethical conclusion that there are universal moral rules, even if it were accepted that it is a fact that all societies that survive have exactly the same moral rules. Whether or not a society survives or falls apart (or would survive or fall apart) due to the moral rules it follows is an empirical question, and deciding it does not decide the metaethical issue between the relativist and the universalist. Moral rules may still be context-dependent, even if certain moral rules are always followed by successful societies. If Rachels is correct about his empirical claim, then a society that does not follow the three rules he mentions will eventually fall apart, but this in itself would not show that such a society's rules were not moral rules or that moral principles are not context-dependent. Such a society would perhaps be unfortunate, but not necessarily immoral (or amoral).

It may be that some normative lesson can be extracted from the considerations Rachels advances, but if so his three rules would be universal prudential rules rather than universal moral rules. To achieve his goal of supporting moral universalism Rachels must mount a separate argument showing that it is morally wrong (rather than imprudent) for members of a society to act in such a way that the society is unlikely to endure. A society that did not care for infants would quite possibly die out, but this fact by itself does not show that it is morally wrong (rather than imprudent) not to care for infants. Similarly, it may indeed be difficult for anyone to feel secure in a society in which murder was not prohibited, but again this seems to say more about the prudential desirability of such a general rule, and does nothing to

establish that it is the universal moral rule that Rachels wants to say that it is. In the same vein, people who constantly asserted what was not the case would be incomprehensible to each other perhaps, but it is not clear from this alone that they would be immoral. Lying might seem on the face of it to be morally questionable because it involves deception rather than because it may result in a failure of communication. A disabled person who may have great difficulty communicating with others is not held to be morally wrong because of their difficulty in communicating.

The upshot of this is that it does not seem to be too much of a stretch to charge Rachels with committing something close to the mirror image of the fallacy he accused the cultural relativist of in section 2.4. He has failed to distinguish facts about the likely causal origin of the three rules he mentions from arguments that establish them as universal moral rules. This could be considered the moral universalist's version of the genetic fallacy; the causal genesis of a rule is mistakenly taken to provide a metaethical justification for it. It would be an open question to ask: 'This rule is necessary for society to exist, but is it a moral rule?' The genetic fallacy is certainly an error to be avoided. Rachels' argument is a good example of how the genetic fallacy is not restricted to relativists but sometimes finds its way into the reasoning of universalists as well. I now return to the discussion of Harman and Wong, to see whether they have committed the genetic fallacy as well.

2.8.2 Locating morality.

Both Harman and Wong see metaethics as something like a branch of or adjunct to science; it is the science of morality. A metaethical theory like moral relativism is something akin to a scientific hypothesis. Whether moral relativism is the case or not may be currently undecided, but a careful examination of the observable evidence will allow us to draw a reasonable inductive inference that will best explain the available data. The most plausible explanation of this data is likely to be that moral relativism is the case. According to this approach, any acceptable metaethical explanation must

take a naturalistic form that finds a place for moral principles and other features of morality within the causal chains of the world of non-moral facts.

Harman spells this out explicitly in his paper: "Is There A Single True Morality?" where he makes what is for him a significant distinction between "two different ways to do moral philosophy" characterised by "a difference in attitude toward science." (1984, p 29) His favoured approach is naturalism. The naturalist is a theorist who, "takes it to be of overriding importance in doing moral philosophy actually to attempt to locate moral properties." A naturalist explanation that succeeds will, "be able to locate value, justice, right, wrong, and so forth in the world in the way that tables, colors, genes, temperatures, and so on can be located in the world." (1984, p 33) Now it may be the case that a version of moral relativism could be constructed that was inspired by the details of Harman's theory and that did not require this sort of naturalistic explanation, but Harman himself thinks that being able to locate moral properties in the world in the way described above is the primary task of the metaethicist. He also believes that the naturalist approach provides significant support for his version of moral relativism. Conviction R, with its emphasis on explanation of moral diversity, would find a likely ally in naturalism.

The opposing approach, one that Harman calls "autonomous ethics", does not consider the main goal of metaethics to be this task of naturalistic location. Instead, according to the follower of autonomous ethics, "[w]e begin with our initial moral beliefs and search for general principles … an important aspect of the appeal of such principles will be the way in which they account for what we already accept." (1984, p 31) It is easy to see how someone impressed by Conviction U might be drawn to autonomous ethics. The "initial moral beliefs" would include those non-negotiable elements of morality that the universalist accuses the relativist of threatening. During the course of his paper, Harman sets the two metaethical strategies of naturalism and autonomous ethics against each other, noting how and where they disagree, and trying to fairly present the attractive features of each position, while in the

end declaring that he finds naturalism to be ultimately the more plausible of the two. Harman is not interested in the coherentist notion of trying to make two seemingly opposing strategies fit together in a way that does not privilege either of them.

Harman is aware of the sensitivity of the issue of reduction in metaethics. According to him, naturalism need not involve a "serious" reductive definition of moral terms in non-moral terms. However, the naturalist's metaethical project of the locating of moral properties will undoubtedly have to involve specifying in some way those non-moral properties of the world that constitute the moral properties, or with which the moral properties are identical. The naturalist metaethicist trying to locate moral properties in the world is always going to find it tempting to reduce them to non-moral properties. Harman himself acknowledges and gives in to this temptation. In discussing moral beliefs he says: "our having the moral beliefs we have can be explained entirely in terms of our upbringing and our psychology, without any appeal to an independent realm of values and obligations" (1984, p 32). Although he does not try to offer a strict definition of the former in terms of the latter, he accepts that this sort of explanation involves "naturalistic reduction" and that he is reducing moral beliefs to facts about upbringing and psychology (1984, p 33). For Harman's naturalist, moral properties such as values and obligations drop out of metaethical explanations as unnecessary, because they do not have an obvious naturalistic location. Harman does not try to give an entire explanation of a sample moral belief in terms of upbringing and psychology. The naturalist's point is that an entire explanation in terms of upbringing and psychology could always in principle be given, and this sort of explanation would satisfy the requirement that Harman considers so important in metaethics of naturalistically specifying location.

2.8.3 A naturalist argument for moral relativism.

Harman believes that the naturalist approach supports his moral relativism and he gives a two premise argument to this effect. The first premise he states

80

thus: "If [a] person does not intend to do something and that is not because he or she has failed in some … empirically discoverable way to reason to a decision to do that thing, then, according to the naturalist, that person cannot have a sufficient reason to do that thing." (1984, p 36)

Harman couples this premise with the further premise that for any moral principle that might be alleged to be universal, there will be some people who will not intend to follow this moral principle and this is not because such people have failed in some empirically discoverable way to reason to a decision to follow the principle. As an example of such people and such a principle, Harman cites, "certain professional criminals, who do not act in accordance with the alleged requirement not to harm or injure others." (1984, p 36)

The conclusion of the argument is that these professional criminals do not have a sufficient reason not to harm others. If they do not have a sufficient reason not to harm others then they cannot be said to have a sufficient moral reason to follow an allegedly universal moral principle. The upshot is that if we take a naturalistic approach to metaethics then moral universalism is not the case with respect to this principle. Harman believes that a similar argument could be run for any other principle the universalist might try to put forward as a candidate for a universal moral principle.

Harman provides some brief discussion of the two premises of this argument, in which he says that the follower of autonomous ethics is able to reject one or both of the premises by appealing to grounds other than scientific ones (in particular, assumptions about the nature of ethics), but that the naturalist appeals to scientific grounds only, and so is not willing to relinquish either premise. The resulting situation is presented as a stand off, in which each side seems to the other as if it were begging the question.

For the purposes of that particular paper this may be all that Harman is prepared to say, but the naturalist argument he gives for moral relativism would have to be more extensively defended if one's intention was to use it to decide the issue between the moral relativist and the moral universalist. For

instance, his first premise is a reflection of the earlier mentioned assumption about reasons for action that Harman takes as the starting point for his metaethics: "I assume that the possession of rationality is not sufficient to provide a source for relevant reasons [for action], that certain desires, goals, or intentions are also necessary" (1982, p 194) (see section 2.5.3). Demonstrating that this assumption is warranted would be part of a suitable defence of the naturalist argument for moral relativism. Furthermore, the second premise would be on shaky ground unless it could be convincingly argued that it was not possible to be a naturalist moral universalist. It could be considered to be a live question as to whether the universalist could present empirically discoverable moral principles grounded in universal facts about, for instance, the evolution of human nature, such that these principles would provide a naturalistically acceptable sufficient reason for any (human) moral agent to follow these principles.

However, I do not propose to discuss in depth the minutiae of this particular argument that Harman presents because it seems to me that the argument is just an instance of a more general style of argument for moral relativism that follows from taking the naturalist approach to metaethics. According to the naturalist approach, a naturalistic explanation must be given as to why we have the moral principles that we do, and this explanation must locate the moral principles in the world. In the case of Harman and Wong, this naturalistic explanation consists of presenting a more or less abstract causal explanation detailing the origin of moral principles. Harman has a story about the process of tacit bargaining that leads to morality and Wong has a story about the emergence of morality from a set of rules that had no normative content. A feature of this sort of naturalistic explanation is that it attributes to our moral principles some useful nonmoral function, so that the adequate fulfillment of this function will be the way to naturalistically locate moral principles in the world. This approach to metaethics I will characterise as functional role naturalism.

2.8.4 Functional role naturalism.

Functional role naturalism could be thought of as being the metaethical counterpart to the position in philosophy of mind known as functionalism. Roughly speaking, according to functionalism a mental state is distinguished in terms of the fulfillment of causal roles. For instance, a pain state might be said to be caused by tissue damage and the cause of avoidance behaviour. There is a distinction here between role functionalism and occupant functionalism (Papineau, p 807). According to role functionalism, a pain state is whatever fulfils the causal role given above. It might be the firing of a certain set of neurons or it might be electricity moving down a wire, the important thing is whether it fulfills the causal role, not what it happens to be. According to occupant functionalism, a pain state is the physical state that actually realises the causal role. To the occupant functionalist pain states are first-order properties, while to the role functionalist they are second-order properties. As the name that I have given it indicates, it seems to me that functional role naturalism is a species of role functionalism. Thought of in this way, functional role naturalism can provide theoretical backing for moral relativism.

According to functional role naturalism, moral principles receive their metaethical justification, their reason for being, by fulfilling a certain nonmoral function that is given in naturalistic terms. For Harman the function of moral principles is to provide sufficient reason to motivate action in accordance with them, for Wong the function of moral principles is to resolve conflict and achieve dispute resolution within a group. Following role functionalism, the set of principles that fulfills the requisite nonmoral function for a group will be the moral principles of that group, whatever those principles happen to be. It just so happens that the general rules going under the label of 'morality' that we at present follow in our society fulfil the function of moral principles for us, although different general rules could fulfil this function just as well for other groups. Functional role naturalism thus leads to moral relativism in circumstances where it is the case that

differing sets of principles fulfil the same nonmoral function for different groups. Fulfillment of a nonmoral function is the means of naturalistically locating moral principles in the world, so that whatever fulfills the function counts as a moral principle.

An analogy here would be relativism with regard to nutritional principles. The functional role approach works very plausibly here. Nutritional principles are justified by fulfilling the function of promoting health in the person who follows them. Because the physical makeup of some people's bodies is different, the nutritional principles that will fulfil the function of promoting health will be different for some people. For instance, for those who are diabetic, the nutritional principles that fulfil the function of promoting health will be different than for those who are not diabetic, and so nutritional relativism follows.

Taking Wong as an example to show how this would work in the metaethical field, if moral principles receive their metaethical justification by fulfilling the function of conflict resolution, then because the situation or composition of some groups or societies are different, the moral principles that will fulfil the function of resolving conflict will differ across these groups. For instance, the principle, 'Lying is morally wrong,' may be a principle that, for one reason or another, resolves conflict in one society but not in another society. According to functional role naturalism, 'Lying is morally wrong,' would fulfill the function of moral principles and thus be a moral principle in the first society, but would not fulfill the function of moral principles and thus would not be a moral principle in the second society. Hence, it would be morally wrong to tell a lie in the first society but not in the second society, and moral relativism has apparently been demonstrated.

The functional role naturalist's argument here would be a more general version of Harman's two premise argument mentioned in the previous section:

N1) Something is a moral principle only in so far as it fulfils a certain nonmoral function.

N2) For any allegedly universal moral principle, there will be some group or groups for whom the principle does not fulfil the function, but some other principle fulfils this function.

Conclusion) There are no universal moral principles.

Premise N1 is just the statement of functional role naturalism. Premise N2 is an empirical claim about moral diversity. The nonmoral functional role that moral principles are supposed to fulfil will vary from theorist to theorist. Although Harman and Wong make different claims about what this function is, they both try to give the nonmoral functional role of moral principles by providing a speculative causal explanation of how moral principles arose.

2.8.5 The naturalistic fallacy.

An objection to the above argument is that premise N1 is false. The use of functional role naturalism by moral relativists to support their moral relativism commits the naturalistic fallacy. It appears that it would be an open question to ask: 'Such-and-such a principle fulfils a certain function, but is it a moral principle?' (Or the opposite question: 'Such-and-such a principle does not fulfil a certain function, but is it a moral principle?') Because the functional role of a moral principle is determined (at least by Harman and Wong) by providing a speculative causal explanation of how moral principles arose, the version of the naturalistic fallacy being committed here is the genetic fallacy. Functional role naturalism (as Harman and Wong apply it) adopts a reductionist approach that conflates the distinction between causal explanation and metaethical justification. The moral principles we follow obviously arose due to some cause or causes and of course empirical investigation into just what these are will be an interesting and worthwhile exercise and may even have implications for metaethics. However, according to this objection, even if accurate causal explanations were available it would be going too far to seek to use them to extract a functional role that will in turn serve to provide metaethical justification.

A recent example by G. A. Cohen brings out what does seem to be a real distinction between being caused to accept a moral principle and being justified in accepting a moral principle. Take the moral principle, 'Abortion is morally wrong.' Someone may be caused to accept this principle only by having the particular experience of finding that she cannot bring herself to have an abortion. But even if she had not had this experience she might still have been justified in accepting the principle (Cohen, p 232).

The genetic fallacy is prominent because of Harman's and Wong's emphasis on causal explanation, but the more general concern is that functional role naturalism commits the naturalistic fallacy. It is of course the case that sometimes empirical discoveries will make us revise our moral behaviour and that facts have an effect on our values. The discovery that an activity previously thought to be benign, for instance smoking in an enclosed space, is in fact dangerous to others, may lead to a revision in moral standards or values concerning smoking. But the objection being considered is that functional role naturalism has a far more serious implication than that. If fulfilling a functional role is what *makes* a principle a moral principle, then any principle that is found to fail to fulfil this role is thus not a moral principle. The result is that in following such a discredited principle we are mistakenly treating something as a moral principle that is in fact not a moral principle at all. Functional role naturalism leaves open the possibility that it could 'turn out to be the case' that, for instance, lying is not wrong, as it turned out to be the case that the earth is not flat. The functional role naturalist's strategy of tying the justification of moral principles to the function they fulfil leaves our current moral principles hostage to fortune.

It is merely contingent that a principle such as 'Lying is morally wrong,' fulfils the function that theorists like Harman and Wong speculate that it does. Furthermore, just because this principle appears to currently fulfill a certain nonmoral function in our society, this is no guarantee that it will continue to do so. Taking Wong as an example again, it is possible that new empirical research by sociologists or psychologists or the like could demonstrate that in

fact an apparent moral principle such as 'Lying is morally wrong,' no longer fulfils the function of resolving conflict. New research may show that lying as a matter of fact resolves conflict far more effectively than does telling the truth. According to functional role naturalism such a discovery would show not only that we must now be suspicious of the principle 'Lying is morally wrong,' if we want to act morally, but furthermore that we are at present mistaken if we continue to believe that this principle is a moral principle at all, and that we may have been mistaken all along in treating it as a moral principle.

This emerges when we look at how the functional role approach works in the field of nutrition. There are no qualms here in allowing that if new empirical research shows that an apparently sound nutritional principle in fact does not fulfil the function of promoting health in the person who follows it, then we should not only abandon that principle, but accept that we were regrettably mistaken all along in following it. The chance of wholesale error about nutritional principles is an unfortunate but accepted possibility that is a part of putting forward such principles. The nutritional principles we follow are acknowledged to be correct only so far as we currently know. This would also be the situation with moral principles under functional role naturalism. Our moral principles will always have an air of uncertainty about them. They must always be thought of as 'our current moral principles so far as we know', because it will always be a contingent matter that they as a matter of fact fulfil the function that justifies them as being moral principles.

Under this regime moral discussion would take on a completely different character. Our moral principles must continually be reexamined as conditions change. A question such as: 'Is 'Lying is morally wrong' still a moral principle or not?' will have to be asked in the same spirit as a question about the current exchange rate, with the expectation that today the answer may not be the same as it was last week. Those who follow out of date moral principles are not just being old-fashioned, they are morally wrong, so there will be a high premium on having access to the very latest empirical research about

matters surrounding the nonmoral function that moral principles fulfil. The functional role naturalist might say that this would actually be a superior way to conduct moral discussion, and any uneasiness about it just exposes a commitment to what Harman labeled autonomous ethics.

There is a parallel between this criticism of functional role naturalism and the sort of criticism of consequentialism that proposes a situation in which certain consequences, such as maximising pleasure for example, are achieved and yet it still seems that something morally wrong is occurring. We might be asked by the critic of consequentialism to imagine a society that contains a racial minority that is so despised by the majority of the population that the pleasure that large numbers of them would gain from mistreating this minority would outweigh the displeasure suffered by this minority as a result of the mistreatment. The consequentialist may simply bite the bullet here and say that despite what our intuitions may tell us, to mistreat the minority is actually not wrong in this situation. Likewise, the functional role naturalist may bite the bullet and insist that the moral principles we are currently committed to could turn out to be in error (although I am not suggesting that the functional role naturalist has anything else in common with this sort of consequentialist).

To me it seems more plausible to say that empirical discoveries about conflict resolution or any other non-moral function that might be alleged to justify morality are not the sort of reasons that would force us to reject a principle like, 'Lying is morally wrong.' The moral principles we follow may or may not resolve conflict more effectively than anything else, but this is irrelevant to the question of their metaethical justification. So, while these considerations may not be conclusive, I believe they raise difficulties for the plausibility of the naturalist strategy followed by Harman and Wong. In the next section I will look at a version of moral relativism that does not take naturalism as its starting point.

2.9 Arrington's conceptual relativism
2.9.1 Introduction.

Robert Arrington develops a version of moral relativism that is very different from the versions of Harman and Wong. He calls it conceptual relativism. In a nutshell, this is the view that moral principles are relative to a concept of morality. Arrington's conceptual relativism is something of a hybrid in that it accommodates universalism to the extent that he believes that there is one and only one universally correct set of moral principles. He calls these "rules of moral grammar" (p 269). An example of a rule of moral grammar is: 'Lying is morally wrong.' According to Arrington, although the rules of moral grammar are universal, they may not necessarily be subscribed to by everyone, and a lack of subscription to these rules does not indicate any sort of irrationality or any error about factual matters. I will explain the notion of a rule of moral grammar more fully below, after first explaining some of the terminology Arrington uses in situating his conceptual relativism.

Arrington employs a series of complex distinctions to create a theoretical space within which he positions his theory and from which he opposes other metaethical theories, including other sorts of moral relativism. He distinguishes different varieties of absolutism from different varieties of objectivism, and further distinguishes all of these from different varieties of subjectivism (p 193-199). I will not rehearse all the details of these divisions, but merely present the ones that are most relevant to my discussion.

