

EVIDENCE IN MORAL THEORY AND THE GROUNDING OF REFLECTIVE  
EQUILIBRIUM

by

AVI CRAIMER

B.A. McGill University, 2003

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF THE REQUIREMENTS  
FOR THE DEGREE OF

MASTER OF ARTS

in

THE FACULTY OF GRADUATE STUDIES

(Philosophy)

THE UNIVERSITY OF BRITISH COLUMBIA

June 2005

© Avi Daniel Cramer 2005

## **Abstract**

This work articulates and defends an evidence based account of moral epistemology called Case Evidentialism. From this account, we can explain and justify the bi-directional revision of moral judgments and moral principles that is advocated by the promising methodological doctrine called Reflective Equilibrium. The thesis begins with a survey of meta-ethics in order to find compelling views to represent the meta-ethical camps of realism, subjectivism, and non-cognitivism. I then argue that these three meta-ethical viewpoints imply a common account of moral evidence called Case Evidentialism.

In this account, the contents of our moral judgments are evidentially supported without inference from another claim. In contrast, moral principles are only supported as generalizations inferred from the contents of our moral judgments. Inference from well supported moral generalizations may overrule moral judgments. Therefore, the bi-directional revision of moral principles and moral judgments advocated by Reflective Equilibrium can be grounded in Case Evidentialism.

In the third chapter, I address a series of challenges to the use of moral judgments in moral theorizing which originates in the work of R. M. Hare. I argue that Case Evidentialism has a powerful response to each challenge. In Chapter 4, I look at the doctrines associated with moral particularism. Some of these doctrines imply a scepticism about moral principles that is at odds with Reflective Equilibrium and Case Evidentialism. I explain and criticize the particularist's arguments against moral principles, concluding that particularist scepticism about moral principles is

unsustainable. I also show that there may nevertheless be moral theoretic insight within the particularist's arguments.

Finally, in Chapter 5, I consider whether Case Evidentialism would be acceptable to advocates of Reflective Equilibrium. Although there are some differences between the methodological implications of Case Evidentialism and those of classic Reflective Equilibrium, I argue that Case Evidentialism does a better job explaining most of Reflective Equilibrium's methodological precepts than the standard interpretation. Further, those elements of Reflective Equilibrium which are not captured by Case Evidentialism are epistemologically unattractive.

## Table of Contents

Abstract.....	ii
Table of Contents.....	iv
Acknowledgements.....	vi
Dedication.....	vii
Introduction.....	1
Chapter 1: Meta-Ethics .....	4
1.1 Moral Realism.....	4
1.2 Moral Subjectivism.....	7
1.3 Quasi-Realist Non-Cognitivism.....	14
Chapter 2: Case Evidentialism.....	17
2.1 First Thesis: Non-inferential evidence for particular moral judgments.....	18
2.2 Second Thesis: General moral claims not supported non-inferentially.....	23
2.3 Third Thesis: Moral principles supported through inference .....	27
2.4 Fourth Thesis: Evidence from generalizations to particular claims .....	31
2.5 Grounding Reflective Equilibrium in Case Evidentialism .....	37
Chapter 3: Hare's Attack on Moral Judgment.....	40
3.1 Moral Disagreement .....	41
3.2 Levels of Moral Thought .....	49
3.3 Preferences for Paraworlds .....	52
3.4 Universality of Moral Language.....	56
Chapter 4: Taking Account of Moral Particularism .....	60
4.1 Observational Realism.....	61

4.2 Moral Holism Defined .....	63
4.3 Case Evidentialism and Moral Reasons.....	67
4.4 Arguments for and Against Holism .....	71
4.5 Beyond Contributory Principles .....	78
4.6 Moral Principles.....	83
4.7 Infinite Holism and Scepticism.....	86
Chapter 5: Moral Evidence and Reflective Equilibrium.....	91
5.1 Epistemological Privilege .....	92
5.2 Wide Reflective Equilibrium .....	97
Conclusion .....	105
References.....	108

## Acknowledgements

This research was supported by the Canadian Social Sciences and Humanities Research Council Master's Scholarship. I would like to sincerely thank my research supervisor, Catherine Wilson, for her dedicated support and incisive comments. This work has benefited from discussions with colleagues, family, and friends including: Scott Anderson, John Beatty, Sylvia Berryman, Tim Christie, David Cramer, Daniel Elstein, Kerah Gordon-Solomon, Tim Hanson, Christina Hendricks, Sebastian Hew, Monte Johnson, Rachel Ptashny, Patrick Rysiew, Julian Smith, Daniel Star, Kate Steinmann, Chris Stephens, Dustin Stokes, and Jen Young. A special thank you goes to Marilyn Samuels who graciously proofread the final draft and to Aliza Cramer who gave me tips on formatting in APA style. I am grateful to Green College at the University of British Columbia, which provided a wonderful living environment during the research phase of this project. Green College also gave me the opportunity to give a presentation for their member speaker series from which Chapter 2 of this work has been developed. During the crucial writing phase of this project I stayed at the University of Cambridge. I must thank the Philosophy Department of the University of Cambridge for granting me library privileges during my stay, without which this work could not have been completed. Finally, thank you to all my wonderful teachers at McGill and the University of British Columbia who instilled in me the passion for philosophy from which these ideas have sprung.

*To my parents,  
David and Marilyn,  
whose love made this thinkable*

## Introduction

In the mid-20<sup>th</sup> Century many philosophers doubted that philosophy or any field of theoretical inquiry had special authority to speak on moral matters. For example, Monro's 1967 book *Empiricism and Ethics* opens with the remark,

Moral philosophers may not always agree about whether it is their business, as philosophers, to try to answer such [normative] questions. A good many of them would be shocked at the very idea. ...any systematic attempt to answer such questions will raise some peculiarly difficult problems of method, and...it is with these methodological problems that the moral philosopher is principally concerned. Most philosophers nowadays will add that they are his *sole concern*" (Monro, 1967, p. 3, my emphasis).

If we take Monro at his word about the prevailing opinion of 1960s moral philosophers, we are struck by a strange conundrum into which they have fallen. On the one hand, moral philosophers had rejected the very project of attempting to offer systematic normative moral theories, while on the other hand they were concerned with addressing "methodological problems" with such theories. Yet, there would seem to be no point in addressing methodological problems in a given theoretical domain without having any would be theorists in need of a methodology. Further, if the moral philosophers themselves had decided to give up on the project of normative moral theory, it is difficult to see why they should think that anybody would be interested in their insights into the foundational problems of ethics. If the problems with normative moral theory were, indeed, so dire, then it seems that philosophers should have given up ethics all together or else contented themselves with purely sceptical arguments.



Fortunately, the 1970s saw a renaissance in normative moral theory. A landmark in this philosophical about-face was the publication of John Rawls' *A Theory of Justice* in 1971. In this work, Rawls was able to sidestep the prevailing pessimism about normative moral theory by using a method which he called 'Reflective Equilibrium.'<sup>1</sup> Reflective Equilibrium holds that moral theorizing must rely on both our considered moral judgments about particular cases and our previously held moral principles. The moral theorist seeks to bring her particular judgments and her moral principles into a state of equilibrium, by sometimes modifying her judgments and other times modifying her principles.

In suggesting this method, Rawls made little attempt to address the semantic and metaphysical questions about morality that had occupied philosophers for most of the century. He simply suggested a method which sounded plausible enough on its face and then launched into a detailed presentation of his own moral theory. This turn away from meta-ethics toward substantive moral inquiry was surely warranted given the philosophical climate of the time. When many moral theorists accepted the simple yet ungrounded method of Reflective Equilibrium, the discipline of moral theory flourished, reaching a level of nuanced complexity which was impossible under traditional methodologies such as intuitionism, utilitarianism, Kantianism, and contractualism.

Despite the apparent success of Reflective Equilibrium in facilitating the development of normative moral theory as a discipline, its lack of explicit philosophical

---

<sup>1</sup> The same term, reflective equilibrium, is used in the literature to denote two quite different things. First, it is used to refer to a methodological doctrine which tells us to test moral principles and moral judgments against one another. Second, it is used to refer to the state of optimal coherence between judgments and principles at which this methodological doctrine tells us to aim. To avoid confusion, I will always capitalize the term, writing it as 'Reflective Equilibrium,' when speaking of the methodological doctrine and will use the lowercase 'reflective equilibrium' or just 'equilibrium' to refer to the state of coherence aimed at by somebody who follows Reflective Equilibrium.

grounding led many philosophers to call the method into question. As a result, normative moral theorists continue to be in doubt about the legitimacy of their most promising methodology. Although this has not stopped the rapid development of moral theory in the last 35 years, it has resulted in a continuation of the strange ambivalence towards methodology which Monro identified in 1967. This ambivalence, can only slow the progress of the otherwise promising developments in moral theory. It also runs the risk that inappropriate empirical and mathematical methods from the social or natural sciences will be substituted for distinctively moral theoretic methods in order to achieve a false sense of scientific respectability. The central goal of this thesis is to develop and defend an account of moral theorizing which preserves and clarifies the methodological insights of Reflective Equilibrium, while also providing this methodology with an epistemological and meta-ethical grounding.

Unlike many works in meta-ethics and moral epistemology, I hope this work will be of primary interest and usefulness to moral theorists themselves. I aim to provide moral theorists with a better understanding of the distinctive methodology of their discipline. Such an understanding will enable moral theorists to proceed with confidence in the face of those who argue that moral theory rests on a mistake.

## **Chapter 1: Meta-Ethics**

Any acceptable account of moral theorizing must be consistent with our best understandings of the nature of morality. As such, meta-ethics, the study of the nature of morality, is a reasonable starting place for developing such an account. It would be helpful to be able to start from widely accepted answers to questions such as what type of semantic content do moral sentences have?, what is the relation of moral discourse to truth and falsity?, and to what, if anything, do moral terms refer?

However, the debates over these questions have not yet given us a single set of answers which is clearly superior to all alternatives. We are left with a range of mutually incompatible meta-ethical positions among which it is difficult to decide. Given this situation, it would be best to find an epistemological account of moral theorizing which is neutral among plausible meta-ethical positions. In preparation for this task, I will examine the major positions that have been put forward, rejecting those that are clearly inadequate to capture the phenomena of moral discourse. This process leaves us with three mutually inconsistent meta-ethical views.

### **1.1 Moral Realism**

The most important meta-ethical disagreement in contemporary philosophy is over moral realism. The moral realism debate originates in the rise of the naturalistic world-view embraced by most philosophers. This world-view is seen as debunking the great metaphysical myths of the past. Many philosophers lump "moral values" as among the metaphysical myths in need of debunking, often citing Hume as their inspiration.

Hume's idea that values must be located in human sentiment stems from the conviction that values lie in the human psyche rather than the dictates of God, the demands of pure reason, or the properties of strange entities such as Plato's Forms. Put another way, anti-realism stems from the claim that "[p]roperly understood our natures contain within themselves the ingredients from which we construct the objects of our concern and care, including moral concern and care" (Blackburn, 2002, p. 133).

In contrast, contemporary moral realism "denies the existence of that sharp and significant distinction between fact and value" needed to get anti-realism off the ground (McNaughton, 1988, p. 39). Instead, the realist claims that "moral opinions are beliefs which, like other beliefs, are determined true or false by the way things are in the world" (*ibid.*). The most prominent contemporary version of moral realism, takes an empiricist stance towards our knowledge of moral truths. On this view, our moral judgments about particular situations are observations of moral reality and "moral observation must not be thought of as in a special category of its own, quite unlike any other kinds of observation" but rather on par with other kinds of observations that go beyond simple sensory reports such as "that this cliff is dangerous, that Smith is worried or that one thing is further away than another" (*ibid.*, pp. 56-57). McNaughton argues that since indispensable non-moral observations such as these are not reducible to descriptions of sensations, we should not let the fact that moral judgments are not reducible to descriptions of sensations make us doubt the legitimacy of moral observations. Also, since we accept that non-moral claims can be supported non-inferentially by the fact that we observe them, we should also accept that our moral judgments about particular situations can be supported by the fact that these judgments involve observation of the moral features present in those situations.

Moral realism of the sort advocated by McNaughton is attractive because it naturally makes sense of our moral discourse. Moral judgments often take a propositional form, suggesting that they express beliefs. Further, the observational understanding of concrete moral judgments gives us a familiar empiricist account of our epistemic access to moral truths, helping to ease the fear of metaphysical strangeness which motivates anti-realism. As a result of the naturalness with which observational realism is able to interpret our moral discourse, less explication of the view is needed than for other meta-ethical views. Therefore, having given this brief account of realism, I will now turn to examine some of the many species of moral anti-realism.