The most important term that Arrington introduces is "nonobjectivism" (p 193), as his conceptual relativism is presented as a variety of nonobjectivism. The nonobjectivist is distinguished by what he rejects, and that is both moral objectivism and moral subjectivism. I have already shown above how Arrington's preferred line of criticism against other moral relativists such as Harman and Wong is to argue that their positions slide into either objectivism or subjectivism (see section 2.7). Despite their less than compelling nature, Arrington believes his criticisms of Harman and Wong are sufficient for him to reject the entire foundation that their versions of moral relativism are built upon. For Arrington, the approach taken by Harman and Wong is too close to epistemological objectivism for a moral relativist to credibly hold, and

Arrington is determined to come up with a theory of moral relativism that is nonobjectivist.

When characterising the opponent of the moral relativist Arrington does not at first glance seem to recognise that a distinction can be drawn between moral absolutism and what I have called moral universalism. Instead he lumps together a commitment to "absolute moral truths" with a commitment to "universal standards of moral value true for all human beings," (Arrington, p 193) and calls that moral absolutism. It is only the second sort of commitment, a commitment to what I have called moral universalism, that necessarily makes one an opponent of the moral relativist. As I noted in section 1.5, it would not be inconsistent to hold that moral truths were absolute by virtue of being unable to be overridden, while also at the same time being relative, or context-dependant.

However, Arrington is perhaps aware of this possibility as he later distinguishes between three forms of absolutism. These are moral absolutism, epistemological absolutism and metaphysical absolutism (p 194). The crucial distinction for Arrington is between moral absolutism and epistemological absolutism. What Arrington calls moral absolutism is close to what I have called moral universalism, that is, the commitment to the idea that there is one and only one universal set of moral principles, whereas what he calls epistemological absolutism concerns the question of whether these moral principles could be rejected without error or irrationality. This is a vital distinction for Arrington to make as his position is to accept what he calls moral absolutism while at the same time rejecting what he calls epistemological absolutism (p 195).

He thinks that different groups or individuals may be completely aware of all the relevant nonmoral empirical information and be completely rational, but be using a mistaken concept of morality that can lead them to reject the universally correct set of moral principles (the rules of moral grammar). Clearly Arrington must be committed to a reasonably sharp distinction between nonmoral factual information and what might be called moral

information (for instance, what the one correct concept of morality is) for him to be able to claim that a person or group could be in possession of all the relevant information of the first kind, and yet not be in possession of information of the second kind.

The role/occupant distinction made above in section 2.8.4 is useful in bringing out the significant difference between Arrington on the one hand, and Harman and Wong on the other. One way to put it is to say that while Harman and Wong focus on the functional role played by moral principles, believing that the occupant of this role may vary, Arrington focuses on the occupant of this role, believing that it may not vary. For example, according to Arrington the ancient Greeks were morally wrong to keep human slaves, although they were not making any factual errors or being irrational when they practiced human slavery. There was a concept that occupied the role in ancient Greek society that the concept of morality occupies in our society today, but the ancient Greek's concept was not the concept of morality. They may have reasoned impeccably about human slavery, but because for them a concept other than morality occupied the role of the concept of morality, they inevitably came to the incorrect conclusion about the moral status of human slavery. (p 252) I say more about this example below.

Looking at another example, one of the two opposed sides of the abortion debate is morally wrong, although this side is not thereby making any factual mistakes or being irrational. As each side has a different concept occupying the role of the concept of morality, they come to different conclusions about the moral status of abortion. However, only one side (or perhaps neither side) actually has the concept of morality occupying the relevant role, so at most one side draws the correct conclusion about abortion (although Arrington himself is circumspect about which side he thinks this is (p 290)).

2.9.2 Moral principles have moral goals.
One of the most distinctive features of Arrington's conceptual relativism is his insistence that the point of moral principles is not to resolve conflict or create

a stable society or any other sort of non-moral goal, but rather it is simply to encourage actions that are morally right and discourage actions that are morally wrong. As he puts it: "moral rules have moral ends, and whether in fact they promote still other ends is irrelevant to their justification" (p 250). He asks a number of rhetorical questions designed to build up the credibility of this assertion: "Do not moral rules have their own intrinsic end – namely, the promotion of morally correct behaviour? Many societies have rules prohibiting cruelty, but is this because cruelty is seen as promoting conflict and instability or because it is seen as an immoral way for one person to treat another? Is injustice denounced simply because it upsets a balance of interests or rather because it is unfair and wrong?" (p 250).

This is a major tenet of Arrington's conceptual relativism, the doctrine that moral principles have moral goals. The key claim here is that moral principles do not stand in need of justification in terms of some other non-moral goal in order for them to achieve their purpose as moral principles. Now there are at least two different ways that such a conception of moral principles could be understood. One could say that it is not *necessary* for moral principles to receive a non-moral justification in order to achieve their purpose as moral principles, or one could go further and say that it is not *possible* for moral principles to receive a non-moral justification. Arrington's commitment to develop a nonobjectivist version of moral relativism leads him to adopt the second, far stronger and more stringent interpretation of the formula that moral principles have moral goals. As it was the spectre of objectivism through the rational justification of moral principles that he believed was so awkward for Harman and Wong when critiquing their versions of moral relativism, Arrington asserts that moral principles must be "relative to something for which a rational justification is not possible" (p 250). If it is not possible to give a rational justification of moral principles then there is no danger of falling into objectivism and Arrington's goal of producing a nonobjectivist version of moral relativism will not be jeopardised. The implications of the claim that it is not possible (as opposed to not necessary)

92

to justify moral principles resonate throughout Arrington's theory and create what I consider to be some strongly counterintuitive implications. These will be addressed below

First I must fill in some further details of Arrington's notion of rules of moral grammar. Arrington paints a picture of the nature of moral discourse that derives from broadly Wittgensteinian notions about language that Arrington relocates to the field of metaethics. The most notable of these is Arrington's claim that the moral principles that jointly define our concept of morality are "rules of moral grammar" (p 269). He also characterises them as "rules for talking and thinking about morality" (p 275), and says: "The rules of moral grammar are universal, a priori, and necessary, not contingent, debatable generalizations that are matters of mere belief" (p 291). Because of the significance of this idea for Arrington, I will go into some detail in explaining it.

2.9.3 Rules of moral grammar

When Arrington talks about moral principles being rules of moral grammar he takes these rules to be defining assumptions forming the linguistic background that makes moral discourse possible. A rule of moral grammar such as: 'Lying is morally wrong,' is not something the truth or falsity of which can be argued about using moral language, as rules of moral grammar provide the very grounds for moral discussion. "The rules of moral grammar ... function as an element of the presupposed background of moral discourse and not, for the most part, as ingredients in its content" (p 282).

This is not to deny that it may be possible for a moral dispute to arise over whether an agent is permitted to lie in a particular situation. This might occur because the two sides in such a dispute may disagree over just what factors are most important when applying the moral principle in this situation, whether or not the principle may properly be overridden or outweighed by countervailing considerations, and so forth. What they are not disagreeing over, however, is the general applicability of the moral principle dictating that it is wrong to lie.

(p 271). It is a commitment by both sides to the rule of moral grammar embodied in this principle that is what allows the dispute to even proceed; common agreement on the moral principle forms the background without which the disagreement could not meaningfully arise.

There is no difficulty with imagining a case in which it was arguably the morally correct thing to do to tell a lie, perhaps because some other moral rule is judged to take precedence on this occasion. However, making such a decision in a case like this would not amount to a rejection of the general moral principle: 'Lying is morally wrong.' It would instead be a reluctant recognition that "unfortunately, we had to tell a lie on these occasions, that is, we were forced to do something that remains morally dubious" (p 272). Arrington notes that while moral disagreement may arise over a wide variety of topics, we never disagree over whether it is in general morally wrong to lie, or to break promises or commit murder (p 273).

According to this conception of moral discourse, we cannot reject the rules of moral grammar within the context of a moral discussion because these rules are the presuppositions that enable us to use moral language. One could not participate in the practice of moral discourse while at the same time rejecting the indispensable background assumptions that make this discourse possible. Someone who rejected the moral principle that says that it is wrong to lie would not be making a move within the moral disagreement described above, they would instead be showing by this rejection that they are, as far as Arrington is concerned, operating with a concept other than morality. The rules of moral grammar must be accepted as presuppositions by those who wish to engage in moral discourse, and a rejection of these rules amounts to a refusal to engage in moral discourse. Any such rejection does not threaten the appropriateness of the rules themselves as the foundation of moral discourse, but rather places the one who rejects them outside the domain of moral discourse. This would not prevent them from participating in other sorts of normative discussion such as prudential discussion, but participation in moral discourse would not be open to them.

Arrington's analogy between moral principles and grammatical rules could do with some clarification. Presumably the idea is that when it comes to a grammatical rule such as where to place the possessive apostrophe, there is no real answer to the question: 'Why is the apostrophe placed where it is in this case?' The answer is simply: 'Because that is what the rule dictates.' There is nothing further that needs to be said, or that can be said. No justification of the placement of the possessive apostrophe is necessary, and nor is it possible. Someone who does not place the apostrophe where the grammatical rules dictate is simply being ungrammatical. This sort of situation with respect to morality is what Arrington is trying to create with his idea that moral principles are rules of moral grammar.

When it comes to articulating what the rules of moral grammar are that constitute morality, Arrington says the following: "Morality has to do with personal autonomy and integrity, respect for persons, avoidance of harm to persons, and similar notions" (p 252). He is ready to acknowledge that there is some vagueness around the edges, and backs off from an extensive analysis of our concept of morality, but he does give some examples of what he considers some of the most important rules of moral grammar. These are: "It is wrong to tell a lie.", "One ought to keep one's promises.", "One ought not to harm other people.", "It is wrong to treat others disrespectfully.", and "One must not take the life of an innocent person." (p 302).

Arrington also distinguishes between rules of moral grammar and what he calls "substantive moral principles" (p 276). These latter are "moral principles and rules concerning which there can be legitimate moral disagreement" (p 276). Examples of these are: "Premarital sex is wrong.", "One ought to be patriotic.", "Mercy killing is wrong.", "One ought to lead a simple and frugal life.", and "One ought to give a quarter of one's income to charity." (p 302). The moral discussions that take place when we disagree about substantive moral principles require the rules of moral grammar as background assumptions. (Substantive moral principles should not be confused with moral principles, the term that I introduced in section 1.4. The latter is a general term

95

that covers both Arrington's rules of moral grammar and his substantive moral principles.)

Furthermore, it is sometimes the case that people will take rules of moral grammar as substantive moral principles. Arrington gives as examples: "One ought never to tell a lie.", "One ought always to keep one's promises.", and "One should never kill another human being, in war, peace, or self-defense." (p 302-3). The difference between a rule of moral grammar such as "It is wrong to tell a lie." and a substantive moral principle like "One ought never to tell a lie." is that the former may not be disputed by those who wish to engage in moral discourse, whereas the latter certainly may be. The former identifies a morally relevant dimension of life that must always be taken into account whenever moral discussion occurs, but that does not preclude the possibility that sometimes we may be morally required to lie. The latter is the contingent product of moral discussion and debate, and its rejection need not mean the rejection of morality. Arrington puts it thus: "The grammatical rule is unexceptionable, but the substantive rule requires considerable defense" (p 278).

According to Arrington, understanding the rules of moral grammar is not only necessary to being able to participate in moral discourse, but is also necessary to even understanding what morality is. According to his conceptual relativism there is no logically prior grasp of morality that precedes a comprehension of the rules of moral grammar. "One does not understand *morality* by grasping a general definition of it; one understands it by knowing that we are morally obligated to tell the truth and keep our promises, as well as avoiding harming others and to respect them" (p 283, Italics in original). Morality is what he calls a determinable notion, defined through its determinants (p 284). Arrington draws an analogy with the concept of colour. The following two paragraphs are worth quoting in full as an account of Arrington's notion of rules of moral grammar:

> [Rules of moral grammar] can fruitfully be compared to grammatical
> rules like "Red is a color." The latter defines 'red' by saying that it

designates a color. But equally it partially defines 'color'. ...[W]e have no general concept of color allowing us to discover that red is a color. Color is a determinable property and one that is defined through its determinants – color is red, yellow, green, blue, and so on. These definitions jointly define color and thereby introduce us to the color dimension of experience.

Likewise, "One ought to keep one's promises" and "It is wrong to tell a lie" simultaneously serve to define, on the one hand, 'Keeping one's promises' and 'lying' and, on the other hand, the moral notions of obligation and wrongdoing. "Morally what we ought to do" and "what it is right for us to do" – as well as "what we ought not to do" and "what it is wrong to do" – are determinable notions. They have no identifiable content in addition to such notions as keeping one's promises, telling the truth, breaking promises, lying, and so on. (p 283)

The rules of moral grammar embodied in the tradition Arrington appeals to together constitute the concept of morality. According to Arrington, this is the *only* concept that deserves the name 'morality', and other sets of principles do not and indeed could not count as alternative moralities. Other sets of principles would not have anything at all to do with morality, because the meaning of that term is exhausted by the list of moral principles that make up the concept of morality. Due to the nature of the concept, the single correct morality has no rivals and can have no rivals. "Our moral system constitutes morality: nothing else does." (p 254) While we may have normative discussions other than moral ones, and even consider that the results of these discussions are more compelling when it comes to deciding on our actions, the rules of moral grammar can have no other challengers within the field of morality.

Furthermore, as noted in section 2.9.2, it is a crucial claim for Arrington that it is impossible to justify the rules of moral grammar. For the nonobjectivist such as Arrington no rational justification is possible and to try to look for it is to misunderstand the nonobjectivist nature of the concept of morality. According to Arrington, our concept of morality "has no ground in truth and reason. It is beyond proof and beyond refutation" (p 257).

To the conceptual relativist, those who do not subscribe to the rules of moral grammar are employing a different concept of morality to us, or to be

97

more accurate, because he believes that, "the notion that there could be alternative moralities basically makes no sense," (p 287) those who do not subscribe to the rules of moral grammar are using a concept other than morality when they make certain normative judgements. He thinks that while moral condemnation of all those who do not subscribe to the rules of moral grammar is the proper and suitable response, this condemnation does not mean that it is possible for the rules of moral grammar to be in any way rationally justified to those who do not follow them. Nor does it mean that those who do not follow or who reject the rules of moral grammar are guilty of any sort of irrationality or error. Arrington's conceptual relativism, as an explicitly nonobjectivist theory, does not allow for the rational justification of the rules of moral grammar, so those who follow these rules are not doing so because they have been rationally persuaded to.

As noted earlier, the conceptual relativist believes that the Ancient Greeks, for instance, may have been acting immorally by keeping slaves but they were not acting irrationally or making any sort of logical or factual mistakes, and it would be impossible for us to reason with them about the moral status of human slavery. In Arrington's view those who do not follow the moral rules have not made any errors; "they are simply different, and their lives revolve around a different concept of morality" (p 256), although as I have noted above, it seems that Arrington should be saying here: "their lives revolve around a concept other than morality." Groups with different forms of life to ours may operate with concepts other than morality and there is no rational way of conducting any sort of dialogue with them about this. This strong commitment to nonobjectivism makes Arrington's conceptual relativism a most distinctive form of moral relativism.

Although there is no way that we can rationally justify the rules of moral grammar to those who do not subscribe to them, the conceptual relativist is by no means tolerant of such others. Arrington says: "we are not in logical disagreement with them, for we do not share the same logical or conceptual space. Our practice is pragmatically incompatible with theirs – ours commits

us to actions that physically cannot take place if the actions required by their practice are to occur" (p 261). He is blunt about the possible outcome of any disagreements we may encounter with those who do not subscribe to the rules of moral grammar: "To make our way of life possible, it may be the case that theirs must go. The result would be open conflict, perhaps war." (p260). It may be possible for the conceptual relativist to ignore those who are mistaken about morality so long as there is no possibility of interaction between them and us, for example our society and that of the ancient Greeks. But if there are disputes within our own society between those who subscribe to the rules of moral grammar and those who do not then there is no possibility of any sort of compromise.

Furthermore, under conceptual relativism there are severe restrictions on what forms of amoralism are possible. According to the conceptual relativist there could not be an amoralist who regarded morality as some kind of error or illusion while still having an intellectual understanding of what was involved when one lied, broke a promise, and so forth, because to understand these notions is to realise that they are morally wrong. It is not just that Arrington's conceptual relativist would disagree with such a form of amoralism, but that he thinks such a position could not even be possible. So an error theorist such as Mackie would be attempting to adopt a position that according to the conceptual relativist it is impossible to adopt. The only sort of amoralism possible would be a kind of individualistic amoralism according to which one placed oneself outside morality (presumably by rejecting the rules of moral grammar) while at the same time accepting that actions such as lying and so on were indeed still morally wrong for everyone else who had not rejected the rules of moral grammar.

2.9.4 Criticism of Arrington

I think that Arrington has achieved some genuine insight into the nature of moral discourse with his notion that the rules of moral grammar are presuppositions that must be accepted as a condition of participating in the

practice of morality. I employ this idea of presuppositions myself with the notion of basic moral principles that I introduce below in section 3.3.2. However, I disagree with much of the detail of Arrington's theory. In this section I will point out what I think are two of the most serious flaws in his conceptual relativism, which in my view remove it from contention as a plausible way of dealing with the conflict between moral relativism and moral universalism.

Criticism 1) Conceptual relativism is inconsistent.

The claims about morality being a determinable notion that were made in the two full paragraphs quoted in section 2.9.3 seem to indicate that there is no room for the fact/value distinction within conceptual relativism. For example, someone who might seem to understand a factual description of lying – that it is knowingly deceiving another – but who does not believe that lying is morally wrong, apparently does not really understand what it is to lie. Those who operate with a concept other than morality do not have what might be called 'a different perspective' on lying, murder, and so on – it appears they literally do not understand what these activities are.

This is inconsistent with the epistemological nonobjectivism that is such a vital part of conceptual relativism. Epistemological nonobjectivism entails that the rules of moral grammar may be rejected without error or irrationality. But for morality to be a determinable notion implies that for someone to believe that lying, for instance, was not morally wrong or was morally neutral would mean that they were not only making an error about the contents of the concept of lying, but were also making an error if they believed that they understood what lying was.

Criticism 2) An implausible stance towards others.

Arrington's moral relativism is far from a conventional version of the view. Some of the claims that the conceptual relativist is committed to are claims that a traditional moral relativist would probably be horrified to hear. These

include the claims that those other groups who do not accept our rules of moral grammar are acting immorally whenever they breach these rules no matter what they may think about the rules, or even whether they are aware of them at all. What is more, as such groups cannot be morally reasoned with, it may be that the only appropriate response we can make when encountering such groups is to attempt their extinction. This is a far cry from the moral relativist as traditionally thought of, someone who advocates tolerance of groups that follow different moral rules to ours, because 'our moral rules are correct for us, and their moral rules are correct for them'. Of course, just because Arrington's conceptual relativism develops in a non-traditional direction need not necessarily make it suspect. In his defence Arrington could point to the difficulties that a moral relativist such as Wong has in trying to incorporate a serious commitment to tolerance into his theory of moral relativism.

Still, there is far more than just tradition that a critic of Arrington might appeal to in expressing their alarm at the implications that his view has for our treatment of those groups that may reject or not be aware of the rules of moral grammar. The major problem is not that Arrington classifies the so-called moral practices of such groups as not really being moral, but rather that it seems we cannot even have an understanding of what it is that these other groups are doing when they follow such practices. The difficulty lies in how we are supposed to be able to recognise or identify any practice as being what we might, until we encountered conceptual relativism, naively have wanted to call an alternate morality.

Consider the normative concerns of the ancient Greeks and the question of whether these can be identified as constituting a moral practice. The conceptual relativist seems to have no easy way of talking about what up until now seemed to be merely the exercise of a morality that differed from our own. Arrington says, "We can grant a metaethical statement to the effect that other people like the Greeks *mean* something else by 'morality' than we do" (p 252, his italics). But this is a rather dubious claim. Obviously the ancient

Greeks did not mean anything by the English word 'morality', because they did not use this word. Translations into English of ancient Greek texts such as Aristotle's *Nicomachean Ethics* will contain English moral vocabulary, and while it may be true that something is irrevocably lost in translation when this occurs, Arrington's claim is far stronger. In his view it would be a mistake to use our concepts of morality, right and wrong, and so on, in such translations, because according to his conceptual relativism the Greeks could not possess these concepts.

Given the lack of legitimate moral disagreement within our society about the principle, "Human slavery is morally wrong," it seems that such a principle would be considered by the conceptual relativist to be a rule of moral grammar rather than a substantive moral principle. Following conceptual relativism, subscribing to a rule of moral grammar such as, "Human slavery is morally wrong," is necessary to understanding what our conception of morality (and according to conceptual relativism, the only conception of morality) is. When we are faced with apparent evidence that the ancient Greeks did not subscribe to this rule of moral grammar, it seems that the conceptual relativist is required to admit that the ancient Greeks did not understand what morality is.

While Arrington does concede that "[the ancient Greek's] normative concerns, though not identical to ours, nevertheless do overlap with our own" (p 252), he insists that these concerns "must be identified as a distinct concept" (p 252). But this leaves him with no way to say, using our language, what it was the ancient Greeks thought they were doing when they apparently argued about right and wrong. Arrington could not say that they were doing the same sort of thing that we do when we have a moral argument, because the rules of moral grammar that constitute the concept of morality and that provide the necessary background for engaging in moral argument were not in place for the ancient Greeks.

However, to even raise the question of the normative concerns of the ancient Greeks seems to be some kind of concession that they had some word

or phrase that can be interpreted (even if only very roughly) by using the English word 'morality'. It might seem on the face of it to be a reasonable claim to say that if a word or phrase plays the same sort of semantic role in the language of the ancient Greeks as it plays in our language, then the word or phrase means the same sort of thing (although it would be a matter of controversy within philosophy of language as to just how a vague idea like this should be cashed out). Although emphases may differ across the two cultures, Arrington seems to be drawing a very long bow indeed to insist that the concept employed by the ancient Greeks did not and could not contain anything resembling the concept of morality employed by us today. I suggest it would be hard to find in the works of the ancient Greek philosophers any hint that actions such as murder, lying or breaking promises were considered by them to be anything other than morally wrong and to be condemned by all right-thinking people.

Arrington does mention "value-neutral descriptive conditions (e.g., formal or functional ones)" (p 252) that would assist us in trying to articulate what the ancient Greeks were doing, but it is not clear how helpful this is going to be to him. He does not say what either the formal or functional descriptive conditions might be, and it seems his position does not allow any real room for such descriptive conditions. These conditions could not identify the ancient Greek's concept in terms of the similar role it plays in their language and lives to the role our concept plays in our language and lives, as this renders implausible the claim that the concepts are not the same. Furthermore, description of the ancient Greek concept in terms of non-moral goals such as conflict resolution and suchlike would run the risk of leading Arrington towards objectivism or subjectivism, which is exactly the problem he saw in the theories of Harman and Wong. Arrington's claim that the ancient Greeks had a different concept in place instead of our concept of morality depends for its plausibility on being able to identify some aspects of ancient Greek life as having a recognisably moral dimension. However, this very identification

103

removes credibility from the claim that the ancient Greeks had no concept of morality.

Furthermore, according to conceptual relativism, subscribing to a rule of moral grammar such as, "Human slavery is morally wrong," is not only necessary to understanding what morality is, but is also necessary to understanding what human slavery is. So, not only did the ancient Greeks not understand what morality is, they did not even understand what human slavery is. Despite the fact that they were far more closely acquainted with it than we are, they did not understand what human slavery is, but we do. This seems to be an outrageously parochial stance for the conceptual relativist to take. It does not appear to be something that is entirely out of the question to imagine a representative from our society debating the morality of human slavery with a representative of ancient Greek society.

We might wonder why Arrington does not at any stage consider the possibility that morality may be conceived of as a Wittgensteinian family-resemblance concept. This lack is curious for someone who draws on Wittgenstein as a major source of inspiration. One might have expected Arrington to say something to forestall the objection that while we *may* identify the ancient Greek's normative concerns as a distinct concept from our own morality if it suits our purposes, it is not at all clear that we *must* do so.

A metaphor popular with Wittgenstein was that of the game; and Arrington perhaps ought to have considered the possibility that the concept of morality may be thought to be closer to the concept of game than to the concept of colour. Wittgenstein compares different games with each other, saying: "For if you look at them you will not see something that is common to *all*, but similarities, relationships, and a whole series of them at that." (§ 66, p 32). The differences between activities such as chess and tennis do not prevent us calling both of them 'games'. A similar thing might be said by a Wittgensteinian to happen when we compare the ancient Greek's normative concerns with our own morality - we find many differences, but also many similarities, and there is nothing to prevent us calling them both 'morality'.

Merely pointing out a difference between our normative concerns and the ancient Greek's normative concerns does not seem to be sufficient to justify Arrington's claim that we must call only ours 'morality'. Wittgenstein himself even says, in the only mention he makes of ethical matters in *Philosophical Investigations*, that the word 'good', "must have a family of meanings". (§ 77, p 36)

Of course one does not have to be a Wittgensteinian or to buy into the notion of family resemblance concepts to think that Arrington has not made a convincing enough case for saying that the ancient Greeks did not possess the concept of morality or the concept of human slavery. One might take a view that there are similarities underlying both our normative concerns and those of the ancient Greeks that make it possible for us to identify with the ancient Greeks to an extent that Arrington does not seem to be willing to admit. I think it is not so far-fetched to accept that the ancient Greeks possessed the concept of morality even if they did keep human slaves, and the suggestion that they did not possess or did not understand the concept of human slavery simply beggars belief. In section 4.2.3 below I present my own recommendations for thinking about the issue of human slavery as a part of my attempt to achieve a compromise between the positions of the moral relativist and the moral universalist.

2.10 General afterword on moral relativism

This concludes my excursion into the three most prominent theories of moral relativism currently on offer, those of Harman, Wong and Arrington. Each of these theorists presents a way to accommodate Conviction R, and to answer Q3 with various versions of the claim that different groups are committed to different moral principles because moral principles are context-dependent. For Harman this is because tacit bargaining in the context of differing social circumstances has led different groups to be motivated to make different moral inner judgements. For Wong it is because, also in the context of differing social circumstances, different groups are led to operate with

differing extensions of the phrase 'adequate moral system'. For Arrington it is because, in the context of a commitment by some groups to principles other than the rules of moral grammar, these groups operate with concepts other than morality. I have presented criticisms of each of these varieties of moral relativism, both in the detail of how the theories are worked out and in the general assumptions behind each theory. I do not think that any of these three versions of moral relativism convincingly makes the case for unqualifiedly answering 'Yes' to Q3, and completely ruling out Conviction U. I now turn to examine moral universalism.

Part Three
Moral Universalism

3.1 Introduction

It is unfortunate that the confusion noted above in section 1.5, regarding just which metaethical position is opposed to the moral relativist, extends even to those who do take themselves to be opposed to moral relativism. For instance, in the recent book on this topic in the *Great Debates in Philosophy* series, Harman contributes an essay on moral relativism (from which I have quoted extensively in describing his position), and Judith Jarvis Thomson contributes an essay that according to the format of the book is supposed to be a counter to Harman's position. However, in her contribution Thomson argues at length for her "Thesis of Moral Objectivity: It is possible to find out about some moral sentences that they are true," (p 68) when this is a thesis Harman does not deny. In his response to Thomson, Harman explicitly notes that he is quite able to accept moral objectivity (p 158). It seems that Thomson erroneously believes that to establish her cognitivist thesis is to damage the credibility of Harman's moral relativism. However, she does not actually argue against moral relativism, but only against something she calls "Moral Scepticism", which is simply the negation of the above thesis of moral objectivity.

Harman points out that there is no reason why a moral relativist such as himself would need to be committed to the moral scepticism that Thomson criticises. On the contrary, the moral relativist would want to oppose such a view. The moral relativist certainly thinks that it is possible to find out about moral sentences that they are true (relative to some context of course), by consulting our (as the relativist sees them) context-dependent moral principles. Context-dependence is not at all the same as unknowability, nor

107

does the former imply the latter. To the moral relativist it would be an absurd idea that we somehow might not know what our own local moral standards are, or that we might not be able to find out whether moral sentences were true for us in our own local context. If anything, the relativist might say, it would be easier to find this out than to find out whether moral sentences are true in a universal sense.

The battle that Thomson wants to fight over whether moral objectivity or moral scepticism is the case only bears very indirectly, if at all, on the issue between the moral universalist and the moral relativist. I will not go over Thomson's arguments against moral scepticism here as moral relativism is a different position from moral scepticism. It is unfortunate that, despite the essay-and-response format of the book in which Harman has most lately presented his version of moral relativism, he does not receive the critical attention of a genuine moral universalist.

3.2 Universalism and rationalism

3.2.1 Two separate questions.

One way that a moral universalist might defend Conviction U is to try to show that it is the conclusion of a valid argument with all true premises. The goal would be to demonstrate that there is some rationally compelling argument supporting Conviction U that would allow the universalist to label as irrational anyone who did not accept it. This rationalist route to universalism will be explored throughout this section.

When examining rationalist arguments for universalism, it is important to keep separate two distinct metaethical questions. There is a significant difference between the question of whether rational argument will lead us to one and only one set of universal moral principles and the question of whether being rational necessarily motivates us to follow moral principles. It is only the former question that the moral universalist is concerned to pursue. This is because it is unnecessary to give any argument forcing a moral relativist to be motivated to act morally. The moral relativist is *already*

motivated to follow moral principles, although of course these are context-dependent moral principles rather than universal moral principles. See, for instance, Harman (1982, p 190).

I have argued in section 1.6 that moral relativism is a distinct metaethical position from moral scepticism or nihilism. Unlike proponents of these latter two positions, the moral relativist is firmly committed to doing the right thing. This is important to note because it seems that a great deal of the rationalist effort in metaethics is devoted to showing that the rational agent must be motivated to be moral, and further devoted to deciding just what rational argument or arguments can succeed in showing this. But the moral relativist and the moral universalist do not differ in their motivation to be moral. This is something that they both share. The moral universalist who looks to rational argument for assistance faces only the task of showing that rational argument leads to one and only one set of universal moral principles, and is not required to show that rational agents are necessarily motivated to be moral.

However, it might be thought that the argument or arguments that rationally compel agents to be motivated to act morally also at the same time show that there is one and only one universal set of moral principles that agents are rationally compelled to be motivated to follow. A rationalist might attempt to simultaneously answer the two metaethical questions distinguished in the previous paragraph. But the out-and-out rationalist (this label I give to the rationalist who tries to answer both questions at once) will face extra difficulties that the moral universalist will not. The following subsection will illustrate one of these difficulties.

3.2.2 A difficulty for rationalism.

Consider the sort of things a rationalist might say to someone who claims to 'reject' the following moral principle:

(I) Murder is morally wrong.

109

Pointing out the reaction of fear or disgust we are likely to have towards someone who blatantly rejects (I) would not necessarily show that it is the rejecter who is irrational. Reactions of fear or disgust are not themselves always entirely rational. More promisingly, (I) may be said by the rationalist to follow necessarily from other principles that the person who rejects (I) holds, so that it is inconsistent of them to reject (I) while retaining these other principles. However, this still leaves it open as to whether it is (I) or one or more of the other principles that the person must relinquish if they are to remain consistent. A stronger rationalist line would be to try to prove that (I) follows with certainty by deduction from self-evident premises, or something in a similar vein, so that anyone who can follow an argument is compelled to accept (I). The burden of proof is here obviously being accepted by the rationalist, and the plausibility of their claim will depend entirely on just what arguments of the kind mentioned above they are able to produce.

Alternatively, the rationalist may try a more direct stipulative approach, and state that 'murder' means 'wrongful killing'. Surely it may not be intelligibly denied that wrongful killing is morally wrong. If (I) is true by definition, then the person who rejects (I) is simply ignorant of the meaning of the word 'murder'. But the rationalist is likely to want to say something more than this. To claim that 'murder' and 'wrongful killing' are synonymous just seems to be another way to state the conviction that (I) is the case.

An apparent difficulty for the rationalist is exposed with the idea that there is a ready counterexample to the claim that there is a failure of reasoning going on when someone rejects (I). It seems we have apparently good evidence that there are at least some individuals who reject the claim that murder is morally wrong and who do not appear to be irrational at all. The coldly calculating murderer who carefully plans and executes their murder, weighing various probabilities and taking various contingencies into account, always preferring A over C whenever they prefer A over B and B

over C, and so on (perhaps in the end never being caught for the murder they commit) seems on the face of it to be quite rational. An example of a calculating murderer could be Harman's invention of the successful criminal who is employed by the gang Murder, Incorporated (Harman 1982, p 191; 1984, p 39-40). I look at the implications of Harman's example of the successful criminal below in section 3.2.4.

A first response that the rationalist might make to this alleged counterexample is to say that the term 'rational' has more than one sense. The calculating murderer is 'rational' in one sense, but not in another, specifically moral, sense. However, this sort of reply looks like it is begging the question unless there is some independent reason to think there is a specifically moral sense of the term 'rational'. It is safer to proceed on the assumption that 'rational' is used in the same sense by the metaethicist as it is used by the logician or the economist.

More promisingly, the rationalist might say that in actuality the calculating murderer only appears to be rational. In every instance of calculated murder there is some hidden inconsistency or other error present in the reasoning of the murderer. This is not supposed to be an empirical claim that has to be established on a case by case basis, but rather a general theoretical claim. Nagel is one rationalist who takes this general line and gives the following argument for where the inconsistency is located. He talks about harm in general instead of murder, although if his argument works it will apply to the calculating murderer as well as to other cases:

> This is a matter of simple consistency. Once you admit that another person would have a reason not to harm you in similar circumstances, and once you admit that the reason he would have is very general and doesn't apply only to you, or to him, then to be consistent you have to admit that the same reason applies to you now. (Nagel, p 66)

The calculating murderer is invited to concede that in a comparable situation someone else would have a reason not to harm them (the murderer), and that furthermore this reason is a very general one so that the murderer in

111

their turn has a reason not to harm their victim. The problem for this rationalist line of argument lies in the idea that the calculating murderer would make these two concessions. They may be willing to acknowledge that if another person has a reason not to harm the murderer, then this reason is a very general one, but it may be less plausible to say that the calculating murderer must admit that another person has a reason not to harm the murderer. Indeed, once they become aware of the inconsistency, they may, in order to remain consistent, explicitly reject the idea that another person has a reason not to harm them. They may agree that they could have no moral complaint against someone else who made a plan to murder them and then carried it out before they themselves were able to commit the murder they were planning.

Nagel maintains that there is something like self-evidence attaching to the principle that others have a reason not to harm us, that this is a principle that "most people, unless they're crazy" would agree to. But taking this line may not be so convincing. Of course it is true that most non-crazy people are not calculating murderers, but this does not mean that all calculating murderers are necessarily crazy. Nagel would presumably not want to claim that 'crazy' in this context is simply supposed to *mean* 'irrational', as the irrationality of rejecting this principle is just what he is trying to show.

The rationalist may have other responses to make to the calculating murderer counterexample, but I will not pursue them. The point I would like to make with this example is that there is an ambiguity between two different ways that it is possible to 'reject' (I). Someone (perhaps the calculating murderer) may accept that (I) is required by morality but reject the principle because they are not motivated to act morally, or someone may be motivated to act morally but reject (I) because they do not think that morality requires it of them. An out-and-out rationalist like Nagel is arguing against the first sort of person as well as the second sort of person, whereas the universalist is troubled only by the second sort of person. This is why the calculating murderer counterexample would not pose a problem for the universalist.

While the out-and-out rationalist who is trying to show that all rational agents are necessarily motivated to act morally may have difficulty in dealing with counterexamples like the calculating murderer, the universalist is not trying to show this.

In the background here looms a significant metaethical controversy that threatens to become overly prominent. This is the opposition between internalism and externalism. Internalism is the view that for any rational agent there is an internal connection between the recognition that an action is morally required (or forbidden) and the motivation to perform (or to avoid) that action. Externalism is the denial of this, holding instead that something external, such as a desire to do what is morally right, is required before a rational agent must be motivated to act in accordance with morality. Obviously the out-and-out rationalist is an internalist, but the universalist may be an internalist or an externalist. Subscribing to the notion that there is one single universal set of moral principles does not in itself commit the universalist to a particular view about what a rational agent's relationship to these moral principles must be.

This means that the universalist is able to accept with equanimity the notion that the calculating murderer is simply someone who has abandoned morality. Such an individual fails to act according to the universal moral principle (I) because they refuse to be moral. It is not the case (as Nagel seems to suggest) that they are trying to act in the morally right way but they are somehow making a mistake about what acting in the morally right way entails. The calculating murderer is someone who has rejected (I) in the first sense described above, the sense of declining to be bound by the normative system of which (I) is a part. They are not motivated to follow (I), but this in itself does not entail that they do not believe that (I) is a universal moral principle.

For the universalist the useful result of recognising this distinction is that the counterexample of the calculating murderer (or indeed any other counterexample of an agent who is apparently not motivated to act morally)

does not threaten their position that there is one and only one set of universal moral principles, because the universalist is able to say that there is a sense in which even the calculating murderer can accept that murder is morally wrong. The universalist would perceive no problem in terms of the internal consistency of the position of the calculating murderer who believes that murder is morally wrong and that this is a universal principle, if the calculating murderer has no concern themselves about avoiding doing what is morally wrong. The amorality of the calculating murderer that is being envisaged here could be summed up with the phrase: 'Yes, I know that what I am doing is morally wrong, and indeed morally wrong for everyone, but I don't care about avoiding what is morally wrong.' Of course, an out-and-out rationalist such as Nagel is not going to be able to accept that such a thing can coherently be said, but the universalist need not be troubled by a statement such as this, because it is not an expression of moral relativism.

The upshot of these considerations is that the moral universalist does not have to concern themselves with trying to argue that all rational agents must necessarily be motivated to follow moral principles. Instead they may concentrate on arguing that there is one and only one set of universal moral principles that may rationally be accepted. To say that even a calculating murderer would not deny that murder is morally wrong buttresses Conviction U, the notion that there are some things that are simply unacceptable, morally speaking.

The point could also be put in terms of Humean moral psychology. The moral universalist is concerned only to argue that our beliefs will converge on a single set of moral principles, and is not required to argue that certain beliefs about moral principles will necessarily create the desire to follow them. The out-and-out rationalist will take it upon themselves to argue for both of these conclusions, trying to show that there is one universal set of moral principles and also that these principles are such that all rational agents are motivated to adhere to them, but the question of motivation is not an issue that exercises the moral universalist.

Things would be different if the calculating murderer were to say something like: 'Murder may be morally wrong for you, but it is not morally wrong for me.' This would be an expression of moral relativism rather than amoralism and the universalist must attempt to refute it. They may take a rationalist path here by trying to show that rational argument leads us to one and only one set of universal moral principles, and that (I) is one of these principles, so that if the calculating murderer were to make the above claim then they would be irrational.