What all the forms of anti-realism I will consider have in common is that they argue that moral judgments are not beliefs about the mind-independent world. Invoking the Humean division of the mind into representational states and motivational states, anti-realists claim that motivational states, variously referred to as attitudes, sentiments, desires, or felt reactions, are essential to moral judgment. One reason for this claim is that beliefs by themselves typically do not motivate action, while moral judgments do. For example, if I judge that some action is morally wrong, I am thereby motivated against performing the action, whereas, if I believe that some action will get me money, I am not thereby motivated towards the action, I also require the desire for money. Therefore, moral judgments must bear some intimate relation to action motivating states such as attitudes, feelings, and desires, rather than merely to beliefs about the world. In fleshing out this common claim into an account of moral discourse, anti-realists have taken two main directions. The first, called subjectivism, allows that moral judgments are beliefs which can be evaluated as true or false, but that what makes them true or false has to do

with our own attitudes, desires, reactions, etc. The second, called non-cognitivism, claims that moral judgments express attitudes and are, therefore, not beliefs at all.

## 1.2 Moral Subjectivism

Wiggins traces a progression of subjectivist views from simple to sophisticated. He begins by presenting the simplest formulation of subjectivism as the claim "that 'x is good' or 'x is right'...as uttered by speaker S, says that S approves of x, or that x induces in S a certain sentiment of approbation" (Wiggins, 1987, p. 185). On this view, call it simple subjectivism, moral judgments express beliefs which are true if and only if, the person making the judgment has the relevant attitude. Simple subjectivism creates problems for moral theorizing because it undermines our reasons for engaging in open minded moral debate.

For example, If Sarah says, "Abortion is wrong," and Andrew says, "Abortion is not wrong," then, according to simple subjectivism, Sarah has simply attributed to herself a negative attitude towards abortion, while Andrew has attributed to himself a neutral or positive attitude towards it. However, there seems to be nothing about this situation so understood which generates a disagreement in judgment. There is no logical or factual contradiction generated by the simultaneous belief of the two judgments, because one person's having a given attitude is both logically and empirically independent of another person's having that attitude. All that remains to be identified as a "disagreement" is a brute difference between the two in their attitudes. This difference might lead to a practical conflict when they take actions based on their respective judgments, but it does not seem to be the sort of difference which calls for extended dialogue. Assuming that

neither Sarah nor Andrew is mistaken about their own attitudes towards abortion, then the difference between them does not look as though it could be resolved through debate or inquiry because there are no further facts that could be cited by one as evidence for or against the moral judgment of the other.

Although, if true, simple subjectivism would make it difficult to rationally justify moral theorizing, simple subjectivism is also refutable. Simple subjectivism leaves our actual practices in moral debate utterly mysterious. Although differences in moral judgments sometimes defy our best attempts at dialogue, it is certainly not usually the case that there is nothing more to say than "we have different attitudes towards that." We feel that if our views differ, one of us must be wrong about the matter. In practice we use concrete examples and moral generalizations to try to win over our opponent. There would be no reason to use this method of argumentation if each moral judgment merely attributed an attitude to oneself, because there is no reason to generalize from one attitude to the next.

For example, if I were to hold a judgment that testing cosmetics on animals was okay, but that killing animals for food was wrong, then a friend employing rational argument, could very likely prevail on me to change my views. He could argue that the two actions are morally similar. Once the tension between these two judgments under this generalization was pointed out to me, I would feel a pressure to revise one of the judgments or to articulate some morally relevant difference between the two. This pressure towards revision would be inexplicable if the two moral judgments were reducible to statements about my own attitudes. To see that this is so, consider how fruitless it would be for a friend to try to convince me to start liking green peppers, which

I hate, on the grounds that they do not taste significantly different from red peppers which I enjoy eating. This argument fails because here the standard of significance simply is what I find tasty. In practice, we do not sit around using such arguments to try to convince people to eat vegetables (instead mothers resort to coercive techniques such as withholding dessert). So it looks as though simple subjectivism is a correct account of certain kinds of gustatory judgments and that our actual forms of discourse about food reflect this fact. Conversely, the presence of inexplicable forms of argument in the moral realm indicates that simple subjectivism fails as an account of morality.

Wiggins next explicates a more sophisticated form of subjectivism arising from Hume's 'Of the Standard of Taste.' Here, Hume suggests that when we say that x is good/right/beautiful we mean that "x is the kind of thing to arouse a certain sentiment of approbation," or in terms familiar to us but alien to Hume, "x is such as to make a certain sentiment of approbation appropriate" (Wiggins, 1987, p. 187). This account drops the personal pronoun that attended simple subjectivism's interpretation of moral judgments. Rather than attributing an actual attitude to ourselves, moral judgments make the claim that a certain attitude merits being felt by anyone. This account still views moral judgments as beliefs, but as beliefs about which sentiments are appropriate in response to a situation. This account is immediately more attractive than simple subjectivism because it makes sense of genuine moral disagreement. For example, I might think that right-wing political views merit a reaction of moral contempt while you think they merit a reaction of approval. These two judgments express logically inconsistent beliefs about which one of us must be wrong.



The difficulty for the Humean is to explain, what it could mean for something to merit a response or for a response to be appropriate, without relying on any moral properties to cash out the notion of appropriateness. Wiggins writes that Hume attempted to come to grips with this problem by turning towards the judgments actually made by a competent moral judge under proper conditions. Just as the ultimate standard for deciding something's colour is the way that the object appears to healthy colour sighted people in the daylight, Hume maintains that the ultimate standard of value is the reaction which would be had by a healthy moral judge under conditions which are conducive to good judgment. When people engage in argument about moral judgments, "what they are really seeking to do is to approximate the verdicts of that [ideal] judge" (Wiggins, 1987, p. 190). However, having understood success in moral judgment by comparison to the judgments of an ideal moral judge, the Humean now has the burden of offering us standards for being a good moral judge which do not themselves refer to success in moral judgment. Hume suggests that it is a matter of judging with a cool head, free from bias, and with access to full information. He also makes an analogy with sensory judgments, where illness or a defective sense organ disqualifies a judge.

Let us consider now whether Hume's account (as reconstructed by Wiggins) makes better sense of our practice in moral debate. We sometimes do invoke considerations such as bias, mental illness, or lack of information as evidence against moral judgments. For example, if I am arguing against somebody who judges that our moral duty is to oppose the implementation of the Kyoto Protocol on climate change, and I discover that he is very ill-informed about the scientific case for climate change or that he has friends in the oil and gas industry, I would be right to count these considerations as

evidence against the credibility of his moral judgment in this matter. Hence, in cases such as this one the Humean account seems to give an acceptable explanation of our practice in moral debate.

However, consider Sarah and Andrew's disagreement over abortion. Sarah might argue that abortion is an instance of infanticide, and that it is, therefore, wrong. If Andrew agreed with Sarah's application of the category infanticide to instances of abortion and agreed that infanticide was wrong, then he would feel some pressure to view these instances of abortion as wrong or to provide some special justification why these cases were not wrong despite being acts of infanticide. Yet, the rational force of this sort of argument cannot be understood on the Humean account presented above. Sarah might be happy to grant that Andrew is as unbiased, healthy, and fully informed about abortion as herself. Nevertheless, she still thinks that her infanticide argument holds rational force. Yet, on the Humean view, judgments made by two equally well situated moral judges would be *ipso facto* equally well supported. Since Sarah's argument does not attempt to show that there is some kind of flaw in Andrew's capacity as a moral judge, it does not, on the Humean view, carry epistemic weight. Nevertheless, it seems extremely implausible to say that Sarah's argument is not of the right sort to be taken seriously in moral discussion. To say that would be to allow that Andrew could simply shrug his shoulders in the face of Sarah's challenge, saying, "why should the fact that abortion is infanticide have any bearing on my judgment that it's not wrong?" Faced with this reaction in real life, we would likely think that Andrew was joking (not very funny), morally depraved, or crazy. So although the notion of an ideal moral judge is useful in evaluating moral judgments, it cannot be the whole story.

This brings us to Wiggins' own position. It is still supposed to be a species of subjectivism, but it is one which attempts to avoid the failures of traditional versions of subjectivism to capture essential elements of our moral discourse. Wiggins argues that moral properties seek to capture our felt reactions to the world, but that they are not reducible to statements about these felt reactions. The truth of moral judgments, therefore, depends on facts about our felt reactions, but we can only evaluate the truth of a moral judgment by asking whether the judgment's moral predication is correct.

Wiggins gives a fascinating account of the way in which irreducible value properties might emerge from our felt responses to the world. He envisages a process of cultural evolution in which we begin with the sort of property/reaction pairs envisaged by the simple subjectivist, but eventually complicate these to the point where there is no simple correspondence between property and easily identifiable reaction. For example, a property such as nausea is claimed to be at a "semi-valuational" stage (Wiggins, 1987, p. 198). With nausea, something like simple subjectivism is clearly correct. To say that green peppers are nauseating is just to say that green peppers provoke in me a reaction of nausea. However, Wiggins is explicit in claiming that complex valuational terms such as good/right/beautiful have progressed beyond this stage. For "when the system of properties and reactions diversifies, complicates and enriches itself – there will often be no saying exactly *what* reaction a thing with the associated property will provoke without direct or indirect allusion to the property itself." (*ibid.*, p. 195). At some point, the relationship between the predicate used to capture certain reactions, and the identification of the reactions captured by the predicate will become reciprocal. For example, our claim that 'today's property distribution is unjust' is true if and only if we react to today's

property distribution with a certain mix of indignation, disapproval, etc., but exactly what this mix of feelings is can only be put into words by using the term injustice itself.

Therefore, in order to decide whether something provokes feelings of injustice in us, rather than, say, disgust, we will have to rely, in part, on our unanalyzed judgments that the thing in question is unjust.

Wiggins admits that this account is in some sense circular because we know that the reaction of "rightness" is felt in response to X, if and only if we know that the predicate 'right' is the correct description of our reaction to X. However, the predication 'X is right' is true if and only if feelings of "rightness" are felt in response to X. However, Wiggins' argues that this circularity is acceptable because a subjectivist account should not be expected to provide a reductive analysis of valuational language into non-valuational language (Wiggins, 1987, p. 188). Although, historically, many moral subjectivists have presented their theories as reductive analyses, this is supposed to be non-essential to the subjectivist project. As Wiggins writes, "[t]he point of calling this position subjectivism is that the properties in question are explained by reference to the reactions of human subjects" (*ibid.*, p. 195). Another way that Wiggins might have tried to explain this apparent circularity in his view is to say that metaphysically all we have are humans trying to make sense of their reactions; no additional properties needed. However, at an epistemic level we must recognise that our actual judgments about moral properties are indispensable for identifying the reactions these terms are meant to pick out. Wiggins' view trades on the idea that although value properties arise from human attitudes, we can only discover these attitudes and their relation to value properties by

looking to these properties in the world as they apply to various actions and states of affairs.

By forcing us to look for subjective properties out in the world, Wiggins' view begins to draw quite close to the observational moral realism presented above. Both views claim that our judgments about moral properties are the only way to access moral truths. The difference is that Wiggins' subjectivism gives a metaphysical reduction of moral properties to felt reactions, while the moral realist claims that moral properties are metaphysically irreducible.

### **1.3 Quasi-Realist Non-Cognitivism**

Simon Blackburn has argued against Wiggins' notion of irreducible response dependent moral properties in favour of his own non-cognitivist position. Blackburn challenges Wiggins' notion that the truth/acceptability of moral judgments has anything to do with actual or appropriate felt responses. For example, Blackburn argues that what is wrong with cruelty is not that it provokes in us certain responses that are correctly labelled with the predicate moral wrongness. Rather "[i]t is because of our responses that we *say* that cruelty is wrong, but it is not because of them that it is so" (Blackburn, 1993, p. 172, his emphasis). Cruelty is wrong for a whole host of *moral* reasons, such as the harm caused and disrespect displayed by cruel acts, but never do we say that cruelty is wrong because cruel acts merit attitudes of moral disapproval.

It may look as though the forgoing objection by Blackburn favours the realist alternative to Wiggins' subjectivism in which moral properties are taken to be metaphysically irreducible, but Blackburn is an avowed anti-realist. This tension is

resolved by an understanding of Blackburn's unique meta-ethical position, quasi-realism, which combines an expressivist non-cognitivism, with a commitment to such traditionally realist notions as moral truth, moral knowledge, and moral objectivity. Expressivist non-cognitivism holds that moral judgments do not express beliefs, not even beliefs about attitudes, but rather that they express attitudes directly. The enigma of how such a combination of views could be consistent is unlocked by Blackburn's use of a Carnapian distinction between internal and external questions about morality. When answering internal moral questions, such as what it is that makes cruelty wrong, we do not refer to our attitudes or appropriate reactions, but to morally salient features of cruelty. Of course, in so doing we are expressing attitudes, but this fact is irrelevant to answering the moral question. Blackburn rejects entirely the request to give some non-moral metaphysical account to answer moral questions.