3.2.3 A jumble of issues.

The prospects for rational arguments in support of universalism are unfortunately not easy to determine when universalism is only addressed as a part of an out-and-out rationalist program. The claims of the out-and-out rationalist in metaethics have been much contested, with numerous different attempts to provide the argument or arguments required, and also opposing efforts to refute these attempts. For instance, Arrington and Wong between them devote many pages to criticism of the endeavors of a number of different out-and-out rationalist philosophers. Arrington presents a detailed attack on the theories of Bernard Gert and Alan Gewirth (Arrington, p77-116), while Wong offers less comprehensive but more wide-ranging criticisms of the theories of Gewirth (again), Nagel, Kantians in general, McDowell, Foot, Moore, and Platts (Wong, p80-104). Unfortunately, in this barrage of criticism neither the distinction between moral universalism, moral absolutism and moral objectivism made in section 1.5, nor the distinction between the two separate metaethical questions made in section 4.2.1, are recognised. Many issues are conflated and that is why this arena is one in which it seems to me the debate has a great potential for becoming bogged down in what are, from the perspective of my project, side issues. Many of the metaethical questions that are raised by the out-and-out rationalist arguments considered by Arrington and Wong are not relevant to the issue that I am investigating. I have chosen not to enter into the details of these controversies as it seems to me that this would

merely contribute to the impasse I see in the dispute between the moral relativist and the moral universalist.

3.2.4 Smith's rationalism.

Despite the reservations noted in the previous section, it would be remiss not to include at least one example of how an out-and-out rationalist might argue against moral relativism. The theorist I will examine is Michael Smith, whose version of out-and-out rationalism is more recent and arguably more plausible than those versions considered and dismissed by Arrington and Wong.

Smith takes a distinctive approach to metaethics with his view that it is dominated (and metaethicists are divided) by what he calls "the moral problem". This problem is the difficulty of accepting two ideas that seem to pull in opposite directions. The first is the idea that moral judgements express our beliefs about matters of fact, and the second is the idea that moral judgements express our desires by being intimately connected with our motivation to act. It seems to be too much to expect moral judgements to fulfill both of these expectations, with the result that: "nothing could be everything a moral judgement purports to be" (Smith, p 11).

Smith's solution to the problem is to say that the concept of a reason for action is ambiguous. He makes a distinction between motivating reasons for action and normative reasons for action (p 95). Motivating reasons for action *explain* an agent taking a certain action and are psychological states that are reasons from the perspective of the agent themselves. In contrast to this, normative reasons for action *justify* an agent taking a certain action and are reasons from the perspective of a fully informed rational outside observer. For instance, an agent may have a motivating reason to turn left at the next intersection because he believes that turning left will get him to his destination most quickly, but the agent would only have a normative reason to turn left if a fully informed observer would agree. If, from the privileged perspective of a fully informed observer, turning right would actually get the agent to his destination most quickly, then the agent would have a motivating reason to turn

left but a normative reason to turn right. We could explain his turning left but we could not justify it.

Smith then draws out the implications of this distinction, noting that teleological/causal explanations of actions that cite motivating reasons will rely on claims about what an agent desires (p 131), while rational justifications of actions that cite normative reasons will rely on some other kind of attitude. Smith characterises this attitude as valuing, or what it is desirable to do (p 132). I will not give the details of Smith's discussion of the difference between what we desire to do and what is desirable for us to do (what we value), but Smith eventually says: "what it is desirable for us to do is what we would desire that we do if we were fully rational" (p 151).

This formula provides Smith with a mechanism to solve the problem he has set himself. He analyses the concept of moral rightness as, "the concept of what we would desire ourselves to do if we were fully rational" (p 185). Using this concept, our moral judgements about what is right can be seen as both expressing our beliefs about facts, that is, facts about what we have normative reason to do, and also as connected with our motivation to act. It is beyond the scope of this thesis to evaluate whether Smith's arguments in support of this mechanism are successful, but I would like to focus on the use he makes of this mechanism to criticise moral relativism.

He asks the question: "whether, if we were well-informed, cool, calm and collected, we would tend to *converge* in the desires we have" (1993, p 408, Italics in original). Smith thinks that we would. He thinks that moral discussion under idealised conditions of full information and perfect rationality would result in a massive convergence of desires, and furthermore that this convergence is not contingent but is due to the "privileged rational status" (p 408) of these desires. Smith does not give an argument for this claim that the desires of fully-informed rational agents would necessarily converge but he simply thinks it is more plausible to say that they would than to say that they would not. Such convergence of desires would be problematic for the apparent counter-example to out-and-out rationalism that arose earlier, the notion of the

calculating murderer. Smith is concerned to debunk this notion and he directs particular criticism against Harman's example of the successful criminal.

As we saw earlier, Harman created the example of a successful criminal employed by the gang Murder, Incorporated (Harman 1982, p 191; 1984, p 39-40) in order to help make moral relativism seem more plausible. This criminal does not mind in the least harming or killing anybody who is not a fellow member of their gang, and Harman takes such an example to support moral relativism. In the context of the criminal morality of the gang, moral principles are followed that are different from those moral principles followed by ordinary citizens. (See section 4.2.4 where I give my own account of what I think is really going on with this example) After describing the example, Harman says, "Someone can be completely rational without feeling concern and respect for outsiders" (1984, p 40). Smith's response, however, is to reject this utterly: "[Consider] the claim that fully rational creatures would want that, if they find themselves in the circumstances of the successful criminal, then they gain wealth no matter what the cost to others. And the successful criminal's opinion notwithstanding, it seems quite evident that we have no reason to believe that this is true. Fully rational creatures would want no such thing." (Smith, p 195) It is interesting and significant that Smith and Harman have exactly opposite views about this. Smith goes on, "[the successful criminal] sticks with this opinion despite the fact that virtually everyone disagrees with him." Smith's belief is that there is (or would be) a massive convergence of opinion about this but that some perversity on the part of the criminal keeps them from partaking in this convergence.

However, it seems to me inaccurate to represent this as a situation in which the opinion of a criminal is pitted against the opinion of "virtually everyone", when really the only opinions present are the opinions of Smith and of Harman. I must confess I find it difficult myself to arrive at an opinion about whether a fully rational creature who was in the shoes of a successful criminal would want to gain wealth no matter what the cost to others. Harman is absolutely convinced that they would, whereas Smith is absolutely convinced that they

would not. Smith might say that Harman is not fully appreciating the requirement that the successful criminal is supposed to be imagined as a fully rational creature, while Harman might counter that Smith is ignoring the constraint that the fully rational creature is supposed to be in the circumstances of a successful member of a criminal gang, and not in the circumstances of, say, an ordinary member of society. Harman could make the point that while there may indeed be a massive convergence of opinion that fully rational creatures in the circumstances of ordinary members of society would want that successful criminals not pursue wealth no matter what the cost to others, this is not what is under consideration.

Rather than massive convergence, it seems that Smith's approach to this issue potentially opens the door to massive disagreement instead. It faces the difficulty that nobody has ever been in the position of a fully informed, perfectly rational agent. Inevitably any metaethicist who takes this approach must speculate on what such an agent would desire, based on abstractions from those situations in which rational argument or becoming more fully informed resulted in the modification of desires. Speculation creates the potential for disagreement. The speculative nature of this enterprise is compounded in this case by the further abstraction required in placing the fully rational agent in circumstances which are unlikely to much resemble the circumstances of many metaethicist's lives (being a member of a successful criminal gang). I think it is fair to say we really have very little understanding of what the desires of such an agent might be. We could certainly guess, but it would be no more than a guess. As we see, Harman's guess is very different from Smith's guess. It seems likely that a moral relativist impressed by Conviction R would agree with Harman about what fully rational creatures in the shoes of a successful criminal would desire, while a moral universalist impressed with Conviction U would agree with Smith, so the situation seems to be once more at a stalemate. It appears that Smith's version of out-and-out rationalism has not taken us any further than the clash of convictions with which we began.

3.3 Coherentism

3.3.1 Distinct types of moral principle.

The coherentist project of defusing this clash of convictions begins with the claim that there are distinct types of moral principles. The dispute between moral relativists and moral universalists rests upon a key underlying assumption made by both parties. This is the assumption that all moral principles have the same sort of use within the practice of morality. Different moral principles look alike on the surface, and so they are taken by both the universalist and the relativist to be alike. I wish to argue that not all moral principles are of the same type. I believe that there are at least two distinct types of moral principle that are distinguished by being used in two importantly different ways within the practice of morality and that a conflation of these disparate uses leads to the conflict between relativists and universalists.

In Part One I argued that the main point of contention between the moral relativist and the moral universalist could be articulated as the following question:

Q3: Are moral principles context-dependent?

A context-dependent moral principle is one that ineluctably requires reference, either explicitly or implicitly, to a context. A context is, roughly speaking, the situation in which a principle is being used to arrive at a moral judgement about what to do. The moral relativist answers 'Yes' to Q3, although, as we have seen, different moral relativists have differing views on exactly how to characterise the context. The moral universalist answers 'No' to Q3, taking the position that all moral principles apply universally, so that a reference to context is not required. The answer to Q3 that I propose is: 'There are two types of moral principles, those that are context-dependent and those that are not.' Explanation and elucidation of this answer will comprise the bulk of the remainder of my thesis.

Despite their opposing viewpoints, the moral universalist and the moral relativist both have a similar goal and this goal is to say something general about moral principles considered as an overall group. Moral principles are seen through the lens of relativism or universalism as belonging to a homogenous class in which every member of that class has a similar status with regard to its context-dependence, that is, they are all either context-dependant or they are not. In pursuit of a general theory of the status of moral principles there seems to be little recognition of what could be called the different facets of the use of moral principles. The moral universalist tends to focus on just one of these facets and the moral relativist on another and then each of them generalises from their favourite kind of case to moral principles as a whole. Both the moral universalist and the moral relativist are able to give a not implausible account of things with respect to their favoured type of moral principle, and this further entrenches them into their own respective positions. I think it is often overlooked that the type of moral principles that the universalist favours as examples have a different use from the type of moral principles that the relativist favours as examples. It seems to me that carefully distinguishing between some of the different uses for moral principles within the practice of morality will go a long way towards achieving a coherent fitting together of Conviction U and Conviction R in a way that shows how the core convictions of universalism and relativism may be regarded as compatible. This is what I will attempt in what follows.

In order to investigate these different facets of the practice of morality, I will for the sake of simplicity restrict myself to constructions made with the predicate, '_ is morally wrong.' When it has as its subject a general class of actions this predicate can be used to make a large number of general moral principles. This predicate has the convenient feature that it is able to express either that a type of action is morally forbidden, or that a type of action is morally required, the latter by using the locution: 'Failing to do action of type such-and-such is morally wrong'.

The universalist will tend to point to moral principles such as 'Murder is morally wrong' and 'Lying is morally wrong' to support their view that moral principles are not context-dependant. The relativist will tend to point to principles such as 'Human slavery is morally wrong' and 'Abortion is morally wrong' to support their view that moral principles are context-dependant. Each of these principles has the same form and seems to make use of the same predicate, but the form of a particular moral principle is not necessarily going to be a reliable guide to its use. I will distinguish some significant differences in the way that moral principles are used and argue that an observation of these distinctions is a way to achieve the coherentist goal I have set myself. This will also go some way towards redirecting the discussion between relativists and universalists into more fruitful areas.

3.3.2 Basic moral principles.

With this overall coherentist goal in mind, I continue the examination of universalism by taking another look at the universalist's complaint against the relativist, a complaint that is encapsulated within the notion of Conviction U. The universalist maintains that it is a given of normative moral discourse that there are some elements that are, so to speak, not negotiable. The universalist's contention is that there would be something more than peculiar about the stance of someone who was trying to participate in normative moral discourse without acknowledging that there are some topics, such as whether or not someone has committed murder, that are always of relevance in a moral discussion. The universalist would question whether such an odd participant had any real understanding of how to take part in a moral discussion at all.

The view of the universalist is that no matter what the context or circumstances, an act such as murder is always of moral import and our moral discussions and our moral judgements must reflect this. It is important to note that while the universalist says that murder is always of relevance, this does not force the universalist's hand about the particular moral

judgements they may make about particular instances of murder (so long as the universalist is not also an absolutist). There are some situations in which mitigating factors may (although of course they also may not) excuse murder.

For example, consider the decision to deliberately target civilians during wartime as part of a policy of weakening the infrastructure of an enemy country and thus damaging their ability to wage war, with the overall aim of ending the war more quickly than might otherwise have been the case. Before a decision like this is made, moral discussion is called for, and various arguments might be considered. It might be argued that while the killing of civilians in this manner is murder, in some circumstances acts of murder may be excused, even if they may not be justified (See section 1.7). Alternatively, it might be argued that mitigating factors influence what counts as murder, so that in some circumstances the killing of civilians does not count as murder. Deciding whether the circumstances in question in this particular case are the sort of circumstances these arguments mention would require much more discussion before any moral judgements about the case could be arrived at. I do not mean to undertake such discussion here, but merely to point out that on the universalist conception, murder (even murder that can be excused) is a matter that is always of relevance when moral judgements are being made.

The moral universalist emphasises that this is very different from the situation that seems to be possible under moral relativism, whereby in a certain context it could be possible for it to be completely morally irrelevant that something is an act of murder. The moral relativist appears to the universalist to be willing to allow context to overshadow absolutely everything else. So in the above example an examination of any possible mitigating factors that might or might not excuse the murder of civilians may be unnecessary under moral relativism. The context may be such that the murder of the civilians, while it is acknowledged as being murder, does not even need to be excused because in this context murder does not matter morally.

The moral universalist (at least, the moral universalist who is not also an absolutist) is able to take context into account, as I have indicated above. However, they believe that there is a significant difference between saying: 'In this context, murder may be excused,' or: 'In this context, these acts are not acts of murder,' and saying: 'In this context, it is morally irrelevant that these acts are acts of murder.' In the first two cases the moral principle, 'Murder is morally wrong,' is accepted as being prior to the context in importance, so that the moral principle must be acknowledged even if it is outweighed, while in the latter case the context has usurped the priority of the moral principle. To the universalist this latter situation seems to be an unacceptable implication of moral relativism, and this is seen by the universalist to count strongly against the possibility of moral principles being context-dependent, and in favour of the view that moral principles are universally applicable. To the universalist, context merely determines how moral principles are to be applied, while for the relativist context determines the moral principles themselves.

There is some helpful terminology that captures the notion the moral universalist is interested in when making this point against the moral relativist. The universalist who is not also an absolutist sees moral principles as pro tanto principles. A pro tanto moral principle expresses a consideration that is always of concern to us and that must always be given some genuine weight when engaging in normative moral discourse. Although it is possible for mitigating factors to outweigh the concerns raised by the pro tanto moral principle, this does not mean that such concerns are negated, they have merely been outweighed in this instance. For example, that an act is an act of murder is always something that is of moral concern to us, even in those circumstances where it is possible to excuse the act.

A pro tanto moral principle may be contrasted with a prima facie moral principle. The latter is a moral principle that appears to have genuine weight, but that may on further examination prove to actually have no weight at all. It is possible that the concerns that are raised by the prima facie moral principle

124

are not relevant and may be ignored. (This terminology was first used in this way by Shelly Kagan in 1989. For his similar account of "pro tanto" and the distinction between this term and "prima facie", see Kagan p 17) We could say that the moral universalist fears that moral relativism turns pro tanto moral principles such as, 'Murder is morally wrong,' into prima facie moral principles.

The moral absolutist, of course, does not believe in pro tanto moral principles (although the absolutist may believe in prima facie moral principles), so the universalist who also happens to be an absolutist will not use this term. However, as I have already noted in section 1.5, a commitment to absolutism is not something that necessarily distinguishes the universalist from the relativist, so there is nothing about universalism in itself that prevents moral principles from being pro tanto moral principles.

My suggestion is that in making the point of the last but one paragraph against the moral relativist, the moral universalist is partially, but only partially, correct. Some, but not all, moral principles are pro tanto principles that apply universally. These universal moral principles I will call *basic moral principles*. It is by means of the notion of basic moral principles that Conviction U is accommodated. In the following three sections I will explain what basic moral principles are before turning to the question of what degree of justification it is appropriate to ask for them.

Basic moral principles are pro tanto principles. They indicate the fundamental concerns of any moral discussion and hence the sort of things that must be taken into consideration whenever moral discussion occurs. However, the basic moral principles do not dictate what the results of any given moral discussion must be nor do they dictate what moral judgements we will arrive at. An example of this has already been given above. 'Murder is morally wrong,' is a basic moral principle, and this means that murder is always a relevant concern when undertaking moral discussion, but this does not mean that murder may not sometimes be morally excusable.

125

As a part of deciding what we should do in various situations, we undertake many different sorts of normative discussion, both moral discussion and normative discussion of other kinds. It is not the case that just any kind of normative discussion counts as a moral discussion. Prudential discussion is an example of normative discussion that does not count as moral discussion. There must be some way to separate moral discussion from other sorts of normative discussion such as prudential discussion. The way we achieve this separation is by recognising that there are different fundamental concerns that shape the nature of each of the practices that these discussions are part of. Understanding what these fundamental concerns are is just the ability to appropriately engage in the relevant kind of discussion. The concerns in question must be universal to the practice because they are what makes the practice the practice that it is rather than some other practice.

The basic moral principles establish the sort of features of a situation that are of moral concern to us. These are the features that we must take into account when we make a moral judgement as opposed to some other kind of judgment. They are presupposed whenever a discussion that counts as a moral discussion is taking place. If a discussion is taking place in which the basic moral principles are not acknowledged by the participants, then whatever kind of discussion it may be, it is not a moral discussion. An acknowledgement of the basic moral principles is a necessary (although not sufficient) condition for participation in moral discussion. That is why the basic moral principles must be universally present in all moral discussion.

Despite their importance, basic moral principles do not often have to be bought up as a part of a moral discussion. This is because they work as presuppositions that allow a moral discussion to occur. The participants in a moral discussion must already be well aware of the basic moral principles and furthermore be committed to them as one of the conditions of their effective participation in the discussion. It is this that explains the universalist reaction described above. All effective participants in normative moral

discussion will understand that the basic moral principles are not something that can be questioned from within the precincts of such discussion.

Although they will not figure as a part of the content of a moral discussion, basic moral principles may sometimes be voiced as a reminder to participants that the current discussion is a moral discussion and not some other kind of normative discussion. It could be that a participant is perhaps allowing prudential considerations to overshadow moral ones and is thus straying from moral discussion into prudential discussion. For example, someone may need to be reminded that speculation on the financial benefits that may accrue to oneself from the killing of a wealthy elderly relative would be out of place in a moral discussion about whether it was right to do so, although it would be appropriate as a part of a discussion about whether it was prudent to do so.

3.3.3 Principle H.

According to my account, the sorts of things that are of fundamental moral concern are already recognised by anyone who understands how to participate in moral discussion, so the basic moral principles will not be anything novel or unexpected but should rather come as no surprise. 'Murder is morally wrong,' is merely the most obvious one. Other basic moral principles may be articulated differently by different commentators but the thrust behind different articulations will be consistently similar concerns. Despite other flaws in his approach, it seems to me that Arrington has expressed these moral concerns as well as anyone has in the examples he gives of rules of moral grammar. Using the predicate I recommended above, '_ is morally wrong,' Arrington's list is:

'Lying is morally wrong,'

'Breaking a promise is morally wrong,'

'Harming other people is morally wrong,'

'Treating others disrespectfully is morally wrong,'

'Taking the life of an innocent person is morally wrong.' (p 302)

Comparable lists are given by other universalist commentators, such as Pojman, who also adds to the above list:

'Rape is morally wrong,'
'Depriving another person of their freedom is morally wrong,'
'Failing to treat equals equally and unequals unequally is morally wrong,'
'Failing to help other people is morally wrong,'
'Failing to obey just laws is morally wrong.' (p 514)

Arrington's and Pojman's lists could possibly have other principles added to them that also deserve to be considered as basic moral principles, but rather than do that I would prefer to try to get at an important notion that seems to lie behind the basic moral principles. This is the notion of harm to others. The question of whether harm would be or has been done to others is always something that is of relevance in a moral discussion. Some immediate clarification of what it is to do harm to others will be offered here, but more important claims about this notion will be made in Part Four.