Blackburn does, however, think that we can answer external questions about ethics, and puts forward an expressivist non-cognitivism to do just that. However, despite Blackburn's insistence on the separation between internal and external questions, we might still worry that Blackburn's answer to the external question of what moral judgments are could have some implications for how we conduct ourselves in moral debate. The identification of moral judgments with expressions of attitude might make us worry that Blackburn's account is in danger of having the sorts of implications about moral disagreement and moral debate that we found difficult to swallow when considering simple subjectivism. However, this danger can be avoided by Blackburn's rigid adherence to the external/internal distinction. He argues that for the same reason that we cannot explain why cruelty is wrong in terms of our attitudes, neither can we use

our story about the expressivist semantics of moral language to derive moral conclusions. In moral debate, we are using moral language. In use, Blackburn claims, moral language functions in precisely the same way as non-moral language. It is only when we step outside of moral debate to ask meta-ethical questions that we are able to see that moral language is expressive of attitude. Thus, Blackburn neatly reconciles the weighty legacy of Hume's location of value in human sentiment with the Wittgensteinian commitment to respect the functioning of language in its proper context.

Blackburn's quasi-realism is tantalizing because it seems to capture the best of the realist position while remaining true to the motivations behind anti-realism. However, quasi-realism's apparent immunity from objections may be due to its explanatory poverty. Older versions of non-cognitivism were evaluated (and usually rejected) according to their ability to offer compelling explanations of our use of moral language. Blackburn's view circumvents this form of evaluation by refusing to give an actual analysis of moral sentences into sentences which are transparently attitude expressive. While this refusal makes the view harder to criticize, it also counts against the view's explanatory success because the claim that moral judgments are really expressions of attitude is uncashable. Therefore, although quasi-realism is a strong contender for the correct meta-ethical account, it does not knock observational realism and Wiggins' subjectivism out of the running. The latter accounts attempt to give more substantive explanations of our moral language. Thus, even if they are somewhat more vulnerable to criticism than quasi-realism, we should not be quick to reject them.

## Chapter 2: Case Evidentialism

We have arrived at three meta-ethical positions which have been deemed plausible candidates for the truth about the nature of morality: observational realism, Wiggins' subjectivism, and quasi-realism. It is possible that further debate in the field of meta-ethics will reveal that one of these views is clearly better than all others. However, at present, each view appears to have both strengths and weaknesses and we have seen no decisive arguments in favour of one over the others. Therefore, it would be best for our moral epistemology to be neutral, if possible, between the three competing meta-ethical views. In this chapter, I will articulate an epistemological account which I call Case Evidentialism and argue that this account is supported by all three of the meta-ethical views we have been considering. I capitalize the term 'Case Evidentialism' so that it will not be confused with other doctrines.<sup>2</sup>

Case Evidentialism has four central claims: (1) When anyone makes a judgment about the moral status of a particular action, the content of that judgment is supported by non-inferential evidence, (2) general moral statements are never supported by non-inferential evidence, (3) moral principles and other general moral statements may be evidentially supported by inductive and abductive inference from the contents of our moral judgments about particular actions, and (4) evidentially supported general moral statements can support or undermine particular moral claims through inference.

All four of the central claims of Case Evidentialism invoke the concept of evidence. Evidence is defined as follows: A claim is supported by evidence E if and only

---

<sup>2</sup> In developing Case Evidentialism I have been greatly influenced by Susan Haack's work on evidence in empirical inquiry (Haack 1993; 2003). I will not try to say exactly how I have modified her views in adapting them to give an evidentialist account of moral inquiry.



if E gives us reason to think that the claim is more likely than not to be correct. In other words, evidence gives us a defeasible reason to accept the claim supported by the evidence.

## **2.1 First Thesis: Non-inferential evidence for particular moral judgments**

Recall that the first claim of Case Evidentialism states, when anyone makes a judgment about the moral status of a particular action (hereafter particular moral judgment), the content of that judgment is supported by non-inferential evidence. Non-inferential evidence is a degree of likelihood/epistemic weight that gets assigned to a sentence because the sentence is produced through our engagement with the world or with our own mind. For example, the empirical observations such as 'the cup is on the table' is non-inferentially supported by the fact that somebody makes this observation of the world. Similarly, the introspective claim, 'I feel sad' is non-inferentially supported by my introspective judgment that the predicate sadness best expresses what I am currently feeling. Therefore, both subjective claims such as 'I feel sad' and objective claims such as 'the cup is on the table,' are supported by non-inferential evidence. This non-inferential evidence consists in our *act* of introspection and observation respectively, while the claim supported by this evidence is the sentence which is produced by the act of introspection or observation.

Introspection and observation are both actions which consist in our attempting to put a non-linguistic experience, such as a feeling or a perception, into a sentence. We hold that such acts give non-inferential evidence to the sentences they produce because we assume that every person has some ability to accurately interpret their own

experience, unless we have evidence of incompetence. Without this presumption in favour of our ability to describe our own experience, we would be led to complete scepticism about empirical and introspective knowledge.

Case Evidentialism claims that the contents of moral judgements about particular actions are non-inferentially supported by the act of moral judgment. The particular action which we consider may be encountered as an actual dateable event. For example, we might consider a moral choice that we or somebody we know has actually faced. Alternatively, we might imagine a moral choice, which to our knowledge has never been faced by a real person, but which is faced by a fictional person. Such fictional cases may be drawn from literature or they may be invented by philosophers as thought experiments. When considering fictional cases, we imagine the moral choice as one which is taken by a unique moral agent facing concrete circumstances, so judgments about fictional cases are still particular judgments.<sup>3</sup>

For clarity and ease of use, I will call the act of making a moral judgment, a moral judging and call the sentence produced by the judging a judged content. Each moral judging produces a judged content, and the two together constitute a single moral judgment. To illustrate this terminology, consider Sharon who reads about a police raid at a gay bathhouse. Upon reading the article, she is prompted to assert that the activities in the bathhouse were morally wrong – a moral judging has occurred. The judged content of this moral judging is whatever is expressed by the sentence, 'the activities in the bath

---

<sup>3</sup> The semantic question of how fictional actions are to be understood is beyond the scope of this work. However, on one plausible understanding, fictional agents and actions are logically particular. The character Hamlet is not a definite description as is the present king of France. Rather, Hamlet is a fictional person who performs various actions and is faced with various choices. The present king of France may be analysed as a quantified variable ranging over logical objects, whereas Hamlet must be analysed as a logical object.

house at the time of the police raid were morally wrong.' Case Evidentialism holds that Sharon's judging gives her some evidence in favour of this moral sentence.

Defined as such, the judged content does not depend on the meta-ethical position held by the moral judge. A meta-ethical position will determine what one *takes* moral sentences to be expressing and, therefore, what one *takes* the judged content to be.

However, the actual judged contents do not depend on what meta-ethical position we presently favour, but on what meta-ethical position is actually true. So, for example, even if a moral realist intends for his moral judging to yield a belief about moral properties, the judged content of his judging would not be such a belief if realism is false. If, for example, quasi-realism is true, then the realist's judged content would be the endorsement of an attitude. The first claim of Case Evidentialism, that moral judgments give non-inferential evidence to judged contents, can be established by showing that moral judging is an action like observation and introspection – an action which produces a claim that has initial plausibility given the manner in which it was produced. To establish this, we must examine each of our three meta-ethical accounts in turn.

Observational realism provides the most straightforward basis for the first thesis. On this account, moral judgments are observations. Therefore, we can understand moral judgments on analogy with observations in science. Just as the laboratory scientist's acts of observation give us non-inferential evidence for data-points and observation statements, so too, our moral judgments give us non-inferential evidence for judged contents. On the realist view, a moral judging is an attempt to put into words the experiences we have of objective moral properties. Given the anti-sceptical assumption that we must place some trust in our interpretations of experience, we arrive at the

conclusion that judged contents are evidentially supported by moral judgments. How much evidential weight a moral judging gives to a judged content will depend on how well informed the moral judge is about the case in question, but any moral judging gives at least some weight to the judged content which it produces. This establishes the first thesis of Case Evidentialism on the basis of observational realism.

The first thesis of Case Evidentialism can also be established from the two forms of anti-realism we are considering. Wiggins' subjectivism claims that we experience felt reactions in response to actions and events in the world and that through the correct application of moral predicates to these felt reactions, we come to know moral truths. Applying a moral predicate to reactions felt in response to a concrete situation is simply the subjectivist analysis of a particular moral judging. The application of a subjective predicate is an attempt to interpret one's own feelings through the lens of culturally meaningful terminology. As such, Wiggins' understanding of particular moral judgments displays the same evidential structure as the realist understanding. Moral judgments, understood as applications of subjective predicates, provide non-inferential evidence for judged contents understood as culturally meaningful valuations. Therefore, although the concepts used to analyze particular moral judgments and moral principles differ between subjectivism and observational realism, both support the first thesis of Case Evidentialism.

It is a bit trickier to assess whether Blackburn's quasi-realism allows us to support the claim that particular moral judgments provide non-inferential evidence for judged contents. Blackburn states, "our conception of moral truth is entirely exhausted by our understanding of the equivalence that  $p$  is true if and only [if]  $p$ . To allow that a moral

judgement is true is to endorse the attitude it expresses" (Blackburn, 2002, p. 134).

Elsewhere he clarifies that to say that a moral statement is true is just another way of endorsing that statement (Blackburn, 1999). Yet, if endorsing a moral view is all that it takes for the view to be true, then a moral judging entails the "truth" of judged contents. For Blackburn, a moral judging is an endorsement of an attitude expressed by a sentence M. If saying 'M is true' is just another way of saying M, then to endorse M through a moral judging entails the truth of M. For example, if Drew judges, 'Nancy is morally obligated to give money to Oxfam,' then Drew endorses the claim that Nancy is morally obligated to give money to Oxfam, and, therefore, *it is true that* Nancy is morally obligated to give money to Oxfam. This analysis would make us all into infallible judges of the moral truth. Simply by endorsing a moral sentence we guarantee that the moral claim which that sentence expresses is true. Clearly, there is something wrong with Blackburn's analysis of moral truth, even given his own expressivist account of moral judgment. Blackburn does not want to accept the conclusion that our moral judgments are infallible, he thinks they can be questioned and revised in light of moral considerations. The problem for Blackburn, is to find a way to maintain the quasi-realist commitment to expressivism while allowing a fallibilist account of moral judgment.

The solution to this problem is already implicit in Blackburn's idea that moral claims are revised in light of other moral claims and in light of better information about the case at hand. The truth condition for a moral sentence M, cannot be merely that we endorse M here and now. Rather, the truth condition for M is that we *would endorse* M given perfect knowledge of non-moral facts and in light of all our other attitudes.

Obviously, we do not have perfect knowledge of non-moral facts nor do we have perfect

knowledge of all our attitudes and their relations. Therefore, our judging that M does not entail the truth of M. Infallibility is avoided.

Although we can never claim 'M is true' infallibly, because of our inevitably imperfect knowledge of the non-moral facts and our own attitudes, we can claim that M is more or less likely to be true. Since we do have a limited knowledge of non-moral facts and of our own attitudes, which knowledge informs our moral judgments, it is sensible to claim that judged contents are more likely to be true than particular moral statements which are not judged contents. Therefore, the quasi-realist picture does support Case Evidentialism's first thesis that our moral judgments give us evidence for judged contents. Although quasi-realism differs considerably from both observational realism and Wiggins' subjectivism, it agrees with both that particular moral judgments give non-inferential evidence to judged contents. The first thesis of Case Evidentialism is supported by all three meta-ethical accounts.

## **2.2 Second Thesis: General moral claims not supported non-inferentially**

The second claim of Case Evidentialism is that general moral statements are never supported by non-inferential evidence. Once again, I will argue that all three meta-ethical views support the claim. Examples of general moral statements include traditional moral principles such as 'murder is always wrong' or 'the action which produces the greatest happiness or the greatest number is morally obligatory.' General moral statements may also be more specific moral claims such as 'it is never morally permissible to kill a baby' or 'wealthy Canadians ought to give some of their money to overseas poverty relief.'

In the case of observational realism, particular judged contents are supported by their status as observations, where observations are understood as the putting into words of our non-linguistic experience of moral reality. Crucial to this account is the fact that the consideration of concrete moral cases (whether real or imagined) provides us with moral experience. In contrast, the general claim 'murder is always wrong' is not made in response to experience, or it is obscure what that experience could be. When we imagine a man murdering somebody, we have an experience which prompts our particular judgment that the man's action is wrong, but this experience does not directly give us the general claim that murder is always wrong. In considering whether it is wrong for a doctor to help a suffering patient commit suicide, we might invoke the principle that murder is always wrong to help form a judgment. However, in this case, the general claim about murder is an input to the judgment rather than an output. It seems that moral experience never supports a general moral claim directly. One might say that it is possible to *just see* that murder is wrong, but what we really mean here is that faced with acts of murder, we observe that they are wrong without inference from a general principle. We never *see* that a general moral statement is correct. The concepts of perception and observation are stretched past their breaking point when we start talking about seeing that a general statement is true.