Harm may be straightforwardly, if somewhat vaguely, analysed as making someone worse off than they would otherwise have been. Such an analysis does not make it clear whether not helping should also be considered as harming, when not helping leaves someone worse off than they would otherwise have been. I will not try to say here whether not helping is to be considered as harming or the moral equivalent of harming.

It might be argued that if harming is making someone worse off than they would otherwise have been, then the concept of harm is not sufficient to capture the full extent to which it is possible to do to others what is morally wrong. For instance, in a recent article Rahul Kumar discusses a particular case that he thinks illustrates this. He believes it is possible that parents may have done something morally wrong to their child if the child is born in

128

circumstances that severely restrict it's quality of life, when these restrictions could have been avoided by the parents using contraception and preventing the birth of the child. It seems that the parent's failure to do this has not made the child worse off than it would otherwise have been because the child would not otherwise have existed. So it is arguable that the child has not been harmed by the parents but the parents have done something morally wrong to the child. Kumar's solution is to propose the notion of wronging as replacing the notion of harming. To wrong someone is to have "violated certain legitimate expectations" (p 107) of the wronged party. I will not try to evaluate Kumar's contention here, but I raise it in order to recognise that to some harm may not be an adequate characterisation of what underlies basic moral principles. However, harm can still be a universal value even if Kumar's counterexample means that it has narrower scope than initially thought. So, the more general universalist position is not threatened by arguments such as Kumar's that would replace one universal notion (harming) with another (wronging). A different kind of criticism of the adequacy of the notion of harm, a criticism advanced by the pluralist, will be examined in the next section.

To clarify what is meant by 'others' it is necessary to make the distinction between *moral agents* and *moral patients*. (The latter term was first used by N. Fotion in 1968. For a fuller account of the history of the use of the term 'moral patient' see Hajdin, p 180) This distinction, in particular the concept of the moral patient, will become quite important in Part Four when the moral relativist's Conviction R is looked at more closely.

Moral agents are those to whom moral judgements are addressed – those who can meaningfully be said to be in a position to do or to avoid doing what is morally wrong and who, because of this, may be said to have moral obligations or duties and to have followed or contravened their moral obligations. Moral patients are those who are supposed to be protected by the judgements of morality – those who can meaningfully be said to have been harmed (or wronged) if a moral obligation is contravened. Moral patients are

those entities that are morally considerable. The moral agent could be considered the subject of a moral judgement, while the moral patient could be considered the object of a moral judgement. This distinction does not necessarily preclude cases in which the moral agent and moral patient happen to be the same individual, so it is conceivable that one may have moral obligations or duties to oneself. For instance, the moral judgement may arguably be made that one has a moral obligation to oneself not to take one's own life. However, even if we accept such cases, the two roles remain logically distinct.

The distinction between moral agents and moral patients allows me to put forward a technical principle that I contend is a very general basic moral principle. I will call it Principle H:

Principle H) To harm a moral patient is pro tanto morally wrong.

This principle identifies a universal moral concern. A discussion in which there is no acknowledgement that to harm a moral patient is pro tanto morally wrong could not be a moral discussion, no matter what the context in which this discussion occurred.

It is notable that we can derive many of the other less general basic moral principles from Principle H, including 'Murder is morally wrong', 'Lying is morally wrong' and others on Arrington's and Pojman's lists. Some basic moral principles could be presented as the conclusion of an argument with 'To harm a moral patient is pro tanto morally wrong' as one of its premises. For example:

M1) To harm a moral patient is pro tanto morally wrong.

M2) An act of murder always harms a moral patient.

Conclusion) Murder is morally wrong.

It might seem to be tempting to assert that Principle H is *the most* basic moral principle, and that all the other principles on Arrington's and Pojman's lists could be extrapolated or constructed from this one, using arguments such as the one given above. This issue must be dealt with before continuing so in the following two sections I will investigate this claim.

3.3.4 Monism and pluralism.

The assertion that Principle H is the most basic moral principle touches upon the issue of monism versus pluralism. The monist believes that there is one single most general moral principle, perhaps something like Principle H, and that all of our other moral principles may ultimately be reduced to this one. The pluralist believes that there are a number of incommensurable general moral principles underlying morality, and these are not reducible to one single most general moral principle. The assertion considered above, that Principle H is the most basic moral principle, would be a monist assertion.

The pluralist might argue against this that not all of the basic moral principles on the lists above may be derived from 'To harm a moral patient is morally wrong.' It may be objected that while the murder or rape of a moral patient will always involve harm to the moral patient, there are other basic moral principles, such as, 'Lying is morally wrong,' and, 'Breaking a promise is morally wrong,' that do not fit this ideal. The pluralist may argue that it is not always the case that lying to a moral patient or breaking a promise made to a moral patient harms them. It might happen that being told the truth would in some cases harm a moral patient more than being told a lie. The pluralist could say the same about breaking promises. Breaking a promise made to a moral patient might not always harm the moral patient, and keeping a promise to a moral patient might in some cases harm them more than breaking the promise would do.

According to this pluralist line we could not come up with a sound argument about lying that paralleled the one given above about murder, thus:

L1) To harm a moral patient is morally wrong.

L2) An act of lying always harms a moral patient.

Conclusion) Lying is morally wrong.

This argument is unsound, the pluralist would urge, because the second premise is false; an act of lying does not always harm a moral patient. The same problem would arise again in giving such an argument about breaking promises, and some of the other basic moral principles may create similar difficulties. However, even if the pluralist is correct, this does not affect the moral universalist's position. So long as either the single most basic moral principle (if monism is the case) or at least one of the plurality of basic moral principles (if pluralism is the case) are acknowledged as universal, then the question of monism versus pluralism is not of direct concern to the moral universalist.

The important point for the universalist is that Principle H at least is a universal moral principle. In every possible context in which a moral discussion occurs, it is always relevant to consider whether the action under discussion would harm any moral patient, and furthermore, harm to moral patients always weighs morally in the same way, that is, against the action being discussed. It may be that the pluralist is correct and that not all of the basic moral principles in Arrington's and Pojman's lists are able to be derived from Principle H. Perhaps another very general moral principle, something such as: 'Failing to respect a moral patient is morally wrong,' would be required in order to derive basic moral principles about things such as lying and breaking promises, but in so far as my account of moral universalism is concerned, to accept that at least one moral principle, Principle H, is universal is sufficient.

3.3.5 Basic moral principles and rules of grammar.

There is some resemblance between my basic moral principles and Arrington's rules of moral grammar, especially in the idea that the basic

moral principles must be presupposed by participants in the practice of morality. However, although it seems to me that Arrington's notion is flawed, as I have indicated in section 2.9, nevertheless I think that the grammatical rules of a language can provide a convenient analogy that elucidates the way that the basic moral principles work. Despite this, basic moral principles are not literally grammatical rules. Explaining why will help to clarify the notion of basic moral principles.

The general rules of grammar provide us with a linguistic background or set of common presuppositions by means of which we can express ourselves through language and communicate with one another. Basic moral principles play roughly the same role in relation to moral language use that general grammatical rules play in relation to general language use. Basic moral principles provide us with a set of common presuppositions by means of which we are able to express our views on moral matters through moral discussion and communicate these views to others. We are also able to seek moral advice from others on matters about which we may be undecided, and also at times we may change our moral judgements as a result of moral discussion. These things are only possible against the background provided by the basic moral principles.

Taking the analogy one step further, basic moral principles have an instructional aspect, as do grammatical rules. A basic moral principle such as 'Murder is morally wrong' may be employed when teaching children about what it is to be moral, as a pedagogical aid to help them form an idea about what it is to do something morally wrong. Exposure to basic moral principles may serve as an introduction into the practice of moral discussion by marking out the limits within which moral discussion takes place.

Despite these similarities, the analogy between basic moral principles and grammatical rules has its limitations. There is a strong sense in which speaking in a grammatically correct manner has some optionality about it, whereas acting morally does not. Verbal language use very frequently violates the rules of grammar and yet people are able to communicate quite

satisfactorily despite this. Condemnation of ungrammatical language use is only a feature of certain formal situations or special circumstances in which something important is depending upon precise grammatical accuracy. But in ordinary speech someone may deliberately choose to ignore the rules of grammatically correct English and still have little or no difficulty in communicating effectively with others.

However, violation of basic moral principles, or the rejection of these principles altogether, is regarded as a much more serious matter. Moral condemnation is considered to be a matter of far more gravity than condemnation for ungrammatical speech. A deliberate rejection of the basic moral principles may even be thought by some (such as the out-and-out rationalist considered in section 3.2) to be an indication of irrationality, whereas deliberate ungrammaticality may be viewed as no more than laziness. Furthermore, when children are taught the basic moral principles a very important part of this instruction is that the moral dimension of life is one that they are expected to become a part of. Acting morally is not merely a convenient way to get along with others, but something far more important, whereas speaking grammatically does not have the same degree of magnitude attached to it.

It may be objected that this is all just a matter of degree. It might be said that someone who completely disregarded grammatical rules to the extent that they were constantly speaking nonsense would be quite likely to be regarded by many as irrational, and such complete lack of grammaticality may be considered to be more serious that the occasional lie or breaking of a promise. While this may be true, we need to compare like cases with like to get an accurate picture. The less serious moral transgressions such as lying should be compared to the less serious grammatical transgressions such as uttering an incomplete sentence, something that is an extremely common grammatical lapse. I suggest that the individual who lied as frequently as they uttered incomplete sentences would be a target of moral condemnation long before they were a target of grammatical condemnation. It is also difficult to

think of an act of ungrammaticality that would be comparable in seriousness to an act such as murder. When like cases are compared, I think it is evident that the breaking of moral rules is regarded as more serious than the breaking of grammatical rules.

Furthermore, linguistic rules of grammar are not pro tanto principles. Rules of grammar are not sometimes judged to have been outweighed by other rules of grammar. We may sometimes face moral dilemmas in which moral principles are weighed against one another, but we never face grammatical dilemmas in which rules of grammar must be weighed against one another. Of course, this point does not apply if one is an absolutist. To the absolutist, moral principles are far more like rules of grammar in this respect.

So while the grammar analogy may be a useful tool to help to clarify the role that basic moral principles play in moral language, it must be borne in mind that it is disanalogous to the extent that it does not fully capture the significance and importance to us of following the basic moral principles and acting morally, nor does it reflect the way in which moral principles (if one is not an absolutist) must sometimes be weighed against one another.

3.4 Anti-universalist objections

3.4.1 Cultural construction.

The role played within the practice of morality by the basic moral principles can be seen as being what is behind the universalist's Conviction U. This claim is the first part of my coherentist project of showing how Conviction U and Conviction R can fit together within the same theory. However, before proceeding with this project I will look at two significant lines of objection to the very idea of universal basic moral principles.

The first line of objection is that it is not possible to identify any so-called central concerns of morality, such as harm to moral patients, as being somehow 'given' in a way that is free from cultural variation. The objection maintains that so-called central concerns are differently constructed by

different cultures or different groups within cultures. (The objection is cited in Nussbaum p 40-2) Furthermore, it would be possible to imagine a form of life that did not contain these central concerns at all, in any form.

My response to this line of objection is straightforward. I would question the extent to which central concerns of the practice of morality could be differently constructed across different cultures. Nussbaum is suspicious of how far such difference extends, saying: "Relativists tend ... to underestimate the amount of attunement, recognition, and overlap that actually obtains across cultures ... Despite the evident differences in the specific cultural shaping of [people's experience], we do recognise the experiences of people in other cultures as similar to our own." (p 46) An example of underestimation of similarity is the view expressed to Nussbaum by an anthropologist that the British eradication of smallpox in India may have been morally suspect because it resulted in the collapse of the cult of the deity thought to control the disease. For Nussbaum, this sort of extreme sensitivity to cultural variation overlooks the obvious point that preventing Indians suffering from smallpox is morally commendable, and that: "we are all subject to certain physical constraints and liabilities, whatever our historical and social location" (This example is given in Antony, p 23-4). Nussbaum also emphasises the way in which it is possible for us to relate to works of art from cultures very different to our own, suggesting that cultural variation only goes so far.

A variant of this criticism from cultural construction would be to say that what counts as a moral discussion is itself culturally constructed, so that in the context of a culture other than our own, what counts as a moral discussion could be something different from what it is in our culture. In our culture we have moral discussions in which the central concern (or one of them, if pluralism is the case) is the avoidance of harm to moral patients, but there could be a culture where moral discussions occur in which this is not a concern at all.

136

This suggestion is similar to the conceptual relativism advanced by Arrington, in which different cultures have different concepts of morality, and it faces the same sort of difficulty that he faces. Whatever evidence might be given for claiming that in another culture certain discussions occur in which there is no concern about harm to moral patients, is just as much evidence for the claim that these discussions are not moral discussions at all. Conversely, any evidence that could be advanced in favour of the view that what is occurring in some other culture is indeed moral discussion would be evidence that the central concern (or one of the central concerns) of such discussion is the avoidance of harm to moral patients.

Of course, it would be possible to take from another culture a kind of discussion that was utterly different from moral discussion and insist on giving it the label, 'moral discussion', but this would merely be empty terminological manipulation. To simply stipulate that in another culture something such as a discussion about the effect of weather on agriculture, for example, is actually a moral discussion seems to be taking a path towards moral nihilism. To say that practically anything could be a moral discussion is not far from saying there is no such thing as distinctively moral discussion.

Moreover, the moral universalist is not claiming that everywhere that humans interact then the basic moral principles will be found, or that following the basic moral principles is in any way essential to the human species, or anything of that nature. The moral universalist (according to my interpretation of universalism) claims only that the basic moral principles are universal to the practice of morality. It is quite compatible with this claim that there be a culture or form of life where the practice of morality is not followed or recognised. Naturally, such a form of life would be different from the one we are familiar with.

Finally, while it may be possible to imagine a form of life that did not embrace any of the concerns central to morality, this does not mean that it is possible to imagine the practice of morality occurring within such a form of life. The critic might ask: 'Isn't this admitting that the basic moral principles

are merely contingent on the fact that we are beings of a certain kind?' It is, but I suggest that there is nothing untoward about this. There is undeniably a sense in which the basic moral principles are contingent on some of the characteristics of the beings that use them, and that these characteristics might have been different. For instance, if the resurrection of the dead was commonplace and involved no inconveniences then perhaps, 'Murder is morally wrong,' would not be a basic moral principle. If it was possible to telepathically read the minds of others then perhaps, 'Lying is morally wrong,' would not be a basic moral principle. The basic moral principles we have are contingent in so far as it is logically possible that things may have been very different from the way they in fact are and such difference would effect, among other things, the practice of morality.

Once again, terminological sleight of hand is fruitless here. Someone might imagine beings that engaged in an alien practice and stipulate that the practice is to be called 'morality', but this would achieve little. No matter what word we decide to use as a label for it, the practice that we actually engage in that we happen to call by the name 'morality' has some central concerns that are universal to the practice, and while these concerns may not be relevant to some kinds of beings, this is not the same as saying that they could be irrelevant to us.

In accordance with my coherentist standpoint I think that the most reasonable stance to take towards this line of criticism is to advance a theory that allows room for the recognition of variation and differences of emphasis between cultures as well as recognition of some similarity between cultures (at least, between cultures in which the practice of morality occurs). The notion of basic moral principles is supposed to accommodate Conviction U, which is behind the recognition of similarity. I will introduce mechanisms in Part Four that I believe sufficiently allow for Conviction R, which is behind the recognition of cultural variation.

3.4.2 Moral scepticism.

138

A second line of objection is the challenge of scepticism. The sceptic may allege that before the notion of basic moral principles can be accepted some argument is needed to show that it is the principles that appear in lists such as Arrington's and Pojman's quoted earlier, and not some other principles, that are the basic moral principles. The question may be asked: 'Why are those principles rather than other principles the basic moral principles?' This question may be asked as an empirical question that demands a causal explanation, but I take it that a causal explanation is not what is being sought here, and not just because a causal explanation would be in danger of leading towards the naturalistic fallacy (see section 2.8.5). Instead the demand is for some kind of argument that demonstrates why the basic moral principles are the principles that I have identified, and not some other principles.

As we saw in section 3.3.3, there is a limited sense in which we may be able to justify some of the basic moral principles from within the borders of a moral discussion, by pointing out that while they are all general principles, basic moral principles differ in the degree of generality that they have. I have put forward Principle H as a very general basic moral principle. Because a principle such as 'Murder is morally wrong' is less general than Principle H, it might be said that the former principle may be justified by appeal to the latter. This means that we can give a reason within a moral discussion as to why murder is morally wrong. It is wrong because murder harms a moral patient and it is (pro tanto) morally wrong to harm a moral patient.

However, relying on this feature of the differing generalities of the basic moral principles only pushes the question of justification back a step. Justifying a basic moral principle by appeal to a more general basic moral principle will provide only a temporary reprieve. If it were asked: 'How does one justify Principle H? Why is it morally wrong to harm a moral patient?' then it seems we have to simply admit that something like Principle H is a presupposition that must be made in order to participate in moral discussion. We are able to engage in moral discussion to justify moral judgements, but we cannot engage in moral discussion to justify a very general basic moral

principle such as Principle H, as this principle is a presupposition that makes moral discussion itself possible. A concern about whether or not harm is done to moral patients is one of the characteristics that makes a moral discussion a moral discussion and not some other kind of discussion. We could not have a moral discussion with a person who does not accept the presuppositions of moral discussion, and the person who refuses to be concerned about whether harm is done to moral patients seems to be such a person. So, it is not possible to justify the basic moral principles at the level of moral discussion by doing anything more than pointing to Principle H, and this principle may not itself be justified at the level of moral discussion.

A determined sceptic would still have room to object against the notion of basic moral principles. Their objection would be that no argument has been given that forces anyone to accept that Principle H is the most general of the basic moral principles (or one of the most general, if pluralism is the case), and so the question has been begged. The sceptic may say that I have merely assumed that Principle H is a very general basic moral principle, and that nothing so far has been said that is sufficient to answer a question such as, 'Why is it (pro tanto) morally wrong to harm a moral patient?' As I have helped myself to an assumption about Principle H, a critic could do the same by making their own assumption that there is a context in which moral discussions occur without any concern at all for harming moral patients (or for any of the other things identified by the other most general basic moral principles, such as respect for moral patients, if pluralism is the case). For instance, it may be assumed by the critic that there is a context in which there are moral discussions where the sole concern is the degree of financial profit accruing from one's actions. However, I would maintain that an assumption such as this would again appear to be mere empty stipulation resembling that discussed in the response to the first objection in the section above, whereas my assumption about Principle H I take to reflect something that is there to be seen by all who care to examine what goes on in the practice of morality.

Nevertheless, it must be acknowledged that there is difficulty in answering a question such as, 'Why is it pro tanto morally wrong to harm a moral patient?' to the satisfaction of a determined sceptic, because any possible answer will be subject to a further demand for justification. For instance, it may be said that it is morally wrong to harm a moral patient because this takes away from the dignity of a moral patient and it is wrong to take away from the dignity of a moral patient. This answer may in turn inspire a demand from the sceptic for the reason why it is wrong to take away from the dignity of a moral patient, and so on. When presenting any philosophical theory it seems there is some point where an assumption or two must be made and it is here that the sceptic will direct their attack.

When the sceptic demands the premises that are supposed to support the assumption that it is pro tanto morally wrong to harm a moral patient, I concede that it may not be possible to produce any. In this case perhaps the practice of morality could be accused of hanging unsupported in the air with no solid foundation. Maybe such a lack of justification would even strike a blow against the reasonableness of undertaking moral discourse rather than some other kind of discourse when deciding on our actions. At this point the universalist might turn to an out-and-out rationalist such as Smith to come to the rescue wielding rational arguments that might defeat the sceptic, but the possibility of such rescue would depend on accepting the details of the argument presented by the out-and-out rationalist.