For similar reasons, Wiggins' subjectivism does not allow non-inferential evidence to be assigned to general moral claims. General moral statements are correctness conditions for our moral predicate applications; they are not, themselves, applications of the predicates. When we claim 'murder is always wrong' we are claiming that whenever faced with an act of murder, a well situated moral judge has certain

feelings which are best captured by the predicate wrongness. The truth of this correctness condition depends entirely on the appropriateness of predicate applications in particular cases. Therefore, while we might infer the truth of this statement from predicate applications over a wide range of cases (more on this in the next section), we cannot claim that this statement gains any degree of plausibility directly from our experience of felt reactions.

An objector might say, "surely we can simply have a strong feeling towards murder in general. This feeling might be best described as moral wrongness." The objector is trying to claim that we somehow simultaneously feel that all acts of murder are wrong, without considering each act in particular. However, we must ask what it would mean to have feelings towards murder in general. It might mean that all particular actions which are properly described as murder provoke in us feelings properly described with the predicate 'wrong.' This means that the feelings are always directed towards a particular act of murder, not towards the general category. On the other hand, speaking of a feeling towards murder in general could be meant to indicate a feeling we have towards the category of murder. However, we cannot have feelings of "wrongness" towards the category of murder, because wrongness is a property of actions, while the category of murder is not an action but an abstract type. At most, we could feel that the category of murder was morally bad or that the *use* of this category was morally wrong. Such feelings do not give us reason to think murderous acts are wrong; rather, they give us reason to avoid using the category in our thought and discourse. Another example makes this point vivid. We might judge that describing an action as harlotry is morally offensive, and judge, 'Harold's use of harlotry as an action-type is morally wrong.' Here we are certainly



not claiming that all acts of harlotry are morally wrong, nor could this be inferred from our claim. The fact that we can have feelings about a category of action cannot be used to establish the moral status of actions which fall under that category. Therefore, Wiggins' subjectivism implies the second thesis of Case Evidentialism, that general moral claims are never supported by non-inferential evidence.

On the quasi-realist account, all that it takes for a moral statement to be true is for it to express an attitude which we endorse at the limit of perfect knowledge and self-transparency. If we endorse any moral claim, this gives us some reason to think that we would endorse it in this limiting condition. Therefore, quasi-realism may seem to support a homogeneous account of moral evidence in which both particular judgments and general claims are said to be, in part, non-inferentially supported. For example, if Joan judges that meat-eating is always wrong, she could be said to be endorsing an attitude towards all actions which involve eating meat. By virtue of this endorsement, she has some reason to think that she will continue to endorse the attitude on further reflection. This seems to allow, contrary to the second thesis of Case Evidentialism, that her general moral view is evidentially supported without inference from particular moral judgments.

The claim that the statement, 'meat-eating is wrong' is non-inferentially supported would be intelligible if the statement could be reductively analyzed as the expression of a single attitude towards all acts of meat-eating such as 'yuck, all acts of meat-eating!' However, quasi-realism rejects the use of expressivist analyses for the purposes of discourse about internal moral questions. Therefore, in moral theorizing, quasi-realism must work with the propositional understanding of 'meat-eating is wrong' as, for any action, if it is an act of meat-eating, it is wrong. On this understanding, the truth of the

statement depends entirely on truths about the wrongness of particular actions. Therefore, in order for Joan to rationally accept the general claim that meat-eating is wrong, she must make judgements about particular cases of meat-eating from which the general truth can be inferred. If quasi-realism is truly to take the propositional structure of moral language seriously, it must admit that a general moral statement does not express a single attitude. Rather, a general moral statement is understood to endorse multiple attitudes, each of which is directed towards a particular action.

At the level of particular judgments, such as 'Cindy's eating roast chicken today was wrong,' the propositional understanding of moral sentences provides us with no further analysis of their truth conditions. Only at this point can the quasi-realist legitimately step in with his metaphysical picture and say that we are endorsing an attitude. Therefore, despite appearances, the quasi-realist is committed to restricting non-inferential evidence to particular judged contents, resulting in a heterogeneous epistemology. This commitment brings quasi-realism back into harmony with Case Evidentialism.

### **2.3 Third Thesis: Moral principles supported through inference**

As already hinted in the previous section, general moral statements may be inferentially supported by judged contents. Moral principles are held to be a special kind of general moral statement. Although we will not develop a rigorous account of the distinction between moral principles and other general statements until Chapter 4, we can imperfectly indicate the difference by saying that moral principles tend to be highly general claims covering a wide range of different actions, whereas other general moral

statements might be highly specific claims which cover only a small range of cases. For example, 'it is always wrong to murder' is usually considered to be a moral principle whereas, 'it is always wrong for Allison to murder her nasty aunt' is a highly specific general moral statement which is not a good candidate for a moral principle. For now, I will focus on the evidential support for general moral statements, of which moral principles are one kind, and will return in Chapter 4 to discuss the establishment of moral principles over other general moral truths.

We established above that judged contents are supported by non-inferential evidence provided by moral judgments. We can now explain how judged contents can legitimately be used to evaluate general moral claims. Another kind of evidence which is well understood from empirical inquiry is the evidence provided by non-deductive inferential links between statements. One statement counts as evidence for a second statement when the first statement is evidentially supported itself and there is a cogent inference from the first statement to the second statement. The likelihood of the first proposition impacts our assessment of the likelihood of the second. Since judged contents are evidentially supported themselves, inferences from them provide evidential support to general statements. In particular, inductive inference and inference to the best explanation (abduction) from our judged contents provides evidence for the universal statements which generalize those contents. Therefore, I will call general moral claims which are evidentially supported by inference moral generalizations.

For example, my judged contents about many different cases of rape, that all are wrong, give me strong inferential evidence for the generalization 'raping is always wrong.' If some philosopher thinks up a bizarre case in which raping does not seem to be

wrong, such as a case in which raping somebody saves the lives of many other people or prevents many other rapes, this case is inferential evidence against the generalization that rape is always wrong. The judged content 'this unusual case of rape is not wrong' is supported by non-inferential evidence from its moral judging, and some of this evidential support is transferred by inference to reduce the total evidence in favour of the generalization 'raping is always wrong.'

I argued in Section 2.2 that all three meta-ethical accounts treat general moral claims as universally quantified statements ranging over particular actions. For example, the general claim, 'murder is always wrong' is the proposition which states, 'for any action, if it is murder, then it is morally wrong.' Of course, the three meta-ethical positions give different accounts of what it means to say that an act of murder is morally wrong, that is they differ in their account of particular moral claims. Yet, they all agree that general moral truths are universally quantified statements whose truth or falsity depends on the truth of particular moral claims. Therefore, all three views allow that general claims may be inferred via induction and abduction from evidentially supported particular moral claims. Since the three views are also in agreement that judged contents are particular moral claims which are non-inferentially supported by the moral judgments which produce them, the three views support exactly the same account of what may count as evidence for a moral generalization.

One possible confusion about the claim that moral generalizations are inductively and abductively supported by judged contents is to think that such generalizations would be purely predictive/descriptive generalizations, rather than normative generalizations. This confusion arises when we speak casually of moral judgments providing evidence for

moral principles, thereby, failing to distinguish between moral judgments and judged contents. The claim that particular moral judgments provide evidence in favour of moral generalizations is ambiguous between two readings. First, we could be claiming that inference from the *occurrence* of a moral judgment provides evidence for a generalization which predicts/explains that occurrence. Second, we could be claiming that inferences from judged contents provide evidence for generalizations about those contents. Both of these readings are true, but the generalizations described by the first interpretation are descriptive generalizations, while those of the second are normative generalizations.

A theorist who was interested in the first kind of generalization would try to predict the occurrence of future moral judgments by seeking patterns in known moral judgments. For example, from what Lisa knows of Miriam's past moral judgments about eating meat, Lisa could form the descriptive generalization that Miriam is disposed to judge that meat-eating actions are morally acceptable. Such a descriptive generalization could be carried out for a whole group rather than a single individual, giving us a social scientific theory of the moral norms of that group. We cannot infer moral claims from such a social scientific theory.

As should be clear from the thesis defended in this section, Case Evidentialism is interested in the second kind of generalization. Moral generalizations do not range over occurrences of moral judgments, they range over the moral statements for which we have evidence, the most important of which are judged contents. For example, in considering a cop show on TV, I judge that the drive-by shooting in the show was wrong. I also judge that the drive-by shooting is an instance of murder. These judgments give non-inferential evidence for the judged content that the drive-by shooting was wrong and the belief that

it is an instance of murder. In turn, the judged content that the drive-by shooting was wrong, together with the belief that the shooting was a murder, provide inferential evidence for the moral generalization 'all murders are wrong.'

A descriptive theory of moral norms attempts to develop generalizations based on the fact that we made the judgment that a drive-by shooting was wrong. In contrast, a normative moral theory develops generalizations based on the likelihood that the drive-by shooting was wrong. In general, it is not my *judging* 'action X is wrong and X has features F' which inductively supports the moral generalization 'all F actions are wrong.' Rather it is *that* action X is likely wrong and likely has features F which supports the generalization – this is my judged content. My *judging* 'X is wrong and X has F' is non-inferential evidence for the judged content *that* X is wrong and X has F. So the fear that Case Evidentialism led to descriptive rather than normative generalizations was based on a confusion between the moral evidence provided by inferences from our judged contents and the non-moral evidence provided by inferences from the occurrence of moral judgments.<sup>4</sup>

#### **2.4 Fourth Thesis: Evidence from generalizations to particular claims**

The forth central thesis of Case Evidentialism is the claim that a well supported moral generalization can support or undermine particular moral claims through inference.

---

<sup>4</sup> Such a confusion is at the root of Gilbert Harman's argument (1977) that moral judgments cannot count as evidence for moral theories. He argues that a psychological theory of moral judgment offers a perfectly good explanation for why we make the moral judgments we do. Therefore, Harman concludes, moral theory is explanatorily superfluous. However, while psychology may be able to explain why people *judge* that X is wrong, in so doing it does not explain why X *is* wrong. As empirical statements, psychological generalizations are inferentially supported by moral judgments rather than judged contents. So even a perfectly predictive/explanatory psychological theory leaves our judged contents unexplained. Harman might want to claim that successful psychological explanations of moral judgments make us lose interest in explaining the judged contents. However, for the bulk of people who remain interested in evaluating moral claims, moral theory has not been shown to be explanatorily superfluous.

A moral generalization has the logical form 'for all actions X, if X has features F, then X is wrong.'<sup>5</sup> When such a generalization is well supported by judged contents, it provides inferential evidence about all the particular moral statements with feature F because there is a cogent form of inference from the likely truth of a universal statement to the likely truth of all instantiations of that universal statement.

For example, if we attribute likelihood to the claim that all instances of drunk driving are wrong, then we should increase the likelihood we attribute to claims that particular instances of drunk driving are wrong and decrease the likelihood we attribute to claims that instances of drunk driving are permissible. As this example illustrates, when we come across a judged content that contradicts a well supported moral generalization, then the moral generalization gives us inferential evidence against the contradictory judgement. In cases where the evidence given by inference from a moral generalization against a judged content is stronger than the judged content's non-inferential evidence, then we must revise our judged content to bring it in line with the moral generalizations.

The ability to rule out all but the most anomalous judged contents using well supported moral generalizations depends on the use of inference to the best explanation rather than simple induction alone. Explanation is needed to rule out theories which are forced constantly to make *ad hoc* modifications to their moral generalizations in light of new cases. Our best explanations should be as parsimonious and predictive as possible given the evidence. Harman argues that we should not expect a formal calculus of explanatory inference, such as that which can be provided for deductive inference and

---

<sup>5</sup> At this point we are considering only the most basic kind of moral generalizations which are about our overall moral duty with respect to a certain type of action. In Chapters 4 and 5, we discover other kinds of moral generalizations which may be needed to capture the full complexity of our moral thought.

probabilistic inference (Harman, 1999). Nevertheless, the informal norms of explanation provide rational guidance to theorists of all kinds in evaluating evidence.

To illustrate the point about explanation, consider the moral generalization 'it is always wrong to murder' and assume that this generalization is well supported by many of our judged contents. Across a wide array of cases in which an agent has the opportunity of committing murder, we judge that it would be wrong for the agent to do so. However, we come across a case in which the agent in question must kill an innocent person if she is to save the lives of many other innocent people, and we judge that in this case, the murderous action is not wrong. Faced with this new evidence, we wonder whether we should revise our original generalization that murder is always wrong or if we should hold that the contradictory judged content should be overruled by the generalization.

If we were working with a strictly formal procedure of inductive inference, there would be nothing wrong with modifying the original generalization to read, murder is always wrong, except in cases where one must murder to save many innocent lives. Since none of the other cases we have considered previously have involved murdering to save many innocent lives, a strictly formal inductive procedure can find no evidence against this modified principle. However, within the informal norms of explanatory inference, such a modification is unacceptable, because it introduces an *ad hoc* change to a well supported generalization. Roughly speaking, a modification is *ad hoc* if it is a change made in light of contradictory evidence, where the change accommodates that evidence through a highly specific stipulation that accounts for that evidence and little else. Given the informal norms of explanatory inference, a moral theorist considering the potential



counterexample to the generalization that murder is always wrong must ask whether this counterexample can be accommodated by an alternative and explanatorily superior moral generalization.