However, while the determined sceptic may be difficult to defeat, I have already explained in section 1.6 why I think the sceptical route is not a path that the moral relativist would want to go down. It matters to the moral relativist what is right and wrong and the casting into doubt of the practice of morality as a legitimate enterprise is not a part of the relativist's agenda. Scepticism is not the way to support Conviction R. The alternative that I offer is a coherentist package that allows Conviction R to coexist with Conviction U and it is to the details of this I will turn after summing up the thrust of this chapter.

141

3.5 Final remarks on universalism

It is important to reiterate that moral universalism is a distinct position from moral objectivism and from moral absolutism. The moral universalist claims only that the basic moral principles are universal to the practice of morality and to all moral discourse. To say that the basic moral principles are universal ones is not the same as saying that they are objective or absolute ones, and because of this the moral universalist is not required to produce the sort of arguments that a moral objectivist or a moral absolutist may need to provide. This stance is what might be called 'pure' moral universalism. It may be that many who are moral universalists also happen to be moral absolutists and/or moral objectivists. It would not be inconsistent to simultaneously subscribe to all three of these metaethical views. In so far as someone subscribes to one or both of the latter positions they may be obligated to give further metaethical arguments. What these arguments are would depend on the details of the moral absolutism or the moral objectivism in question. The important point in so far as my investigation is concerned is that Conviction U does not stand or fall on these issues. While the universalist *may* believe that basic moral rules are objective and/or absolute as well as believing that they are universal, the former beliefs are not a necessary feature of universalism.

My goal in Part Three has been to present the notion of basic moral principles as a way of thinking about morality that allows the taking on board of the moral universalist's Conviction U. Basic moral principles are not context-dependent, except for in the obvious but trivial sense in that they occur within the context of the practice of morality. It does not matter in what geographical location or period of history that a moral discussion occurs, or between what segments of society, if it is to count as a moral discussion then the presuppositions I have called basic moral principles must be in place. This is the line of thought that I suggest makes the best sense of Conviction U.

However, the moral universalist can take this too far. Not every moral principle is a basic moral principle. Even acknowledging the nature of basic moral principles as presuppositions of moral discussion does not allow us to privilege Conviction U by ruling out Conviction R because moral discussion still contains the prospect of much disagreement indeed, and not just about moral judgements. Legitimate participants in a moral discussion may refuse to accept some moral principles that other participants do accept and this is an indication that at least some moral principles are not basic moral principles. Here is where Conviction R must be given its due.

Part Four
Moral Relativism Revisited

4.1 Introduction

While Conviction U may be accommodated without being overlooked by using the mechanism of basic moral principles, this does not give it any privileged status over Conviction R. The relativist's insistence that many moral disagreements are real and cannot be just explained away as errors or mistakes may be accommodated by building on the notion that not all moral principles are basic moral principles.

The first thing to acknowledge is that a proportion of what might seem to be moral disagreements really are nothing more than arguments about the facts. What is often in dispute in such arguments is whether or not it can be said that certain sorts of actions do or do not cause harm to any moral patients. For instance, an argument about whether to accept the principle, 'Pornography is morally wrong,' may really turn on the empirical question of whether harm to moral patients results from pornography. If this is indeed an empirical question then the methods of empirical investigation rather than moral discussion are appropriate to attempting to answer it. Observations must be made, inductive inferences must be drawn from these observations, and so forth. In cases where an argument really is about an empirical question the ideal situation would be to end a moral discussion with an agreement on just what empirical evidence would decide the issue, or to provisionally end the moral discussion pending empirical evidence that would allow it to meaningfully proceed. Then efforts could be put into obtaining that empirical evidence, rather than mistaking speculation about empirical matters for moral discussion.

The moral relativist's case cannot rest on any seemingly moral disagreements that turn out to be empirical questions in disguise. If there is disagreement about an empirical question then this probably is due to a mistake about the facts on the part of one or more of the parties to the disagreement, or a lack of empirical evidence that leads to unwarranted speculation about what the facts actually are. The moral relativist is not interested in this sort of disagreement and Conviction R does not rest upon it, so in what follows I will be assuming that the disagreements I discuss are genuine moral disagreements and that any real empirical issues have been agreed upon by the participants.

To help clarify what I think is behind the fundamental conviction of the moral relativist, I would like to suppose for the sake of convenience of exposition that monism is the case, and that Principle H could legitimately be called the most basic moral principle, the one that all of the other basic moral principles may be derived from (see section 3.3.4). Making this supposition will allow me to more easily state which facets of moral discussion the moral relativist is appealing to. To recapitulate, Principle H says the following:

Principle H) To harm a moral patient is pro tanto morally wrong.

Briefly, my claim is we can make Conviction U and Conviction R fit together coherently by accepting that Principle H is universal to all moral discussion, so acknowledging Conviction U, and at the same time also accept that there is still much room for disagreement and diversity in the interpretation of key elements of Principle H, so acknowledging Conviction R. The elements of Principle H about which diversity may be entertained are:

1) There may be diversity with regard to who or what counts as a moral patient. Different moral principles may be advanced representing different views about which entities are legitimate moral patients. Some of these moral principles will be examined below in section 4.2.

145

2) Many of the basic moral principles derived from principle H are about different kinds of harm. There may be diversity about which of the many different possible kinds of harm are more morally significant, and there may also be different interpretations of the word 'harm' as it occurs in principle H. Although the participants to a moral discussion must each accept the basic moral principles, there may be disagreement about their order of priority. The order of priority of basic moral principles is something that will be discussed in greater detail below in section 4.3.

It is the possibility of these two separate dimensions of difference in interpretation or understanding of Principle H that I believe Conviction R is a response to.

4.2 The bounds of the moral community

4.2.1 Limiting moral principles.

Many of the apparently context-dependent moral principles that the moral relativist puts forward in support of Conviction R are not basic moral principles but instead concern the difficult problem of trying to establish the bounds of the moral community. This problem often takes center stage in moral discussion, and even when moral discussion is not explicitly about this matter it is frequently just beneath the surface. We establish the bounds of the moral community when we say which entities have moral standing or moral considerability. 'Moral patient' is the label given to those entities that have moral standing (See section 3.3.3)

A crucial claim I want to make is that often in moral discussion a moral principle is put forward that is actually an attempt to say who or what counts as moral patients. I will call these *limiting moral principles* because they are attempts to establish the limits of the moral community. Two examples of limiting moral principles are: 'Abortion is morally wrong' (to be precise this would have to say, 'Abortion of human fetuses is morally wrong', but for convenience I will work with the abbreviated version) or 'Human slavery is morally wrong.' These two principles have a different use in moral

146

discussion from the use that basic moral principles have. Limiting moral principles may not be readily distinguishable from basic moral principles by their form alone, as they both may be expressed in exactly the same form. For instance, 'Murder is morally wrong,' has the same form as 'Human slavery is morally wrong,' but this obscures an important difference in use. The latter principle is saying (in part) that it is human beings that have moral standing and this is something essentially different from what the first principle is saying.

In arguing their case for relativism the moral relativist may be drawn to giving limiting moral principles as examples. The many conflicts that occur over questions of what has moral standing and the controversial nature of some of the limiting moral principles, such as 'Abortion is morally wrong,' makes them ideal illustrations of the relativist's point about moral diversity. However if a moral relativist were to give a basic moral principle as an example of a principle about which there may be moral diversity then this would be a mistake.

In contrast to the basic moral principles, any of the limiting moral principles may be argued about and challenged within the precincts of a moral discussion. As a part of a moral discussion, disagreement and diversity of opinion about which entities do and which do not have moral standing is frequent and genuine. As I have noted above, such controversy seems to in fact constitute a fair proportion of moral discourse, and some of the most controversial areas of moral debate will be about what are the correct limiting moral principles. I suggest that this is the main reason why the moral relativist is drawn to citing limiting moral principles or disputes about limiting moral principles as examples when arguing for the context-dependence of moral principles. It is not implausible to claim that in different contexts different entities have moral standing.

Reviewing the list of examples of moral diversity given by Harman (1996 p 8-11), it can be seen that many of them concern limiting moral principles. This includes the examples of cannibalism, infanticide, caste systems, the

treatment of women or of animals, as well as the two just mentioned examples of abortion and human slavery. These are all moral issues at the root of which is a question about who or what has moral standing. I will look at the two latter examples in greater detail below.

Many different criteria have been advanced to determine which entities have moral standing. These include membership in *Homo sapiens*, potential personhood, rationality, sentience, being alive and being an ecosystem. It would be possible to produce moral arguments for or against any of these criteria, and such arguments could be criticised or defended while remaining squarely within moral discourse. The possibility of such moral argument has the important implication that while all parties to a moral discussion must be committed to the basic moral principles in order to be having a moral discussion at all, this commitment in itself does not entail any particular view about who or what has moral standing. Disagreement about limiting moral principles takes place against the background of agreement about the basic moral principles. Failure to acknowledge this can lead the moral relativist to overestimate the type and scale of the disagreement that is actually taking place whenever there is a divergence of views about limiting moral principles.

Many of the instances of moral diversity that a moral relativist may be drawn to, in which two groups appear to have irreconcilable moral differences, are disagreements about limiting moral principles. The parties to a moral disagreement over what has moral standing may frame the limiting moral principles that they are arguing about in the same form as the basic moral principles that provide the presuppositions allowing the disagreement to take place, and this similarity of form may lead to confusion.

A commitment to the basic moral principle, 'Murder is morally wrong,' is not something that is up for debate within moral discourse, but whether the killing of a particular kind of entity is the killing of an entity that has moral standing, and so counts as murder, certainly may be a subject for debate within moral discourse. This sort of observation is not new (see O'Neill, p

148

536 for historical precedents; see also Hajdin, p 90-1)), but so far as I am aware the impact of this distinction upon the metaethical issue of the dispute between the moral universalist and the moral relativist has not on the whole been noted. The moral universalist may be correct to say that anywhere and at any time moral discussion occurs the basic moral principles must be followed, but this in itself does not prevent the moral relativist from also being correct in observing that there are a very wide range of views about who or what is the proper recipient of moral consideration.

4.2.2 Example one: 'Abortion is morally wrong'.

One of the clearest examples of disagreement about a limiting moral principle is the debate concerning the moral status of abortion. This is one of Harman's examples of moral diversity and he invokes it as an instance of coexisting but radically different moral positions that require the adoption of some variety of moral relativism to adequately explain. In the face of the deep and apparently irreconcilable division within our society between those who believe that abortion is morally wrong and those who believe that it is not wrong, it may seem initially plausible to take this as evidence that within our society there is no real agreement over basic moral principles. But the fact that there can be genuine disagreement (perhaps even irresolvable disagreement) about a moral principle such as 'Abortion is morally wrong' does not mean that there could ever be genuine disagreement about the basic moral principle 'Murder is morally wrong'.

The surface similarity of 'Abortion is morally wrong' and 'Murder is morally wrong' does not disguise the rather obvious differences in the way the two phrases are used by parties to the moral debate about abortion. The latter is a basic moral principle that is used to provide the common

a conclusion, but it is not considered necessary, within the confines of this moral dispute, to give any argument with 'Murder is morally wrong' as the conclusion. Sometimes the slogan 'Abortion is murder' may be used by those who oppose abortion. This attempt to identify abortion with an act that is acknowledged by all to be morally wrong only emphasises the comparatively controversial nature of the claim that abortion is morally wrong.

The great differences between the positions of the two parties to this dispute should not obscure the important respects in which they concur. Both parties to the debate are likely to agree on a number of factual matters - that abortion is the killing of a fetus, that a fetus is an entity with certain biological properties and so on. For this to be a moral debate, they must also agree on the basic moral principles and they do, the most relevant one here being the principle, 'Murder is morally wrong.' The contentious issue is whether the killing of a fetus is or is not an instance of murder and *this* is a question about the moral standing of a fetus.

There is a vast literature on this topic, as we would expect with something about which moral diversity is legitimate. I will look briefly at a representative writer from either side of the controversy. In arguments about abortion it is a common move to distinguish between a merely biological or descriptive category such as 'human', to which a fetus would on most views belong, and a moral category such as 'person' or 'moral patient', that is used to identify those entities that have moral standing. The argument will then turn on whether a fetus does or does not properly belong in the second category. The issue becomes a question of whether a fetus has moral standing.

For instance, Michael Tooley, a moral philosopher who argues that abortion is not morally wrong, devotes some effort to his claim that a necessary condition for moral standing is the possession of certain cognitive capacities, namely: "the concept of a self as a continuing subject of experiences and other mental states" and also a belief that oneself is such a subject. (Tooley, p 44) If this criterion is accepted then it will be a further

150

matter to argue that a fetus does not meet Tooley's condition and so does not have moral standing, but Tooley recognises that a significant argumentative move is to establish acceptable criteria for moral standing.

Naturally enough Tooley's condition is not acceptable to those who argue that abortion is morally wrong. So Don Marquis points out that Tooley, or anyone who seeks to fix criteria for personhood on the basis of possession of psychological characteristics, still has to show that the criterion being advanced is morally significant. (Marquis, p 186) Marquis would prefer to do away with the concept of personhood altogether and instead suggests that the criterion for membership in the moral community is having a future "like ours". (Marquis, p 190-1)

The point of these remarks is not to take a position myself on whether or not the fetus has moral standing or on what is the most plausible criteria for moral standing. Instead I want to emphasise here that the apparently irresolvable nature of the dispute between those who oppose abortion and those who allow abortion does not threaten the universal status of a basic moral principle like 'Murder is morally wrong'. Both parties to the debate are committed to this rule. There is no grounds to claim that this dispute exemplifies the clash of two incompatible approaches to morality, each containing their own disparate basic moral principles. The disagreement about abortion is a disagreement about the criteria for moral standing, and this disagreement occurs within the bounds of moral discourse established by universal basic moral principles.

4.2.3 Example two: 'Human slavery is morally wrong'.
A popular example of moral diversity is the allegedly radical difference between the moral stance taken in our society regarding human slavery and that taken in ancient Greek society. This example is used by both moral relativists such as Harman as yet another case of moral diversity, and also by moral universalists who want to try to support the claim that other cultures

151

can make errors about morality. We have already seen Pojman doing this in section 2.3, and moral universalist Walter Stace is another, insisting:

> If slavery is morally wicked today, it was morally wicked among the ancient Athenians, notwithstanding that their greatest men accepted it as a necessary condition of human society. Their opinion did not make slavery a moral good for them. It only showed that they were, in spite of their otherwise noble conceptions, ignorant of what is truly right and good in this matter. (quoted in Solomon, p 546)

However, in this section I argue that the ancient Greeks were not so radically different from us morally as either the moral relativist or the moral universalist might want to suggest. Those ancient Greeks who were concerned to be moral at all followed the same basic moral principles that we follow and that anyone of any culture must follow if they are to be moral. It cannot be the case, as the quotation from Stace above implies, that the ancient Greeks rejected the moral principle 'To enslave a moral patient is morally wrong.' If they were concerned to be moral then they accepted Principle H, and if they accepted Principle H then they would have to have had some peculiar beliefs about the condition of slavery not to recognise that it is a great harm. Being closely acquainted with the reality of slavery as they were, they clearly recognised the wretchedness of the condition and the harm inflicted upon those subjected to slavery. The fact that they did not consider it morally wrong to keep human slaves is simply an indication that they did not consider all humans to have moral standing. To the ancient Greeks, not all humans were moral patients and it was morally acceptable to enslave those who were not.

There is a great difference between the two moral principles: 'To enslave a moral patient is morally wrong,' and 'Human slavery is morally wrong.' The first may be considered a universally acceptable basic moral principle that derives from Principle H, while in contrast the second puts forward the claim that all humans (at least) are moral patients. The ancient Greeks could subscribe to the first principle and have moral discussions about slavery without necessarily having to accept the second principle, and I suggest that

this is in fact just what they did. Among the ancient Greeks it was not generally thought that all humans were moral patients and it was believed that those humans that were not moral patients could rightfully be enslaved. If I am right about this then we would expect to find some concern among their moral thinkers when slavery was practiced on those humans who were considered to be moral patients. This was indeed the case.

Ancient thought on slavery tended towards one of two views. There were those who maintained that slavery was merely the unfortunate consequence of fate - perhaps being captured in a war, or kidnapped by pirates, or even being the child of slaves. Anyone, even the greatest of kings, might through ill luck become a slave and for this reason being a slave did not necessarily mean one was inferior to one's master. (Garlan, p 120) The fact that someone was a slave said nothing about the moral status of that individual. Perhaps it was morally wrong that they were enslaved, perhaps not – their status as a slave was nothing more than the result of arbitrary fortune. But such a perspective was by no means the most prevalent.

The dominant view was the doctrine of the natural slave. According to this doctrine there was a moral distinction between the natural slave and the naturally free human. Those humans who were not Greek were naturally inferior linguistically, intellectually, and culturally and were thus naturally suited for slavery. There was a distinction drawn by nature itself between those humans who were moral patients and those that were not, and there was nothing morally wrong in treating the latter group as objects to be owned and resources to be exploited, like horses or cattle. In fact, one common term for a slave was *andrapodon*, a 'being with human feet', an object on the same moral level as *tetrapodon*, a 'being with four feet'. (Garlan, p 20)

While there were constraints on how a master might treat a slave, these were not to protect the slave, but were intended to protect others in the moral community from this master, who otherwise might develop into a tyrant. (Wiedemann, p 23) There was condemnation of those who maliciously injured their slaves but only because this was immoderate behaviour that

would attract the wrath of the gods, not because human slaves had any moral standing (Garlan , p 41).

However, it was not impossible for once free Greek citizens to become the slaves of other Greeks. Prisoners of war, for instance, legally became the property of the victor, and could be reduced to slavery if the victor so wished. An Athenian citizen visiting a foreign city that was besieged and captured might become a slave and be later sold in Athens. Such a case would seem to be a counter example to the doctrine of the natural slave. Hence a further distinction was deemed necessary – among those who were actually slaves there was the natural slave and the legal slave. The natural slave was the slave for whom slavery was appropriate and morally unproblematic, whereas the legal slave was the slave for whom slavery was not appropriate and was morally problematic. Most forthright on this distinction is Aristotle. He says: "For the words slavery and slave are used in two senses. There is a slave or slavery by convention as well as by nature ... no one would ever say that he is a slave who is unworthy to be a slave. Were this the case, men of the highest rank would be slaves and the children of slaves if they or their parents chanced to have been taken captive and sold." (*Politics* 1.1255a) Others such as Plato and Xenophon's Socrates also support this distinction. (Garlan, p 50)

The fact that the ancient Greeks felt the need to make this strained distinction between the natural slave and the merely legal slave makes it clear that for them human slavery was in general only morally acceptable when it was practiced on those humans that were not moral patients. It was at the least morally problematic to enslave a human who was a moral patient.

Furthermore, as it is possible for us to intelligibly disagree with the ancient Greek view that not all humans were moral patients, it would seem to be in principle possible for a representative of our society to have a moral discussion about this with a representative of ancient Greek society. Any such moral discussion would take place against the background of the commonly accepted basic moral principles, principles that I suggest were accepted no less by the ancient Greeks than by ourselves. So it seems that our

154

moral perspective is not so radically different from that of the ancient Greeks after all. Again, as in the example of abortion, the real differences in moral view here concern the limiting moral principles, and the contents of the basic moral principles are not put under threat by this difference.

4.2.4 Harman reconsidered.

Of all the three moral relativists I looked at in chapter 2, Harman comes closest to recognising the important role played by the question of who or what has moral standing. For instance, he says: "[People's values] differ with respect to the extent of the moral community: some restrict it to family and friends; others include all people of a certain race or caste or country; some include all people of whatever race or class; others count animals and even plants as part of the moral community to be protected by the moral rules." (1996, p 17) Yet despite this acknowledgement he does not say what he thinks the impact of this is on the case he is trying to make for moral relativism. Using Harman's terminology, people's values about the extent of the moral community are relative to their moral framework, but I have tried to show how this relativity is compatible with the universality of the basic moral principles.