A utilitarian theorist, for example, might argue that the principle that one should always maximize human happiness explains all the cases on which we based our original generalization that murder is always wrong as well as other cases which have nothing to do with murder, while also accounting for the anomalous judgment about the case of killing to save innocent lives. Explaining the acceptability of the anomalous case of murder in terms of the generalization 'always maximize human happiness' is not *ad hoc*, as was the generalization 'never murder except in this case,' because the utilitarian principle has evidential support which is independent of this particular anomalous case. However, the utilitarian principle runs into problems of its own, and we must evaluate whether it is overall more strongly supported than the generalization that it is always wrong to murder. If not, then our original generalization that it is always wrong to murder would overrule the judgment about murdering to save innocent lives.

Having fleshed out the way in which moral generalizations provide inferential evidence for and against particular moral claims, I should make explicit how each of the meta-ethical accounts I have considered would make sense of this use of theory to improve our moral views about particular situations. We must review the way in which each of the three accounts allows for errors in our moral judgements and why these errors may be corrected by considering well supported moral generalizations.

Taking observational realism as our starting point, it is once again fruitful to draw on the analogy with empirical observation. Empirical observations have been shown to be

conceptually laden (Hanson, 1958). This means that our internalized theories about the world influence (but do not determine) the empirical observations that ordinary people and scientists make. The same would be true on a realist account of moral observation. The theory dependence of observation can be a source of both error and accuracy. If our observations are influenced by an inadequate or false theory, we are prone to observational error. On the other hand, by internalizing well supported scientific theories, scientists are able to make precise observations using complex instruments. Therefore, the claim that observations are theory dependent shows us how observational realism can both account for error in our particular moral judgments and account for the usefulness of a well developed moral theory in enabling accurate judgments. Also, just as scientists are sometimes able to infer that an observation is due to experimental error by recognizing that it conflicts with the predictions of a well supported scientific theory, so too the observational realist would hold that non-inferential evidence for particular judged contents may be overruled by inference from a well supported generalization.

Wiggins' subjectivism allows that particular moral judgments can be mistaken because the standards for the correct application of moral predicates are not perfectly mastered by any individual. These predicates arose organically through a process of cultural evolution, and it may, therefore, be "a matter of keen argument what it takes for a thing to count as having the property that the attitude is paired with" (Wiggins, 1987, p. 199). So, for example, for an action to be right it must be such as to provoke the characteristic feeling of "rightness." However, not only can we not determine whether the action in fact provokes "rightness" in us without judging that the action is right, but in order to judge that it is right we might need to "argue about what the marks are of the

property" rightness (*ibid.*, p. 198). Our judging that an action is right gives us some evidence to think that it really is right since we assume that we have a partial grasp of the proper use of the moral predicate rightness. However, since our grasp of the predicate is only partial and since no individual is an authority on the correct use of the predicate, our particular moral judgments can be overturned by consideration of well supported correctness conditions for predicate use. If Wiggins' subjectivism is true, then well supported moral generalizations are such correctness conditions and can, therefore, be used to guide our evaluation of particular cases.

Finally, the quasi-realist can make sense of the idea that our particular moral judgments could be mistaken by noting that the attitudes we are willing to endorse change in response to moral reflection. They do so because we come to better understand the non-moral features of cases and we discover new inferential relationships between the statements which express our attitudes. Further, we may encounter new cases which prompt expressions of attitudes we did not previously recognize. When we change an attitude in light of these new considerations, it is natural to say that the discarded attitude was false, in the same attitudinal sense of truth and falsity which allows us to say that the attitudes we would ultimately endorse are true.

Blackburn (1999) writes that the reason for the propositional structure of moral judgments is "to bring to the business of systematizing and reasoning about attitude the elegant framework of propositional logic" (pp. 213-214). If the propositional form of moral judgments brings with it the norms of propositional logic, then there is no reason to suppose that it will not also bring the norms of inductive logic and the informal norms of explanatory inference. Since Blackburn is committed to accepting the propositional form

of moral judgments for the purpose of answering normative moral questions, then any inferential links which exist between particular moral judgments and general moral statements must be taken seriously by a quasi-realist. Therefore, quasi-realism must hold that inference from supported moral generalizations may lead us towards improving our moral views about particular cases.

## **2.5 Grounding Reflective Equilibrium in Case Evidentialism**

With the exposition and defence of the central theses of Case Evidentialism complete, we are now in a position to see how Case Evidentialism supports the methodological guidelines of Reflective Equilibrium. Reflective Equilibrium claims that moral theorizing involves revision of both our particular moral judgments and our general moral principles in light of one another. In other words, moral theory is neither an assemblage of disconnected "intuitions" about particular cases, nor is it a dogmatic application of moral principles to particular cases. Reflective Equilibrium holds that both principles and judgments have an indispensable role to play in moral inquiry but that neither are dominant or unrevisable. Sometimes, we change a moral principle in light of our judgments about moral situations, other times, we change our judgment about a case as a result of our holding a general principle.

The evidential relationships between judged contents and moral generalizations explain both of the directions of revision. According to Case Evidentialism, a moral principle is a type of general moral claim. We have seen that inference from judged contents provides us with evidence for and against general moral claims. Therefore, a pre-existing commitment to a moral principle may be revised in light of the inferential

evidence of judged contents. The other direction of revision can be understood in terms of inference from well supported moral generalizations to particular moral claims. If the evidence given by inference from a well supported moral generalization against a judged content is stronger than the evidence given to that judged content by its moral judging, then we should revise the judged content.

In the case of judged contents that are persistently produced by moral judgments in the face of theoretical counter-evidence, we may not be able to revise the judged content in the sense of actually judging otherwise, but we can hold that this persistent moral judgment is misguided and attempt to avoid relying on it in moral decision making. For example, we might judge that an unfamiliar sexual practice, such as rubber fetishism, is morally wrong. On examining this judgment in light of our best supported moral generalizations, we decided that the judgment has too much indirect evidence against it to be sustained in cases where rubber fetishism is consensual and not harmful. Nevertheless, we may be unable to shake the sense we have that rubber fetishism is wrong and, as a result, we continue to have moral judgments attributing wrongness when confronted with cases of rubber fetishism. Given our evidence against the contents of these judgments we should hold that they are misguided and attempt not to act on them. However, if a change in our theoretical understanding were to cast the morality of rubber fetishism in doubt, these repudiated judgments would need to be taken seriously as a source of non-inferential evidence against the practice.

Case Evidentialism gives us more guidance in evaluating our moral views than does Reflective Equilibrium alone. The particular judgments which play a role in supporting our strongest set of moral generalizations will also be those which are least in

danger of revision. In contrast, outlying judgments which do not fit into any well supported generalization are more likely to be revised. Such an account has the advantage of offering robust methodological guidance to moral theorists. In contrast, Reflective Equilibrium gives us only the rather vague directive to allow moral principles to *sometimes* trump particular judgments and vice versa.

In this chapter I began with three plausible meta-ethical views, selected to serve as representatives of realism, subjectivism, and non-cognitivism respectively. I then argued that the four theses of Case Evidentialism provide an epistemological account which can be found implicitly in the three meta-ethical positions despite their different understandings of the nature of morality. I then showed how Case Evidentialism allows us to make epistemological sense of Reflective Equilibrium. In the remaining chapters, I defend Case Evidentialism against philosophical challenges and continue to develop it to take account of new complexities as they arise.

### Chapter 3: Hare's Attack on Moral Judgment

R. M. Hare was one of the most serious and persistent critics of the use of moral judgments in the evaluation of moral theories. Hare presented arguments directly against Reflective Equilibrium shortly following the publication of *A Theory of Justice*, and continued to elaborate these critiques in later work. In this chapter, I explore and respond to several arguments against the use of moral judgments which arise from Hare's work.

To avoid confusion in discussing Hare's arguments, we must first address a matter of textual interpretation. When Hare speaks of moral intuitions, he has in mind roughly what I have been calling moral judgments. For our purposes, quotations from Hare which speak of intuitions should be read as speaking of moral judgments.

Hare states that "surprisingly many philosophers, as soon as they turn their heads to a practical question...think that [moral questions]...can be solved only...by a combination of prejudice (called intuition) and rhetoric" (Hare, 1976, p. 114). This scathing criticism of philosophers who make open use of moral judgments in moral argument is directed with particular force against John Rawls, of whose theory Hare writes, "[s]ince the theoretical structure is tailored at every point to fit Rawls' intuitions, it is hardly surprising that its normative consequences fit them too—if they did not, he would alter the theory" (Hare, 1973, p. 147). The implication here is that because Rawls tailors his theory to fit his own moral judgments, we cannot, therefore, take the theory to be well founded. However, this negative implication only follows on the presupposition that moral judgments are not good evidence for moral theories. Similarly, the first quotation implies that moral judgments are no more than unphilosophical prejudices, but

the identification of moral judgments with prejudice is only compelling if we already suspect that moral judgments are not good evidence for moral principles.

### **3.1 Moral Disagreement**

Hare's main reason for thinking that moral judgments should not be used in moral theorizing is the claim that using moral judgments to evaluate moral theories does not provide us with a rational method for resolving moral disagreements. Since there "undoubtedly are cases in which people's moral opinions disagree," appeal to moral judgments must offer some way of rationally settling disagreement in moral judgment (Hare, 1997, p. 88). However, since both parties to the disagreement claim to be basing their views on their own moral judgments, there appears to be no way of settling the disagreement through moral arguments which appeal to moral judgments. For example, Sarah could claim that Andrew's moral judgment is mistaken, but this claim will have no argumentative force because Andrew could equally claim that Sarah's moral judgment is mistaken. Thus, Hare concludes, appeals to moral judgments are argumentatively impotent. Either, we already agree in our moral judgments in which case we have nothing to argue about, or we disagree about them, in which case they give us no basis on which to argue.

Case Evidentialism's reply is that the threat of argumentative impotence only looms when the disputants are limited to a small fixed number of moral judgments or their judgments are not subject to change as a result of inference from supportable generalizations. For example, imagine a moral argument between two people aptly named, Bore and Dull. Each of them has considered the same ten moral cases and each



has formed a different moral judgment about case number seven. Bore says to Dull, "In Case 7, action X is obligatory" while Dull says to Bore, "In Case 7, action Y is obligatory" (and action Y is practically incompatible with action X). Neither Bore nor Dull have enough imagination to think up any new cases for consideration and neither is able to find within their limited judgment set any supportable generalization which would suggest a common verdict in Case 7. Hare's argument succeeds in showing that appeal to moral judgments is argumentatively impotent for moral debaters such as Bore and Dull because, despite their agreement about nine out of ten cases, they are unable to say anything to each other which could resolve their disagreement about Case 7. As should be obvious, however, the kinds of moral debate which are of interest to moral theorists and the wider public, are debates in which the participants do not suffer from the severe limits to imagination and theoretical ability posited for Bore and Dull.

To put this reply positively, Case Evidentialism is committed to there being no limit in principle to the number and kinds of cases which can be considered in the course of moral argument. Disputants may appeal to cases from personal experience, from their imaginations, from literature, from history, from anthropology, etc. Further, new theoretical arguments can always be brought to bear on existing judgments by finding support for a new moral generalization. Therefore, a moral argument can never be pronounced intractable, and moral inquiry can never be pronounced complete because we are never certain that all relevant cases and theoretical arguments have been considered. All we can say is that the more variety of cases we consider and the more that we adjust our view of the cases in light of theory, the more rationally grounded our views will be overall.

Disputants can seek to resolve disagreements by establishing moral generalizations using the judged contents about which both agree. These generalizations then provide a common source of inferential evidence about particular cases. This method is certainly not guaranteed to resolve all moral disagreements, but it does ensure that we never have to stop trying. The method also allows us to claim an improvement in the rationality of our moral views when moral argument produces revisions in the moral claims we accept.

Before moving on to Hare's other arguments, let us return to the debate about abortion between Sarah and Andrew to see how Case Evidentialism can explain rational debate about a concrete moral disagreement. Sarah and Andrew disagree about the moral generalization that abortion is always wrong. If they hold their respective views reflectively, then the views are based on moral judgments that are made with knowledge of the facts of particular cases of abortion. So Sarah and Andrew must have considered difficult cases such as those in which the women who seek abortion are teenagers, have health concerns, or have been the victims of rape. About each of these cases, Sarah judges that abortion is wrong. Andrew on the other hand, considers such cases to be counterexamples to the claim that abortion is wrong. He claims that in such cases abortion is morally permissible. Therefore, Sarah and Andrew are committed to different views about the moral generalization 'abortion is wrong,' in part, because of different moral judgments about particular cases of abortion.