Harman does at times seem to say something quite similar to my account: "When moral prohibitions on harming and killing and lying exist, they are sometimes supposed to apply only with respect to the local group and not with respect to outsiders." (1996, p 9) In a footnote on the same page he even says: "There will be universal truths about moralities ... perhaps all moralities have some rules against killing, harm, and deception ... the existence of universal features of morality is compatible with moral relativism." However, this does not induce him to try to look for any sort of compromise between his version of moral relativism and universal features of morality. Instead he says: "Any universally accepted principle in this area must verge on triviality, saying, for example, that one must not kill or harm members of a certain group, namely the group of people one must not kill or

155

harm" (1996, p 9). The principle Harman accuses of triviality here is very similar to principle H, but if it seems trivial this is only because we are already well acquainted with the moral dimension of life. To those not so acquainted the principle will not be trivial, as it contains the significant information that harming and killing are marked off from other activities in that there are negative connotations attached to doing them to a certain group. Of course the principle is not fully informative unless it is also known who the members of this group are, to which Harman (and I) would reply: That is relative to the context.

Views about who or what have moral standing look to be instances of Harman's inner judgements. Consider the example I have looked at earlier, the employee of Murder, Incorporated, who "was raised as a child to honor and respect members of the "family" but to have nothing but contempt for the rest of society." (1982, p 191) Such an individual clearly considers that only other members of Murder, Incorporated have moral standing, so that only they are moral patients. This is consistent with the individual being committed to the basic moral principles. They may uphold all of the basic moral principles scrupulously in all of their dealings with the other members of Murder, Incorporated. We can intelligibly disagree with this individual about whether the set of moral patients is restricted to the members of Murder, Incorporated, which means we share agreement with them on the basic moral principles.

Harman's theory of inner judgements, shorn of his account of implicit bargaining, is not incompatible with my account. Inner judgements about whether "it is wrong *of someone* to do something" (1982, p 192, italics in original) may perhaps legitimately vary according to that individual's views about who or what is a moral patient. But the sense of the word 'wrong' here is for Harman a different one from the sense of the word as it occurs in a basic moral principle such as 'To harm a moral patient is morally wrong.' (See section 2.5.2 for his fourfold distinction of the sense of this word) The latter is for Harman a normative judgement about what "ought to be" rather

than an inner judgement about what someone "ought to do." In so far as Harman restricts his moral relativism to inner judgements, as he does in the 1975 paper 'Moral Relativism Defended', then generally speaking it lends itself well to the outlook I am proposing, as it does not question the universality of basic moral principles. I would still disagree with some of the details of that paper however, such as the notion of implicit bargaining, as I have explained in section 2.5.

4.3 Limiting moral principles and context-dependence

4.3.1 Diversity of starting points.

It might be said that the view I have been presenting does not achieve a genuine compromise between the moral relativist and the moral universalist, as all I have done is to relocate the problem of the dispute between them from one area to another. Now the dispute is about what entities are moral patients. There may still be disagreement between the moral relativist and the moral universalist about the context-dependence of any answer to this question. The moral universalist may say that there is only one universally correct set of moral patients, whereas the moral relativist may say that the set of moral patients depends on the context. However, it is in the nature of a compromise that both parties must give some ground to the other and I suggest that with respect to this question it is the universalist who must give ground to the relativist, just as it Part Three it was the relativist who had to give ground to the universalist.

It appears that, unlike basic moral principles, limiting moral principles are context-dependent. One reason to think that universalism about limiting moral principles is not plausible is the controversy among universalists as to which are the correct limiting moral principles. The considerations advanced in section 3.3.2 for the universality of basic moral principles do not appear to apply to limiting moral principles. Accepting that, for instance, all and only human beings are moral patients is not an presupposition that must be made before moral discussion can be undertaken, as moral discussion can occur

about whether non-humans are moral patients. It is possible for people who disagree about limiting moral principles to never the less engage in moral discussion, and in fact moral discussion may occur about the very acceptability of a limiting moral principle.

Furthermore, it cannot be said about limiting moral principles that they are presuppositions, as can be said about basic moral principles, because while the correct basic moral principles are already more or less known and accepted, the correct limiting moral principles are not. Among universalists there are numerous different competing views about which are the correct limiting moral principles and choosing between them requires much discussion and evaluation of alternatives. For instance, Marquis and Tooley are both universalists when it comes to a limiting moral principle relating to human fetuses, and yet both disagree about what this principle is. In order for a universalist's position on a correct limiting moral principle to be credible, they would first need to defeat by argument every other universalist position that advocates a different limiting moral principle. It is in the very nature of their positions that two moral universalists who offer different limiting moral principles cannot both be correct. This is why I think that here is a plausible location of (at least part of) the moral diversity underlying the relativist's Conviction R.

This does not mean that there are no limits at all on the kinds of proposals that may be made about who are moral patients. It would be difficult to argue against the idea that, at a minimum, all moral agents are also moral patients. Furthermore, any proposals about moral patienthood must meet basic conditions of intelligibility, consistency, non-arbitrariness, and so forth that are required for any proposals in any field of legitimate discussion. It is not the case that some completely arbitrary collection of things could credibly be put forward as the set of moral patients. There is no reason not to be committed to the highest possible standards of argumentation when engaging in moral discussion about a limiting moral principle.

Conviction R, according to my interpretation, is a response to limiting moral principles being context-dependent. In the context of different cultures or groups the very best arguments and most careful moral discussion will lead to different limiting moral principles. To clarify this I will introduce the notion of a *starting point*. What I mean by a starting point is the complex of fundamental beliefs and shaping attitudes concerning morality (over and above the presuppositions that are the basic moral principles) that we bring to any discussion of controversial moral issues. A starting point will not necessarily be clung to dogmatically and sometimes may be abandoned or revised during the course of a moral discussion, but it will be given up reluctantly and only under argumentative pressure. Some of the beliefs and attitudes that comprise one's starting point may be widely held within one's culture while some may be peculiar to smaller groups or even to individuals. Some of the attitudes may not even be explicitly expressed but may be revealed by one's reactions to thought experiments involving moral situations.

My claim is that different starting points lead to a commitment to different limiting moral principles, even when the route is via the most impeccable argumentation. In saying that limiting moral principles are context-dependent, the starting point is the context I have in mind. As an illustration of this, I will look at another of the hotly contested questions about moral patienthood, whether animals are moral patients. I will examine the positions of two universalist moral philosophers, Paul Taylor and Peter Carruthers, who oppose each other on this question, and show how their different starting points lead them to adopt different limiting moral principles. Although both Taylor and Carruthers have internally self-consistent positions, as moral universalists the conclusions they present exclude each other. I hope to show that the limiting moral principles regarding animals that each proposes depends on the context of their starting points.

4.3.2 Taylor and the inherent worth of animals.

Taylor has been influenced by the approach of Kant, who was of course also a universalist with regard to membership in the moral community. Taylor follows Kant in searching for "something whose existence has in itself an absolute value, something which as an end in itself could be a ground of determinate laws," (Kant, p 90) in order to make this the criterion for moral patienthood. This "something" may more lately be known as inherent worth, or intrinsic value. For Kant rationality was the only feasible criterion for moral patienthood, as only rational beings could exist as ends in themselves and not just as means to be used to further the plans of others, and so only rational beings could have inherent worth.

In contrast to this Taylor argues for a "life-centred system of environmental ethics", the crucial part of which is "the taking of a certain ultimate moral attitude toward nature," (Taylor, p 98) a moral attitude that he calls the attitude of respect for nature. He says that this parallels the Kantian respect for persons as rational beings which are ends in themselves (Taylor, p 103), although Taylor of course is breaking with Kant in not accepting that rationality alone is that which confers inherent worth. Kant himself is unequivocal: "Beings whose existence depends, not on our will, but on nature, have none the less, if they are non-rational beings, only a relative value as means and are consequently called *things*. Rational beings, on the other hand, are called *persons* ". (Kant, p 90-1, Italics in original) This is a major point of difference between Taylor and his avowed influence, and as they are both universalists, if Taylor is correct about the criteria for moral patienthood then Kant must be incorrect. This aspect of universalism recurs every time universalists disagree about limiting moral principles.

Taking Taylor's moral attitude of respect for nature involves accepting that living beings are moral patients due to a "recognition of their inherent worth." (Taylor, p 99) According to Taylor, inherent worth is the criteria for moral patienthood and while it is apparently not the case that all living beings have inherent worth, "wild organisms, species populations, and communities of life in the Earth's natural ecosystems," (Taylor, p 100) do have inherent

worth. For such living beings as these, there is something morally valuable in the realisation of those states of affairs that allow them to flourish. As the flourishing of these living beings may be assisted or retarded by the actions of human moral agents, we are thus morally obliged to consider such possible impacts when undertaking actions that may affect them. This amounts to an insistence that the kinds of living beings specified in the quote above are to be regarded as moral patients

For Taylor the recognition of the inherent worth of living beings follows from the attitude of respect for nature already mentioned above. Providing intellectual support for this attitude is a belief system he calls the biocentric outlook on nature. Possessing this belief system makes the attitude of respect for nature, "the only suitable or fitting attitude." The biocentric outlook on nature is the argumentative heart of Taylor's position and it consists of four main elements:

> (1) Humans are thought of as members of the Earth's community of life, holding that membership on the same terms as apply to all the nonhuman members. (2) The Earth's natural ecosystems as a totality are seen as a complex web of interconnected elements, with the sound biological functioning of each being dependent on the sound biological functioning of the others … (3) Each individual organism is conceived of as a teleological center of life, pursuing its own good in its own way. (4) Whether we are concerned with standards of merit or with the concept of inherent worth, the claim that humans by their very nature are superior to other species is a groundless claim and, in the light of elements (1), (2), and (3) above, must be rejected as nothing more than an irrational bias in our own favor. (Taylor, p 103)

It is in his defense of the fourth notion that Taylor does what is, from my perspective, his most interesting work. I will take the claim to human superiority he mentions as being the claim that humans alone are moral patients. Taylor has a number of arguments against this claim.

First, he argues that it is begging the question to assert that humans deserve to be the only moral patients because humans possess abilities, such as the capacity to do abstract reasoning, that nonhumans do not possess, and that the possessors of such abilities have greater merit than those without

them. Taylor points out that it is only according to some set standard that judgements of merit can be made, and it would be an error to assume that only the sorts of abilities that humans value can be judged to be of merit when it comes to deciding moral patienthood. So while the capacity to do abstract reasoning may be more meritorious according to human standards than the ability to climb trees swiftly, this does not help to answer the question of whether humans are more meritorious than monkeys, unless the answer is going to be merely that humans are more meritorious according to human standards, which begs the question. Merit as a criteria for moral patienthood is only plausible in the presence of some clearly independent standard of meritoriousness. Taylor believes that he has supplied such a standard in the first three elements of the biocentric outlook on nature quoted above, and according to this standard humans are no more meritorious than many other living beings.

Taylor also identifies the conceptual confusion attendant on the notion that only humans may be moral patients because only humans may be moral agents. The idea here is supposed to be that because only humans may act morally this makes humans morally superior to nonhumans, and such moral superiority is a necessary condition for moral patienthood. But as Taylor says, those that are not moral agents may not be morally evaluated at all, and so it is not appropriate to claim a relationship of moral superiority over them. It makes no sense to say that because moral agents may act morally, moral agents are morally superior to those who are not moral agents, any more than it makes sense to say that because moral agents may act immorally, moral agents are morally inferior to those who are not moral agents. If the set of moral patients is to collapse into the set of moral agents then some other argument than this one that Taylor criticises must be given.

Another standpoint that Taylor assesses as faulty is one alleging that humans have greater inherent worth than nonhumans and this is why humans counts as moral patients whereas nonhumans do not. According to such an account the criteria for moral patienthood would be the possession of a

certain amount of inherent worth. Taylor is concerned that the notion of differing degrees of inherent worth derives from non-egalitarian class-structured societies in which humans were seen as having varying amounts of inherent worth. He sees it as inconsistent and an expression of irrational prejudice against nonhumans if such systems are rejected now as being an unjust way of evaluating humans while at the same time being seen as an acceptable way to evaluate nonhumans. The objectionable element seems to be the arbitrary nature of assigning differing amounts of inherent worth. In the absence of any means of measuring such an intangible quality it is left to the accident of one's birth to determine this. Some more principled criteria than this for moral patienthood is required.

Taylor's starting point is taking the attitude of respect for nature. Once someone takes this attitude, they are inclined to recognise the inherent worth of living beings and appreciate the biocentric outlook on nature and they may be led quite painlessly by the arguments Taylor provides to the view that it is a universally correct limiting moral principle that animals are moral patients. There are no apparent flaws in the reasoning Taylor uses to reach this view and each of the three arguments looks to make an eminently respectable point: 1) If moral patienthood depends on merit then some independent standard for measuring such merit is required. 2) If moral patienthood collapses into moral agenthood then some good reason for this must be given. 3) An arbitrary notion such as differing amounts of inherent value is not an acceptable criteria for moral patienthood. If someone shares with Taylor the attitude of respect for nature then their starting point will be very close to his. Within the context of this starting point, they may accept the limiting moral principle, 'To harm an animal is morally wrong,' perhaps also accepting other limiting moral principles such as 'Eating meat is morally wrong.' Needless the say, Taylor's is not the only starting point on an issue such as this.

4.3.3 Carruthers and the indirect moral significance of animals.

Peter Carruthers is a universalist who has a very different starting point from Taylor on the question of the moral patienthood of animals and not surprisingly comes up with very different results. Carruthers proposes two criteria that a moral theory must meet in order to be acceptable. Firstly it must provide a conception of what moral knowledge is and how people come to be motivated to act morally. (Carruthers, p 23) Secondly it must deliver normative principles that do not stray too far from, "our considered moral beliefs". (p 24) These considered moral beliefs are the essential ingredient of Carruthers' starting point.

Carruthers argues that the approach that best fulfills his two criteria is a contractualist approach. He characterises this as, "an attempt to justify a system of moral principles by showing that they *would* be agreed upon by rational agents in certain ideal circumstances." (p 35-6, Italics in original) According to Carruthers, the contractualist approach meets the first criterion above because he believes that humans have an essential need to be able to rationally justify their actions, including their moral behaviour, to one another, and that this need is "innate (inborn), in such a way as to emerge gradually at a given stage in maturational development" (p 45). He thinks the second criterion is met because he believes that the main normative principle that seems to result from contractualism, a principle of respect for autonomy, in other words not interfering with one another's plans and projects except in so far as the plans and projects of others are interfering with one's own, has "a powerful intuitive appeal" (p 40). Much as Tayler's biocentric outlook on nature was the theoretical side of his attitude of respect for nature, contractualism is the theoretical side of Carruthers' considered moral beliefs.

For Carruthers, taking a contractualist approach results in animals not being accorded the status of moral patients, because they are not rational agents. "Since it is rational agents who are to choose the system of [moral] rules, and choose self-interestedly, it is only rational agents who will have their position protected under the rules." (p 98-9) He considers and rejects two ways in which animals might be accorded status as moral patients under

164

contractualism. The first way would be to assign to some rational agent or agents the task of representing the interests of animals in the negotiations during which moral rules are formulated and agreed upon. This is rejected because it would lead to animals being accorded equal moral status to that of humans, and Carruthers regards this as not acceptable to the considered common sense views of ordinary people. (p 99, 72) It should be obvious by now that Carruthers' starting point is in complete contrast to Taylor's starting point.

The second way a contractualist might allow animals as moral patients would be for the rational agent to regard knowledge of their species as one of the pieces of information about themselves, along with their race, gender, level of intelligence, and so forth, that is 'bracketed' as irrelevant whilst negotiating the moral rules to abide by. Carruthers rejects this because it "would destroy the theoretical coherence of Rawlsian contractualism," that he is ultimately relying upon. For the contractualist, "[m]orality is viewed as constructed *by* human beings, in order to facilitate interactions *between* human beings." (p 102, Italics in original) According to this understanding of a contractualist morality, being human is a necessary condition for being a moral patient.

Carruthers also confronts the question of whether being human is a sufficient condition for being a moral patient and concludes that it is. Some humans, for instance the very young, the senile, and those with mental disorders, are not rational agents and yet Carruthers argues that they are never the less moral patients. He has two arguments for this.

The first is a slippery slope argument. According to Carruthers there is a slippery slope that leads from the position of denying moral patienthood to those humans who are not rational agents down to the position of denying moral patienthood to some humans who are rational agents. (p 114) This is because there are no clear boundaries between non-rational humans and rational humans. For instance, it is difficult to say at what point the non-rational human infant becomes the rational human adult. Unless there is an

unambiguous declaration that all humans are moral patients there will be opportunities for the unscrupulous to abuse the imprecise nature of the distinction between non-rational humans and rational humans. (p 115-6)

Carruthers' second argument for according moral patienthood to all humans regardless of whether they are rational or not is an argument from social stability. Many moral agents are psychologically attached to non-rational humans to such a degree that a moral system that with-held moral patienthood from non-rational humans would create social instability, "in that many people would find themselves psychologically incapable of living in compliance with it." (p 117)

So, according to the results of Carruthers arguments animals are not moral patients, but this does not mean that the way they are treated is of no moral concern at all. Because some humans are animal lovers who care about what happens to animals, "causing suffering to an animal would violate the right of animal lovers to have their concerns respected and taken seriously." (p 107) He acknowledges that this position conflicts with two common views about the suffering of animals, firstly that causing suffering to an animal would be morally wrong even if no animal lovers were ever to find out about it, and secondly that causing suffering to an animal is morally wrong "because of what is done to the animal," (p 108) and not because of the distress it may cause to animal lovers.

Carruthers accounts for these qualms by maintaining that the suffering of animals has indirect moral significance, "by virtue of the qualities of character that [animals] may, or may not, evoke in us." (p 154) Sympathy for animal suffering is "merely a side-effect" of the sympathy that moral agents are required to develop for genuine moral patients. Animals themselves are not moral patients so cruelty to them is not strictly speaking morally wrong, but cruel treatment of animals demonstrates flaws in moral character that may also manifest itself in cruel treatment of actual moral patients, that is, humans. It is because of this possibility that cruelty to animals is to be

condemned and why it is right to feel sympathy for any animals that are treated cruelly. (p 156-7)

Throughout his book Carruthers presents many different thought experiments designed to appeal to a person with a similar starting point to his. A great deal of the weight of his position is borne by Carruthers' assumption throughout that his reader will share his intuitions regarding the different situations he portrays, hence sharing his starting point. Two of the most important of Carruthers' thought experiments are described below.

It is one of the central tenets of Carruther's starting point that, "most ordinary people do not seriously rate animal suffering at all, in comparison to the sufferings of human beings." (p 72) He presents the thought experiment of Saul the sadist to show this. (p 67) Saul the sadist has imprisoned and is constantly torturing a number of different animals and one human being. If given the opportunity to rescue one and only one creature from Saul's dungeon, Carruthers thinks it incontrovertible that everyone would accept that there is a moral obligation to rescue the human rather than any of the other animals. Even if the human would not undergo any more suffering than any of the other animals because of their predicament, nor would the human live any longer if rescued than any of the other animals, never the less, "the intuition remains that it would be unforgivable to do anything other than rescue the human." (p 68)

In a related thought experiment designed to show that humans are considered more morally important than animals, he says that if faced with a choice between saving from a fire one unconscious, elderly human with no friends or family, or saving many conscious, caged dogs from the fire, "[o]ur common-sense, pre-theoretical view is that it would be very wrong to place the lives of many dogs over the life of a single (albeit old and friendless) human." (p 96)

Obviously the way in which one reacts to these thought experiments depends on one's starting point. Although I do not want to comment on the plausibility of Carruthers' starting point, I do not share his confidence that

everyone has the same starting point as himself. It may seem inevitable to Carruthers that his own starting point represents unqualified common sense, but I would suggest this is just the illusion of perspective. Of course, many people would share Carruthers' starting point, and react in the same way that he does to the thought experiments he describes. Within the context of this starting point, someone would reject the limiting moral principle, 'Animals are moral patients.'