In the course of the debate, Sarah pressed Andrew with the argument that abortion is infanticide and is, therefore, wrong. Sarah's argument can be understood to have rational force because Andrew believes that the moral generalization 'infanticide is

always wrong' is supported by his judgments about a wide range of different cases. Thus, if Andrew's judgments that some cases of abortion are not wrong imply that some cases of infanticide are not wrong, then these judgments are called into question by the inferential evidence against counter-instances of the well supported generalization that infanticide is always wrong.

Case evidentialism can make sense of several strategies of counter-argument for Andrew to use. One response to the infanticide argument would be for Andrew to attempt to come up with a variety of non-abortion cases in which infanticide was not wrong. This would create broader support for the generalization that some infanticides are not wrong, allowing a modification of the original generalization that all infanticide is wrong which is not *ad hoc*. Another response would be to claim that the generalization about infanticide is subsumed under the larger generalization about murder, and then try to find cases of murder which are not wrong but which resemble abortion in some way. This would have the same effect as the first strategy of providing broader support for the existence of a category of exception to the general prohibition. Finally, Andrew could attempt to deny that abortion is infanticide, by arguing that the foetus is not yet an infant. As this example illustrates, the use of moral judgments as evidence does not preclude reasoned debate about controversial moral questions.

Hare pushes his argument about moral disagreement further, claiming that the "equilibrium" reached amongst moral judgments "is one which might have been generated by prejudice, and....[i]t would be possible for two mutually inconsistent systems to be defended in this way; all that this would show is that their advocates had grown up in different environments" (Hare, 1981, p. 12). What is new here is the worry

that methods which rely entirely on the evidence of moral judgments could yield mutually inconsistent moral systems, each of which would be equally justifiable given different cultural or individual starting points.

The possibility of two self-contained mutually inconsistent moral systems would be a genuine threat to Case Evidentialism only if we were at liberty to ignore the moral judgments of other people and cultures. Some moral philosophers who invoke moral judgments do talk as if the only moral judgments that matter are those inside one's own heads. For example, Rawls writes at one point, "[f]or the purposes of this book, [*A Theory of Justice*,] the views of the reader and the author are the only ones that count" (Rawls, 1971, p. 50 as cited in Hare, 1973, p. 145). This passage implies an implausibly solipsistic view of moral evidence.<sup>6</sup> As social creatures of limited epistemic means, it is unwise for us to restrict ourselves to the evidence of our own experiences in any domain. Insofar as we believe that other people are expressing well intentioned moral viewpoints different from our own, we should not ignore these viewpoints in our evaluation of the evidence for our moral views.

In general, my believing that another person has evidence for a claim gives me evidence for that claim. For example, if I read a newspaper article claiming that the Prime Minister has resigned from office, I come to have evidence for the claim that the Prime Minister has resigned. However, my evidence for this claim does not come from any observation I have made, nor from any inference from an observation I have made. Rather, the evidence arises from my belief that the journalist who wrote the story is trustworthy and has reliable sources for her claims. Similarly, my belief that Sally judges

---

<sup>6</sup> As Hare points out on the same page, Rawls himself does not stick to this solipsistic line consistently. More often he appeals to the moral judgments shared by people generally.

that the Prime Minister was (morally) wrong to resign at just this moment, gives me evidence that the Prime Minister's action was wrong. Of course, for Sally's judgment to give me *weighty* evidence I must hold other supporting beliefs about Sally. For example, I must also believe that Sally is trustworthy, sincere, knowledgeable about politics, and not confusing moral wrongness with political imprudence. Because these supporting beliefs may be difficult to evaluate in their own right, there is often more uncertainty about the quality of another person's moral evidence than there is about our own evidence. For this reason, we need not call our well established moral views into question every time we learn that somebody disagrees with them.

However, sometimes we rationally accept another person's moral view as better supported than our own as a result of supporting beliefs about that person (or group). For example, when Josh reads about a scientific experiment on human subjects that sounds unpleasant for the subjects and does not seem to be addressed towards a significant research question, he judges that the experiment is immoral. However, Josh later learns that the experiment was approved by an ethics board at the university where he studies. Josh has previous experience with this ethics board and he believes the board to be rigorous, well intentioned, intelligent, and caring. As a result, Josh changes his view on the experiment; coming to think that his original judgment about the case must have been mistaken. The partial success of the Humean ideal judge view of moral argument that we considered in Chapter 1 can now be explained; our beliefs about how good a moral judge somebody is determines the strength of the evidential weight we assign to their moral views.

Imagine now two cultures with mutually inconsistent moral systems such that within each system their particular moral judgments about all the cases they have ever reflected upon have been successfully brought under a set of highly explanatory moral principles. As long as the two cultures have never communicated with each other, then each moral system is strongly supported by that culture's collective moral evidence. For Hare, the possibility of inconsistent moral systems each being strongly supported might smack of relativism. However, this kind of relativism is the same as that present in virtually all forms of rational inquiry, namely, that the view which it is reasonable to hold is relative to the evidence at our disposal. A problematic form of *moral* relativism only follows if after the two cultures begin to communicate, they find that despite earnest attempts at rational debate, neither can rationally critique the other's system because they share so few moral judgments in common.

Case Evidentialism avoids this moral relativism because once each culture begins, through communication, to form beliefs about the other culture's moral views, these beliefs give them new moral evidence against which to evaluate their pre-contact moral systems. Moral theorists from the two cultures would then have to collaborate to search for a new system that best makes sense of the new moral evidence. It is most likely that both cultures will be forced to abandon some of their judgments and principles while having others vindicated and reinforced by the new equilibrium. Rather than providing a justification for cultural prejudice, as Hare claims, treating moral judgments as evidence explains how intercultural dialogue can lead rationally to a better moral understanding. In contrast, Hare's utilitarian methodology seems to result in the irrelevance or irrationality of respecting the diverse moral viewpoints of different cultures.

The fact that real-world moral disagreements between different cultural groups sometimes appear to be immune to rational resolution by appeal to moral judgments can be explained in a number of ways. First, the group or individual might reject moral judgments as a source of evidence. Hare himself is such an individual, and so Hare's belief that moral judgments cannot resolve moral disagreements is *self-fulfilling* because moral judgments will not be able to resolve moral disagreements among people who deny that they have evidential weight. Similarly, a culture with a strictly authoritarian conception of moral knowledge will deny outright any evidential value to the moral judgments of people who are not recognized as moral authorities. Members of such cultures are unlikely to be able to rationally resolve moral disputes with outsiders who do not recognize the culture's special source of moral authority.

Second, even if a culture accepts that everyone's moral judgments have some evidential weight, they may hold beliefs which undermine the total evidential weight they assign to the other culture's moral views. For example, Culture A may believe that members of Culture B are subject to an irrational taboo against reflectively criticizing traditional moral views. Then Culture A would assign a low evidential value to the moral views of members of Culture B because members of Culture A would believe that, overall, the views of Culture B are less rationally grounded than their own views. In such cases, intercultural moral dialogue may be severely hampered until Culture A is led to revise its negative beliefs and accept the rationality of the views of members of Culture B, or if Culture B's moral views really are less rational as a result of a taboo on reflective criticism, then moral dialogue must be preceded by a change in Culture B's taboo.

Finally, even cases where each culture assigns significant evidential weight to the moral judgments and moral generalizations of the other, the intrinsic complexity of a sophisticated moral system may prevent them from discovering a way to integrate the two systems. This last cause of intractability can be likened to the predicament of scientists during a period of scientific revolution. During a scientific revolution, a field of research whose members normally share a broad theoretical framework is split in two. In one camp are members of the old guard who cling to the pre-revolutionary framework, in the other camp are those who endorse the new framework. Although both groups of scientists are committed to the same canons of scientific rationality, their theoretical frameworks are so different and the evidence for evaluating those frameworks is so complex that the disagreement between the two groups becomes intractable (Kuhn, 1970). Similarly, two cultures that have developed complex moral systems in partial or complete isolation from one another will likely face extraordinary difficulty in trying to integrate their moral systems into a moral theory which is rationally acceptable to both cultures. It is a virtue of Case Evidentialism that it allows us to understand how intractable moral disagreements between individuals and groups arise in practice, while at the same time explaining how such disagreements can, in principle, be resolved through rational debate.

### **3.2 Levels of Moral Thought**

In defending utilitarianism against moral judgments that seem to conflict with applications of the utilitarian principle, Hare argues that appealing to such moral judgments in moral debate amounts to a confusion between two levels of our moral



thought. Hare (1976) articulates the idea that our moral thinking has two levels: "Level-1 principles are for use in practical moral thinking, especially under conditions of stress....Level-2 principles are what would be arrived at by leisured moral thought in completely adequate knowledge of the facts, as the right answer in a specific case" (pp. 122-123). Level-1 rules must be simple enough to be readily learned and applied in daily life. They must also compensate for human weakness. For example, it may be that the rule 'never torture a prisoner' is accepted as a Level-1 moral rule even if it is rejected at Level-2. It may be rejected at Level-2 because of the consideration of cases where we judge that it is morally acceptable to torture a prisoner, as in the famous ticking time bomb case. In this case, the prisoner has planted a bomb set to explode very soon which will kill many innocent people. Only he knows the bomb's location and will only tell his captors the location under torture. However, because such cases are extremely rare, and because the natural human weakness for abusing power makes it likely that any permissions with respect to torture will lead in practice to immoral acts of torture, the absolute prohibition on torture may be justified as a practical (Level-1) rule, even if our more elaborate Level-2 moral theory would claim that torture is sometimes allowable. Any acceptable account of methodology in moral theory will have to acknowledge the need for both practical moral guidelines (Level-1) and the nuanced moral theoretic principles from which practical guidelines are to be developed (Level-2).

Hare thinks that considerations about the meanings of moral language favour a version of act-utilitarianism for Level-2. He objects that the "trick of the opponents of utilitarianism is to take examples of [utilitarian Level-2] thinking, usually addressed to fantastic cases, and confront them with what the ordinary man would think. It makes the

utilitarian look like a moral monster" (Hare, 1976, p. 123). For example, consider the case in which seriously harming a child causes a small amount of pleasure to many people, such that overall utility is maximized by harming the child. In cases like this one, ordinary moral thinking is clearly at odds with the injunctions of utilitarianism and presenting us with such cases has been the primary method by which critics of utilitarianism have pressed their challenge. According to Hare, the thinking of the "ordinary man" is not reflective and, therefore, represents the deliverances of Level-1 moral principles. Since Level-1 principles are nothing more than practical tools for approximating the verdicts of Level-2, it is an obvious error in reasoning to appeal to judgments prompted by Level-1 in order to challenge those of Level-2. For example, if utilitarianism is our ultimate answer to the question of why certain types of actions are right or wrong, then we cannot object to utilitarianism by citing the fact that some utilitarian actions seem intuitively wrong to us. Hare claims that our anti-utilitarian judgments about cases such as the child-harming case are a product of internalized Level-1 rules such as 'never harm an innocent,' and that since such rules are justified by utilitarianism at Level-2, there is no way that a judgment based on the rule can count against utilitarianism.

Case Evidentialism needs to allow Hare's distinction between Level-1 and Level-2 principles while avoiding Hare's conclusion that moral judgments cannot be used to evaluate moral principles. To do so, it must deny Hare's claim that our moral judgments are determined entirely by the Level-1 principles we follow in daily life. Going back to Hare's definition of Level-1 as a set of principles "for use in practical moral thinking, especially under conditions of stress," we can point out that the moral judgments to which

we assign evidential weight are those made when we are not under stressful conditions (Hare, 1976, p. 122). The moral judgments that critics of utilitarianism appeal to are those made while striving to arrive at judgment "by leisured moral thought in completely adequate knowledge of the facts, as the right answer in a specific case" – Hare's definition of Level-2 moral thinking (*ibid.*, p. 123). Therefore, Hare is mistaken to identify the moral judgments to which we appeal in moral theorizing with the knee-jerk moral judgments of the ordinary man. To deny this identity, we do not need to claim that moral theoretical judgments are essentially different from those made by ordinary people, only that there is a difference for everyone between a gut reaction and a considered moral judgment. Hare's claim that opponents of utilitarianism confuse two levels of moral thought is now seen to be unfounded. Philosophers who appeal to moral judgments are not confused about the levels of moral thinking, they simply disagree with Hare's methodology for arriving at Level-2 principles.

### **3.3 Preferences for Paraworlds**

Catherine Wilson has developed a theory of moral theory confirmation which begins with Hare's view that "moral decisions involve the application of principles to particular situations" (Wilson, 2004, p. 35). Her theory generates a further objection to Case Evidentialism's claim that particular moral judgments are non-inferentially supported while moral principles are not. Wilson argues that in evaluating moral claims we consider "paraworlds" which are worlds populated by realistic moral agents whose actions are morally idealized (*ibid.*, p. 48). In asserting that swearing in public is wrong, I express my preference for paraworlds in which agents always refrain from swearing in

public, while asserting that swearing in public is permissible means that I prefer paraworlds in which agents sometimes refrain from swearing in public. If this semantic analysis is correct, then a non-cognitivist meta-ethics would be able to treat general moral claims as expressions of attitudes, albeit attitudes towards ideal worlds rather than attitudes towards actions. If this were the case, then the non-cognitivist would assign non-inferential evidence to general moral claims since these express attitudes which might be endorsed at the limit of reflection. Case Evidentialism must reject this modal semantics of moral statements, claiming that general moral statements are universally quantified statements about actions rather than expressions of a preference for a fictional world.