Taylor and Carruthers put forward their considerations from within the context of different starting points. These different starting points lead to a commitment to different limiting moral principles with regard to the question of whether or not animals are moral patients. Their different starting points do not seem to be the result of errors about the facts on the part of either of them. Nor do their routes from their respective starting points to their respective limiting moral principles seem to be the result of faulty reasoning on the part of either of them. Taylor and Carruthers are only two of the many theorists who have proposed various, often incompatible, limiting moral principles with regard to the question of whether or not animals are moral patients.

I believe that the notion of different starting points providing the context from which different theorists reason to different limiting moral principles provides further grounds for my claim that limiting moral principles are context-dependent and not universal. However, context-dependent limiting moral principles are compatible with universal basic moral principles, and this distinction between different types of principles is how I seek to achieve my coherentist goal of accommodating both Conviction U and Conviction R.

4.4 Ordering basic moral principles
4.4.1 Further relativism.

Although disagreements about limiting moral principles form a considerable share of the body of cases of moral diversity presented by the moral relativist

as support for their claim that moral principles are context-dependent, some of the examples of moral diversity highlighted by the moral relativist do not seem to be instances of disagreement about who or what are moral patients.

For instance, Wong provides the example of disagreement concerning whether there is a moral requirement to redistribute assets from the rich to the poor, and considers this disagreement to be something that the universalist is hard pressed to account for. (Wong, p 146-153) It would be a misrepresentation to claim that this was a disagreement about the criteria for moral patienthood. Someone who denies that there is a moral requirement to redistribute assets to the poor is not thereby denying that the poor are moral patients.

Similarly, Harman points out that there is moral disagreement on the question of whether killing is worse than letting die. But the person who says that killing *is* worse than letting die is not necessarily saying that this is because someone who is let die does not have the status of moral patienthood. Similar observations could be made about some of the other moral disputes that Harman marshals in support of moral relativism, such as disagreement about the value of artifacts as compared with human life and about how much help one is morally obliged to give another. (Harman 1996, p 8-11) Disputes about these matters do not seem to be disputes about membership in the moral community, and claiming that they are would distort what is happening when such disagreements occur. It is in these disputes that I see another possibility of genuine grounds for moral relativism.

Earlier I assumed for the sake of convenience that monism with regard to basic moral principles is the case and that Principle H: 'To harm a moral patient is (pro tanto) morally wrong,' is the most general of the basic moral principles. However, the word 'harm', as it occurs in Principle H, is a word that may receive various interpretations, because a moral patient may be harmed in many different possible ways. Often a disagreement may arise about which of a number of different possible forms of harm are more

morally reprehensible. This field of dispute concerns the order of priority of basic moral principles, and I think that such dispute accounts for the examples of moral diversity that Harman and Wong give above.

For instance, a disagreement about whether it is morally worse to kill a moral patient or to let a moral patient die is a disagreement over which sort of harm, killing or letting die, is more morally reprehensible. The question is really whether 'To kill a moral patient is (pro tanto) morally wrong,' and 'To let a moral patient die is (pro tanto) morally wrong,' have the same order of moral priority or whether one of them has a higher order of priority than the other, and if so, which way around this order of priority is. Arguing that one moral principle has a higher order of priority than another does not mean that one is rejecting the latter of the principles as a moral principle. Stating the order of priority among two basic moral principles may be done using a metaprinciple, but this is not a requirement. For example: 'Killing a moral patient rather than letting a moral patient die is (pro tanto) morally wrong,' is a moral principle, although not a basic moral principle.

If the simplifying assumption of monism is not made, then the claim about the context-dependency of the order of priority of basic moral principles becomes stronger. If pluralism is the case and the basic moral principles ultimately derive from more than one general basic moral principle, then questions about the order of priority of the basic moral principles will be more and not less prominent than in a monist system and will take on more importance as a source of moral diversity. For instance, in a pluralist system that contained, as well as Principle H, a principle saying that to treat moral patients disrespectfully was morally wrong, not only will there be disagreements about whether one sort of harm is worse than another, but there will also be disagreements about whether one sort of disrespect is worse than another, and furthermore there may also be disagreements about what sorts of harm are worse than what sorts of disrespect and what sorts of disrespect are worse than what sorts of harm.

The universal nature of basic moral principles is not threatened by acknowledging context-dependence about their order of priority. If, in the context of one group, basic moral principle A is thought to delineate an area of more significant moral concern than basic moral principle B, while in the context of another group B is considered more significant than A, this does not mean that either group has rejected the basic moral principle they consider to be the less important one. Both A and B are still universal moral principles, with pro tanto weight in the moral discussions of both groups.

It may also be the case that in some contexts a limiting moral principle will be considered to have a higher order of priority than a basic moral principle. This does not change the status of the latter as a universal principle that has some weight whenever a moral discussion occurs and the former as a context-dependent principle that has weight only in a certain context. For instance, in the context of a group that considers sentience as the criterion for moral patienthood the limiting moral principle, "Killing a sentient being is morally wrong," may weigh more heavily in their moral discussions than the basic moral principle, "Lying is morally wrong." Members of such a group may agree that it is morally acceptable to lie in order to prevent the killing of a sentient being. Never the less, the first of these principles only has this weight within the context of such a group, while the second principle has weight within any moral discussion.

4.4.2 Arrington on priority.

Arrington seems to take a universalist position on this matter, claiming that his rules of moral grammar, "usually have a built-in index of priority". (Arrington, p 289) In this section I will consider his argument for this index. He does concede that this notion is not very plausible when it comes to many of the more difficult cases in which there may be disagreement about how to judge the order of priority when his rules of moral grammar appear to conflict, such as a situation where we are faced with a choice between telling a lie or breaking a promise. About such cases he says: "The order of priority

171

among our moral rules, then, is limited, flexible, and imprecise, and it is the source of many of our moral puzzles and our most intractable moral dilemmas." (Ibid) Despite this concession, Arrington wants to maintain that over and above his rules of moral grammar there is a set of grammatical metarules about these rules. The purpose of the grammatical metarules is to rank the order of priority of some (but not all) of the rules of moral grammar with respect to some (but not all) of the other rules of moral grammar. He gives two examples of these grammatical metarules: "One ought to tell the truth unless doing so would cause grave harm to others," and, "One ought to keep one's promises unless doing so would involve taking the life of an innocent person" (Arrington, p 302).

However, these two grammatical metarules do not seem to add anything to Arrington's account except an extra layer of complexity. They do not provide a "built-in index of priority" capable of solving the sort of disagreements we are concerned with here, because they are too vague. Phrases such as "grave harm" and "innocent person" provide plenty of room for exactly the sort of disputes that I have been claiming are responsible for moral diversity. Just how grave would the harm have to be in order for it to be morally worse than being lied to? This is a question about the order of priority of basic moral principles. Who counts as an innocent person? This is a question about moral patienthood.

Arrington also seems to have in mind a third grammatical metarule to the effect that 'Murder is morally wrong' has a higher order of priority than 'Lying is morally wrong'. He says: "That killing a person is a more serious offence than lying to him is not a matter we ever debate. It is a given of the moral life" (Arrington, p 291). Arrington seems to be implying here that a discussion about whether murdering someone is worse than lying to them could never reasonably occur, that such a discussion must necessarily be conceptually confused. This appears far too strong a claim, and does not seem to be at all a given of the moral life.

If it happens that circumstances arise in which one is faced with a choice between murder and lying then we may be willing to say that it is likely such circumstances will demand we make a moral judgement that to lie is the lesser of the two evils in this case. But it is the circumstances of the case that we will look to in making this moral judgement, and not an index of priority contained within a set of grammatical metarules about the rules of moral grammar. Despite what Arrington says, there is no doubt that we can conceive of circumstances in which it is possible to conduct a discussion about whether it is morally worse to murder someone than to lie to them. It is not so difficult to invent such a case. Consider a lie about a person's parentage or the circumstances of their birth. This lie may induce in the person such extreme feelings of self-loathing that they honestly believe they would be better off dead, and may fervently wish to be killed. It seems at least intelligible to ask whether such a lie is or is not morally worse than the murder of this person, but if Arrington is correct then this question is not a legitimate one, as we do not ever debate whether it is worse to kill someone than to lie to them. I see no reason why a debate about the above case could not occur. Even if our discussion leads to the moral judgement that in this case it would still be morally worse to murder this person than to tell them such a lie, this moral judgement will be the result of our discussion, and is not a given of the moral life as Arrington claims

Arrington might reply that this objection mistakenly treats his grammatical metarules as substantive principles (See section 2.9.3). Just as his rules of moral grammar may be outweighed in some particular circumstances without this meaning that they lose their relevance in establishing what is of moral concern, so his grammatical metarules may be outweighed in some particular circumstances, without this meaning that they lose their relevance in establishing the order of priority between two rules of moral grammar. However, this reply does not take into account what has happened in the latter instance. Consider again the example in the previous paragraph and imagine that the only possible alternatives available to the

moral agent are to murder the moral patient or to tell the moral patient the lie described above. If moral discussion leads us to make the moral judgement that the murder is the lesser of the two evils in this case, then this moral judgement may allow us to reluctantly excuse this particular murder, but it does not damage the universal nature of the principle that murder is morally wrong. Murder is still morally wrong, even in this particular case where it is less wrong than telling a lie. But this moral judgement would damage the allegedly universal nature of Arrington's grammatical metarule that 'Murder is morally wrong' has a higher order of priority than 'Lying is morally wrong'. We cannot say that murder is still morally worse than telling a lie, even in this particular case in which murder is not morally worse than telling a lie. Arrington's metarules do seem to be substantive principles in this regard.

It seems to me that the metarules that Arrington proposes are no argument against the claim that some instances of moral diversity, such as those examples given by Harman and Wong in the previous section, are really disputes concerning the order of priority of basic moral principles. Given what Arrington says at one point, it seems that he himself ought to concur: "We often disagree over and debate whether a particular rule is to override another on a particular occasion, but this is because there is no a priori order of priority between these two rules. The debate does not challenge either rule as establishing a moral presumption." (Arrington, p 290)

4.5 Conclusion

In this work I have tried to show that the long-standing dispute between the moral relativist and the moral universalist is amenable to treatment through a coherentist approach. Such an approach seeks neither to reject nor to privilege the fundamental convictions of either side but rather tries to accommodate them both within a consistent whole.

The main question that I take the moral relativist and the moral universalist to be in disagreement over is the question I have called Q3: 'Are

174

moral principles context-dependent?' The coherentist answer that I have proposed is: 'There are two types of moral principles, those that are context-dependent and those that are not' Moral principles that are not context-dependent I have called basic moral principles. Underlying many of the basic moral principles (or perhaps all of them, depending on whether monism or pluralism is the case) is the very general basic moral principle I have called Principle H: 'To harm a moral patient is pro tanto morally wrong.' Moral principles that are context-dependent I have called limiting moral principles. These concern who or what counts as a moral patient for the purposes of applying Principle H. I have also suggested that the order of priority between basic moral principles is a matter that is context-dependent.

This approach allows the accommodation of Conviction U and Conviction R - the fundamental convictions behind moral universalism and moral relativism. Conviction U is the conviction that at least one moral principle is universal, that is, it counts whenever moral judgements are made, and it counts in the same way. Principle H is universal in this sense. It is of relevance whenever moral judgements are being made, and it specifies a concern (harm to moral patients) that always counts as morally wrong. Conviction R is the conviction that at least some moral diversity is due to different groups being committed to different moral principles and that not every instance of moral diversity can be explained away as due to mistake or misinformation. On my view some moral diversity is due to different groups being committed to different limiting moral principles, and such commitment is due to these groups having different starting points, rather than making errors or being misinformed.

My approach also allows the major criticisms of each side by the other to be met. The universalist accuses the relativist of promoting an 'anything goes' attitude to morality by letting context override all when making moral judgements. However, basic moral principles provide a place within morality for certain central concerns, most prominently that expressed by Principle H, that may not be relinquished when moral discussion is undertaken. The

relativist accuses the universalist of implying that great numbers of people must be in error about morality and the moral judgements they make. However, on my view, nobody who participates in the practice of morality can be in error about the basic moral principles.

As I noted in Part One, my solution may leave hard-line adherents to either moral relativism or moral universalism unsatisfied. It may be objected that neither Conviction U nor Conviction R have been fully appreciated, but that each have become a mere shadow of themselves at my hands. This may be so, but the position I offer is a compromise between two opposing outlooks, and it is in the nature of a compromise for each side to have to make concessions. It is hoped that the practical benefits to be gained balance the cost of making the concessions that I have argued for. The everyday business of making moral judgements may become a little easier if it is no longer impeded by the prospect of a debilitating confrontation between moral universalist and moral relativist.

Bibliography

Antony, Louise M. 'Natures and Norms' in *Ethics* 111 (October 2000), 8-36.

Aristotle. *The Nicomachean Ethics*. Trans. J. A. K. Thomson. Penguin Classics, 1955.

Arrington, Robert L. *Rationalism, Realism and Relativism: Perspectives in Contemporary Moral Epistemology*. Ithaca, N.Y.: Cornell University Press, 1989.

Baker, G. P. and P. M. S. Hacker. *Wittgenstein: Understanding and Meaning*. Vol. 1 of *An analytical commentary on the Philosophical Investigations*. 4 vols. Oxford: Basil Blackwell, 1980.

Carruthers, Peter. *The Animals Issue*. Cambridge: Cambridge University Press, 1992.

Cohen, G. A. 'Facts and Principles' in *Philosophy & Public Affairs* 31 (Summer 2003), 211-245.

Daniels, Norman. *Justice and Justification: Reflective Equilibrium in Theory and Practice*. Cambridge: Cambridge University Press, 1996.

Darwall, Stephen, Allan Gibbard and Peter Railton. 'Toward *Fin de siecle* Ethics: Some Trends' in *The Philosophical Review* 101 (January 1992), 115-189.

Davidson, Donald. *Inquiries into Truth and Interpretation.* Oxford: Clarendon Press, 1984.

DeGrazia, David. *Taking Animals Seriously.* Cambridge: Cambridge University Press, 1996.

Fried, Barbara H. ' "If You Don't Like It, Leave It": The Problem of Exit in Social Contractarian Arguments' in *Philosophy & Public Affairs* 31 (Winter 2003), 40-70.

Garlan, Yvon. *Slavery in Ancient Greece.* Trans. Janet Lloyd. Ithaca, N.Y.: Cornell University Press, 1988.

Hajdin, Mane. *The Boundaries of Moral Discourse.* Chicago, Ill.: Loyola University Press, 1994.

Harman, Gilbert. 'Moral Relativism Defended' in *Relativism: Cognitive and Moral.* M. Krausz and J. Meiland, eds., Notre Dame, IN.: University of Notre Dame Press, 1982, 189-204.

Harman, Gilbert. 'Is There A Single True Morality?' in *Morality, Reason and Truth: New Essays on the Foundations of Ethics.* David Copp and David Zimmerman, eds., Totowa: Rowman and Allanheld, 1984, 27-48.

Harman, Gilbert and Judith Jarvis Thomson. *Moral Relativism and Moral Objectivism.* Cambridge, Mass.: Blackwell, 1996.

Jacobs, Jonathan. *Dimensions of Moral Theory*. Oxford: Blackwell, 2002.

Kagan, Shelly. *The Limits of Morality*. Oxford: Clarendon Press, 1989.

Kant, Immanuel. *Groundwork of the Metaphysic of Morals*. Trans. H. J. Paton. Published as *The Moral Law: A New Translation with Analysis and Notes*. London: Hutchinson's University Library, 1948.

Kumar, Rahul. 'Who Can Be Wronged?' in *Philosophy & Public Affairs* 31 (Spring 2003), 99-118.

Mackie, J. L. *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin, 1977.

Marquis, Don. 'Why Abortion is Immoral' in *The Journal of Philosophy* 68 (1989), 183-202.

McMahon, Christopher. 'Discourse and Morality' in *Ethics* 110 (April 2000), 514-536.

Moore, George Edward. *Principia Ethica* Cambridge University Press, 1903.

Nagel, Thomas. *What Does It All Mean?* New York: Oxford University Press, 1987.

Nozick, Robert. *Anarchy, State, and Utopia*. New York: Basic Books, 1974.

Nussbaum, Martha C. 'Non-Relative Virtues: An Aristotelian Approach' in *Midwest Studies in Philosophy, XIII: Ethical Theory, Character and*

Virtue. Peter French, Theodore Uehling, Howard Wettstein, eds., Notre Dame: University of Notre Dame Press, 1988, 32-51.

O'Neill, Onora. 'Universalism in Ethics' in *Routledge Encyclopedia of Philosophy.* Edward Craig, ed., London and New York: Routledge, 1998, 535-539.

Papineau, David. 'Functionalism' in *Routledge Encyclopedia of Philosophy.* Edward Craig, ed., London and New York: Routledge, 1998, 805-813.

Plato. *Protagoras.* Trans. C. C. W. Taylor. Oxford: Clarendon Press, 1991.

Plato. *Theaetetus.* Trans. M. J. Levett. Indianapolis, Ind.: Hackett, 1992.

Pojman, Louis P. 'Ethical Relativism versus Ethical Objectivism' in *Introduction to Philosophy: Classical and Contemporary Readings.* Louis P. Pojman, ed., Belmont, CA.: Wadsworth, 1991, 506-516.

Quine, W. V. O. *Word and Object.* Cambridge, Mass.: The M. I. T. Press, 1960.

Rachels, James. 'The Challenge of Cultural Relativism' in *Reason and Responsibility: Readings in Some Basic Problems of Philosophy.* 9[th] ed., J. Feinberg, ed., Belmont, CA.: Wadsworth, 1996, 488-494.

Rawls, John. *A Theory of Justice.* Cambridge: Belknap Press, 1971.

Ross, David. *The Right and the Good.* Philip Stratton-Lake, ed., Oxford: Clarendon Press, 2002.

Singer, Marcus G. 'Universalizability' in *The Cambridge Dictionary of Philosophy*. Robert Audi, ed., Cambridge: Cambridge University Press, 1995, 822-823.

Smith, Michael. 'Realism' in *A Companion to Ethics*. Peter Singer, ed., Oxford: Basil Blackwell Ltd., 1993, 399-410.

Smith, Michael. *The Moral Problem*. Cambridge, Mass.: Blackwell, 1994.

Snare, Francis. *The Nature of Moral Thinking*. London: Routledge, 1992.

Solomon, Robert C. *Introducing Philosophy: A Text With Integrated Readings*. 4th ed., Orlando, Fla.: Harcourt Brace Jovanovich, 1989.

Taylor, Paul. 'Biocentric Egalitarianism' in *Environmental Ethics: Readings in Theory and Application*. 2nd ed., Louis P. Pojman, ed., Belmont, CA.: Wadsworth, 1998, 98-110.

Timmons, Mark. *Moral Theory, An Introduction*. Lanham, MA: Rowman & Littlefield, 2002.

Tooley, Michael. 'Abortion and Infanticide' in *Philosophy & Public Affairs* 2 (1972), 37-65.

Wiedemann, T. E. J. *Slavery*. Oxford: Clarendon Press, 1987.

Wittgenstein, Ludwig. *Philosophical Investigations* Trans. G. E. M. Anscombe. 2nd ed., Oxford: Basil Blackwell, 1963.

Wong, David B. *Moral Relativity*. University of California Press, 1984.

Wong, David B. 'Coping with Moral Conflict and Ambiguity' in *Ethics* 102
(July 1992), 763-784.