The major drawback of Wilson's account is that it cannot explain particular moral judgments at all. This is because when I make a judgment in response to a particular situation, it is often unclear what features of that situation provoked me to make the judgment. For example, if I hear Frank swear in public and judge that Frank should not have done so, how is my judgment to be explained in terms of preferences for paraworlds? If my judgment about Frank depends on very specific aspects of Frank's situation, then the relevant paraworld would be one in which idealized agents refrained from swearing only in situations very much like Frank's. If my judgment is less specific to Frank's situation, then idealized agents might refrain from swearing almost all the time. However, we do not know pre-theoretically which aspects of a situation inform our particular moral judgment.

Therefore, our particular moral judgments about real and imagined cases underdetermine the paraworlds for which we have a preference. Only by considering many particular cases and generalizing from them can we determine which rules our ideal

paraworld agents would be following. However, if we come to know our paraworld preferences through generalization from particular judgments, then the paraworld semantics does not, after all, support the assignment of non-inferential evidence to general moral truths. On this interpretation, paraworlds are simply another way of representing general moral claims, but do not play a role in our moral epistemology.

However, the paraworld theorist might deny that particular moral judgments need to be accommodated by an account of moral language. If moral statements always express attitudes towards paraworlds, then particular moral judgments are illegitimate. On this view, I should never judge that Frank's swearing in public is wrong. Rather, in considering Frank's action, I must first form a judgment about a relevant paraworld and then ask whether Frank's action conforms to what would be done by the idealized paraworld agents. If I prefer a paraworld in which no idealized agent swears in public, then I may infer that Frank's swearing in public is wrong.

Denying the legitimacy of our moral judgments about particular actions is a strongly revisionary account of moral discourse. It tells us that we must never form a moral opinion about a particular case without first forming a judgment about a fictional world of idealized agents. Wilson presupposes Hare's position that moral principles are epistemologically prior to particular moral claims. Yet, the claim that we never, or never ought to, form moral opinions in the absence of a pre-established moral principle is certainly not obvious. Rather, it is a position that we must be driven to through argument. If my defence of Case Evidentialism has been successful thus far, then Wilson's semantic analysis does not give us any new reason to doubt the evidential weight of particular moral judgments.

A final option for the paraworld theorist is for her to grant that particular moral judgments are a legitimate source of moral evidence, but to continue to claim that general moral truths may be supported non-inferentially through the consideration of paraworlds. This is a hybrid of the paraworld account of general moral claims and the Case Evidentialist account of particular moral judgments. However, this hybrid is an unhappy union. General moral claims such as 'murder is always wrong' have a logical surface structure which suggests a universal claim about actions rather than an expression of a preference for a paraworld. Therefore, a theory which respects this surface structure while offering a compelling account of moral argument is preferable to one which introduces a revisionary analysis. Further, particular judged contents inferentially support general truths understood as universal claims about actions, but they do not inferentially support general truths understood as preferences for paraworlds. Therefore, the hybrid position cannot explain how moral evidence is to flow from particular judged contents towards general moral claims. In contrast, Case Evidentialism preserves the surface logical structure of general moral claims, and allows, thereby, inferential evidence to flow from judged contents to moral generalizations.

Although I think Wilson's paraworld theory fails as a semantic analysis of moral claims, it is perfectly legitimate if it is understood as a normative moral theory. Such a theory would claim that our moral obligations are determined by the moral rules that would be followed in a realistic but idealized world in which we would most like to live. This moral theory may be quite similar to rule-utilitarianism. It can be evaluated against our particular judged contents just as any other moral theory. Understood as such, paraworlds do not introduce any difficulties for Case Evidentialism.

### 3.4 Universality of Moral Language

Hare's final argument which speaks against Case Evidentialism is that moral language, properly understood, always makes universal claims. According to Hare, one cannot judge that some unique action is wrong, only that all actions relevantly similar to this one are wrong. Hare argues that whenever one claims that a particular action is wrong, one implies that all actions which are like this one in all their morally relevant features are wrong. Stated as such, the doctrine that moral claims are always universal seems hard to deny. This doctrine creates a puzzle for Case Evidentialism since Case Evidentialism gives a special evidential role to judged contents which are particular moral claims.

Hare's claim that moral language is universal is based on the seemingly undeniable claim that any two actions which share all morally relevant features must have the same moral status. For example, one reads Nabokov's novel *Lolita* and judges that Humbert acted wrongly in sleeping with the young girl, Lolita. The particular claim, 'Humbert acted wrongly in sleeping with Lolita' entails the universal claim 'all actions which share with Humbert's action all morally relevant features are wrong.' For example, if the name of a moral agent is morally irrelevant, then my particular judgment about Humbert implies the universal claim, 'all actions are morally wrong if they are identical to Humbert's action except insofar as the agent has a different name.' Therefore, our particular moral judgments entail general moral claims, the content of which is determined by moral relevance.

Case Evidentialism responds to this line of reasoning by granting that particular moral claims entail universal claims of this sort while claiming that the entailment is epistemologically irrelevant because we do not have any account of moral relevance *in advance* of moral theorizing. Our moral evidence, which originates in our particular moral judgments, is the only ground we have for believing that certain features of an action are morally relevant while others are not. Therefore, even if every particular moral claim entails a universal claim about all relevantly similar actions, we cannot know which universal claim is entailed except in light of the moral generalizations which we establish on the basis of our judged contents. Hence, from the moral theorist's point of view, the entailment of such universal claims by particular claims is irrelevant since there is no way to know which universal claim is entailed until after we finished moral theorizing.

We cannot know prior to moral theorizing which actions are relevantly similar to the particular action about which we make a moral judgment. However, the particular case about which we make our judgment is characterized for us by a set of non-moral features. Since we know that any feature which does not serve to characterize the action for us cannot be informing our moral judgment, there is a universal claim which is entailed by a particular judged content. This universal claim takes the form, 'any action which shares with this particular action all the features by which this particular action is for me now characterized is obligatory/wrong/permissible.' Let us call this universal claim a maximally specific generalization.

A maximally specific generalization is the only universal claim which may be inferred from a judged content. This inference is one of entailment, because whatever evidence we have for our particular claim is also evidence for the maximally specific



generalization. However, the strength of our non-inferential evidence for a particular judged content depends on our non-moral understanding of the case that the judged content is about. The greater our non-moral understanding of the case, the less likely it is that we are ignoring a morally relevant feature. For example, I am told that a man stuck a sharp object into a woman's arm and I judge that the man's action was wrong. However, the fact that this judgment is informed by a very limited set of non-moral features which characterize the case for me means that the judged content is weakly supported by the moral judging. A little imagination will quickly show us why. When I realize that the man in question might be a nurse giving the woman a vaccination, I realize that the non-moral features which originally characterized the case for me left out features of moral relevance.

Maximally specific generalizations are the first step in a chain of inferences by which we ultimately justify moral principles. In Chapter 2, I said that inferential evidence flows between a judged content and a moral generalization in virtue of that judged content being an instantiation of the universal statement which expresses the generalization. However, in order to arrive at highly general moral claims, we must also make inferences between more and less specific universal statements. When there is a common feature F picked out in several generalizations and these generalizations make the same moral predication of the actions with F, this commonality provides evidence for the moral generalization that makes the common predication of all F actions. Going in the other direction, when a moral generalization 'all F actions are wrong' is well supported by inference from more specific generalizations, and a few specific generalizations are outliers which imply that some F actions are not wrong, then there is evidence against

those outlying generalizations. Moral theorizing begins by discovering maximally specific generalizations from judged contents and then works its way up through many levels of generality towards moral principles.

Having shown that Hare's claim about the universality of moral language is unproblematic for Case Evidentialism, we have taken the last wind out of the sails of Hare's critique of the use of moral judgments as evidence. In the process, we have improved our understanding of the inferential relationship between judged contents and moral principles. What has not been argued is why one should prefer Case Evidentialism to Hare's own method. Unfortunately, to launch a critique of Hare's methodology that would do justice to the complexity, rigor, and sheer breadth of Hare's philosophy would be beyond the scope of this work. It can only be noted that Case Evidentialism is more consistent with the actual methods used by most contemporary moral theorists. While Hare's methodology places us in the unfortunate position of having to claim that these moral theorists are busily systematizing prejudices, Case Evidentialism vindicates their project and gives us hope for the continuing progress of the field. This might be called an "argument from hope" against Hare's position. Such an argument would not have force on its own, but coupled with the preceding defence of Case Evidentialism against Hare's attack, the independent support for Case Evidentialism from meta-ethics, and the apparent success of work in moral theory using Reflective Equilibrium, such an argument may prove decisive.

## Chapter 4: Taking Account of Moral Particularism

In the preceding chapter we saw that Case Evidentialism faces a challenge from Hare's attack on the use of particular moral judgments in moral theorizing. Having defended Case Evidentialism against Hare's objections, we must now turn to face a challenge which arises from precisely the opposite direction. The challenge is that *only* particular moral judgments play a legitimate role in our moral thinking, that any appeal to moral principles is misguided and that moral theory is, therefore, a spurious enterprise. This, in its most general form, is the view known as moral particularism.

Moral particularism claims that moral principles may never be evidentially supported. If moral principles are never evidentially supported, then inference from them provides no evidence about particular moral claims. For example, moral particularism claims that a moral principle such as 'lying is always wrong' can never provide evidence for the particular claim, 'it was wrong for Josh to lie to his mother.' In contrast, Case Evidentialism claims that the principle 'that an action is a lie makes the action wrong' can provide evidence for the particular claim, 'it was wrong for Josh to lie to his mother,' if and only if and to the extent that the principle is supported by judged contents. Therefore, moral particularism is opposed to Case Evidentialism because moral particularism denies the evidential role that Case Evidentialism allows for moral principles.

As I have characterized moral particularism, it is a thesis about moral epistemology. Specifically, it is the thesis that moral principles are never evidentially supported and, therefore, cannot support particular moral claims. To distinguish this epistemological thesis from other theses which a moral particularist might endorse I will

call this thesis, principle scepticism. Other claims made by moral particularists, some of which are said to imply principle scepticism, will be identified as we go. We shall discover that some of the moral particularists' arguments may have important implications for moral theory, but they do not support principle scepticism.

There are two main particularist arguments for principle scepticism. The first is that observational realism, the view that moral judgments are observations of moral properties in the world, supports principle scepticism. This argument can be easily dismissed by Case Evidentialism by showing that observational realism is consistent with and supports inference from well supported moral principles. The second and more difficult argument for principle scepticism proceeds from a thesis about the behaviour of moral reasons, called moral reason holism. Presenting this argument and responding to it will occupy the bulk of this chapter. However, sections 4.3, 4.5 and 4.6 contribute to our constructive project by continuing to develop Case Evidentialism in light of legitimate particularist concerns about the behaviour of moral reasons.

#### **4.1 Observational Realism**

McNaughton (1988) writes, "the kind of moral realism which I have been advocating, with its stress on the role of sensitivity to the moral facts of the particular situation, naturally leads to moral particularism" (p. 190). McNaughton clarifies that he means to claim that moral realism leads to principle scepticism when he goes on to say that "[m]oral particularism takes the view that moral principles are at best useless, and at worst a hindrance, in trying to find out which is the right action" (*ibid.*). McNaughton thinks that the best kind of moral realism is one which understands moral judgments on

analogy with observations such as 'the moon is round.' There is obviously no principle we can invoke to tell us what we observe, we simply have to be attentive to the way that the world impinges on our minds through our senses. Similarly, the observational moral realist thinks, our only access to moral knowledge is our sensitivity to the moral features of situations which impinge on us and are captured in our particular moral judgments. Therefore, observational realism might appear to support principle scepticism.

The apparent support of observational realism for principle scepticism can be diffused by pointing out that observational realism only supports the claim that our *non-inferential* evidence is restricted to particular moral judgments. Moral observations are statements about our experience of moral reality, and experience is necessarily particular in that we can only experience one moral situation at a time. Therefore, observational realism implies that only our particular moral judgments are non-inferentially supported and, in this, it agrees with the first and second theses of Case Evidentialism.

In Chapter 2, we took great pains to show that attributing non-inferential evidence exclusively to particular judgments does not imply scepticism about moral principles. We saw that inferences from particular judgments can support moral generalizations. Subject to some elaboration to be spelled out later in this chapter, we may say that our moral principles are our best supported moral generalizations. Hence, the claim that particular judged contents can support moral principles follows from the true claim that moral generalizations can be supported by inductive and explanatory inferences. Once we allow that moral principles can be evidentially supported, then inference from these supported claims provides evidence which can count for or against particular moral claims. Therefore, McNaughton is wrong to claim that observational realism supports principle

scepticism. To establish principle scepticism, we need a further argument to convince us that no moral principles are supported by inference from our particular judged contents.

This further argument is offered by Jonathan Dancy (1993; 2000; 2004) who develops a holistic theory of moral reasons from which he thinks principle scepticism follows. Dancy believes that "undistorted views about how reasons work lead naturally to a particularist account of moral thought and judgment" because such an undistorted view of reasons will counter the "pervasive tendency to think that all reasons are grounded in principles" (Dancy, 2004, p. 15). Dancy argues for a holistic view of moral reasons, which he thinks supports a particularist account of moral thinking which has no place for moral principles.

#### **4.2 Moral Holism Defined**

Dancy (2004) defines holism as the claim that "a feature that is a reason in one case may be no reason at all, or an opposite reason, in another" (p. 7). So, for example, even if lying is found to be a moral reason against an action in a large number of our particular moral judgments, holism claims that in other situations lying may not be a reason against an action. However, the claim that a feature which is habitually one sort of moral reason "may be" another sort of moral reason in another case is either too weak to do the work that Dancy needs or it is intended in a sense which is misleading. The most obvious way to interpret "may be" is as a kind of possibility. Yet, Dancy cannot be claiming that it is merely possible for a feature to change its moral reason giving function, in either the modal sense of possibility (true in a possible world), or in the epistemic sense of possibility (a hypothesis which is open to investigation).

If Dancy's definition was meant to imply modal possibility, then holism would claim that even if lying makes actions wrong for all situations in the actual world, there is some possible world in which lying does not make actions wrong. Dancy's use of examples from the actual world to support holism shows that this is unlikely to be the correct interpretation of his view (Dancy, 1993, pp. 60-61). Further, even if such a modal contingency were generally true of moral statements, it is not obvious why the fact that our principles fail to capture our judgments about every possible world, should make them unacceptable in our world. Empirical generalizations (at least some of them) are modally contingent, but that does not stop us from having reason to believe them to be true of our world.<sup>7</sup>

Holism also cannot be taken to assert merely that we must consider the hypothesis that lying could be a moral reason against some actions but not all. This claim is easily shown to be consistent with the denial of principle scepticism. Case Evidentialism endorses the claim that a moral theorist should be open to considering hypotheses which assert that a feature functions as a moral reason differently in different situations. But Case Evidentialism combines this openness to holist friendly hypotheses with the claim that a well supported moral generalization can give us evidence for or against such hypotheses. Dancy's holism would not allow moral generalizations to be relevant in this way.

Therefore, in order to have the epistemological bite needed to support an argument for principle scepticism, moral reason holism must go beyond the claim that it is merely possible for features to change their reason giving function. Holism must make

---

<sup>7</sup> There is a debate about the evidential value of moral cases which could only arise in possible worlds very unlike our own. I will not enter into this debate, except to say that Case Evidentialism gives us no reason to ignore moral judgments about merely possible cases.

the stronger claim that such possibilities are never improbable; no matter how often a feature counts as a moral reason in a particular way, we are not, thereby, licensed to infer that it is *likely* that the feature will *always* count as a moral reason in that way. For example, no matter how many particular cases there are in which we judge that lying makes an action wrong, we should not infer from this that lying is likely to always make actions wrong. I will take the final and most adequate definition of moral reason holism to be: No matter how many different situations there are in which a feature of an action functions as a moral reason in a certain way, we may not infer from these cases that it is more likely that this feature will always function as a reason in that way.

To properly understand the significance of moral reason holism, we must understand what is meant by 'moral reasons' and how moral reasons fit within Case Evidentialism. A moral reason is a morally relevant consideration that either speaks in favour of doing and action or speaks against doing the action.<sup>8</sup> In Chapters 2 and 3, we glossed over questions about moral reasons and focused on questions about whether an action is morally obligatory, morally permissible, or morally forbidden. These questions require us to give a moral verdict on an action by evaluating its overall moral status. However, equally important in moral thinking are questions about whether a certain

---

<sup>8</sup> There are other kinds of reasons for or against actions, the most prominent being prudential reasons. We can crudely distinguish these kinds of reasons by saying that a prudential reason for an action is a reason which stems from self-interest whereas a moral reason stems from our concern for others. However, providing a rigorous articulation of this intuitive distinction between moral and prudential reasons has proven notoriously difficult. Wilson (2004) has attempted such an articulation by claiming that moral reasons are advantage reducing norms, in that following moral imperatives may involve self-sacrifice. Korsgaard (1996) provides an alternative account of the distinction by arguing that all reasons for action arise from practical identities. Non-moral reasons arise from practical identities which are contingent in the sense that we could cease to hold those identities without ceasing to have any reasons for action at all. In contrast, moral reasons arise from our non-contingent practical identity as reflective creatures. We cannot abandon our identity as reflective creatures without abandoning all other practical identities, and therefore moral reasons are less optional than other kinds of reasons for action. Whichever, account one prefers, precisely how we are to draw the dividing line between moral reasons and non-moral reasons for action is less important than our ability to distinguish them approximately in most cases.



feature of an action makes that action morally better or worse. These are questions about the moral reasons which bear on our actions. Just as we make particular moral judgments about overall questions such as 'Jack's murdering this innocent woman is morally forbidden,' we also make particular moral judgments about moral reasons such as 'the harm caused by Jack to the woman makes his action morally worse.' Similarly, just as there are overall moral generalizations such as 'murder is always morally forbidden,' there are also generalizations about moral reasons such as 'harm makes all actions morally worse.'

Evaluating the balance of moral reasons for each of the actions in a situation helps us to evaluate what we ought to do overall. For example, if Johnny considers whether to roast ants with a magnifying glass on a sunny day, and he judges that the cruelty of this action is a moral reason against doing it, then recognizing this moral reason will incline Johnny towards thinking that roasting ants is morally forbidden. However, our evaluation of moral reasons only informs and does not determine our evaluation of overall moral questions because moral reasons do not tell us how to draw the boundaries of moral obligation. Johnny might judge that the cruelty of roasting ants results in a balance of moral reasons against doing the action. However, he might also judge that the balance of moral reasons against the action is so slight, that it does not generate a moral obligation to refrain from ant roasting.

In general, judgments about the balance of moral reasons do not tell us whether a certain morally unfavourable action is morally forbidden or whether a morally favourable action is morally obligatory. Refraining from permissible morally unfavourable actions and performing permissible morally favourable actions defines the class of

supererogatory actions, actions which are morally meritorious beyond the call of duty. Therefore, overall judgments cannot be reduced to judgments about the balance of moral reasons.

We are now in a position to introduce a few more pieces of terminology which will greatly simplify subsequent discussion. Following Dancy, I will call particular moral judgments about overall questions 'overall judgments' and particular moral judgments about moral reasons 'contributory judgments' (Dancy, 2004, p. 15). The latter term is sensible in light of the fact that moral reasons contribute to the moral worth of an action. For example, 'Frank's action was morally forbidden' is an overall judgment, while 'the harm caused by Frank's action made his action morally worse' is a contributory judgment. I will use similar terms for moral principles. A moral principle of the form 'all F actions are morally forbidden' is an overall moral principle, while a principle of the form, 'that an action is F always makes that action morally better' is a contributory principle. I will also extend the terminology from Chapter 2 when necessary so as to distinguish between contributory moral judgments and contributory judged contents.

#### **4.3 Case Evidentialism and Moral Reasons**

Case Evidentialism's account of the relation between particular judgments and moral generalizations is complicated by the fact that we have (at least) two distinct kinds of particular moral judgment, overall and contributory. Contributory judged contents provide inferential evidence only for contributory principles while overall judged contents provide inferential evidence only for overall principles. However, overall moral judgments are informed by contributory judgments.

We can capture this relation of "informing" by analogy with the way that our non-moral understanding of a case informs our moral judgments of that case. We argued above that moral judgments which are informed by a more complete non-moral understanding of a case provide stronger non-inferential evidence for their judged contents because it is less likely that a morally relevant feature will be ignored. Similarly, *overall* moral judgments which are informed by a more adequate *contributory* understanding of the case provide more non-inferential evidence for their judged contents, because we are less likely to be ignoring a moral reason present in the case.

For example, Sharon makes an overall judgment that it is morally forbidden for her son, Peter, to engage in sex with his boyfriend. Simply as a result of making this overall judgment, she has some non-inferential evidence for the claim that it is morally wrong for Peter to have sex with his boyfriend. However, the strength of this non-inferential evidence will depend on her non-moral understanding of the case, as well as her understanding of the moral reasons relevant to the case. If Sharon has little knowledge about the non-moral facts surrounding homosexuality, generally and little knowledge of Peter's relationship with his boyfriend, then her overall judging gives less non-inferential evidence for its content. Similarly, if she has not considered the balance of moral reasons relevant to the case by assessing contributory judgments about it, then her overall judging also gives less non-inferential evidence for its judged content.

Peter might argue against her overall judgment by pointing to contributory reasons that she has not considered, such as the pain it brings him to suppress his true sexual identity. Becoming aware of Peter's contributory judgment against denying himself sex with his boyfriend gives Sharon new evidence about the balance of moral

reasons in the case. A change brought to her assessment of the balance of moral reasons by this evidence would incline her against her original overall judgment that Peter is morally forbidden to have sex with his boyfriend. On the other hand, Sharon may come to realize that Peter's boyfriend is manipulative and believes that Peter will likely be harmed by his boyfriend's lack of genuine commitment. She then makes the contributory judgment that this potential harm is a moral reason against Peter continuing the relationship. Such contributory factors determine Sharon's assessment of the balance of moral reasons, and changes in this balance will incline either towards or against her original overall judgment that Peter is morally forbidden to have sex with his boyfriend.

Continuing our elaboration of Case Evidentialism, well supported contributory principles are claimed to provide inferential evidence for or against particular contributory judgments which fall outside them. So for example, if the principle 'causing harm is always a moral reason against an action' is well supported by many contributory judgments, then a particular judging that the harm caused by an action is not a moral reason against the action is cast in doubt. For example, Dancy (2004) suggests that we pre-theoretically judge that cases in which harm is justly inflicted on a criminal by the state is an instance where harm is not a reason against the action. However, if this judgment is an outlier which contradicts our well supported generalization that harm is always a moral reason against an action, then, according to Case Evidentialism, we have inferential evidence against the judgment. We must evaluate whether the inferential evidence from the general claim about harm is stronger than the non-inferential and inferential evidence supporting the judgment about instances of criminal punishment. If so, then we should revise our pre-theoretical judgment and claim that the just harm

inflicted by the state on criminals *is* a moral reason against the state's action, but that there are other moral reasons in favour of the action which overrule the reason given by the harm.

If the moral reason holism defined in the previous section is true, then Case Evidentialism's claim that well supported contributory principles provide evidence for or against our particular contributory claims is false. Moral reason holism holds that we never have reason to believe contributory principles because the claim that particular judgments provide evidence for contributory principles rests on the (supposedly false) assumption that atypical cases are made less probable the more we pile up typical cases. Therefore, although we can note that harm is almost always judged to be a moral reason against an action, we are not thereby licensed to assign any plausibility to the contributory principle that harm is always a moral reason against an action. To the holist, therefore, if we pre-theoretically judge that the state's harming Bobby the cat-burglar by imprisoning him is not a moral reason against the state's action, then this is what we should believe about the case. We should not allow our knowledge that harm is usually judged to be a reason against an action to influence our judgment that in this particular case it is not.

This kind of principle scepticism is worrisome to moral theorists because a lot of apparently fruitful work in contemporary moral theory has been couched in terms of contributory principles. Therefore, most of the philosophers who have responded to Dancy's work attempt to argue against holism in order to save contributory principles. In the following section we will examine the argument for and against moral reason holism.

#### 4.4 Arguments for and Against Holism

Dancy (2004) has two main ways in which he tries to establish that moral reasons are holistic. First, he argues that non-moral reasons are obviously holistic and that there is no good reason to think that moral reasons are special in this respect. However, this argument fails to be convincing because, those who challenge holism with respect to moral reasons are very likely to challenge it with respect to prudential and epistemic reasons as well. For example, we could point to the science of game theory as a very well developed theoretical account of the behaviour of prudential reasons. Similarly, in the case of epistemic reasons for belief, the anti-holist will argue that features such as evidential support, entailment, etc., always give us reasons for belief. Therefore, the real case for holism about moral reasons rests on Dancy's second strategy of argument which is to provide cases in which our contributory judgments provide evidence in support of holism.

By considering a wide enough range of situations, Dancy thinks that we can see that wherever a feature seems to act universally as a reason in a certain way, there are cases where our contributory judgments tell us that the feature functions differently. For example, "[i]f I take pleasure in your discomfiture, this is not a redeeming feature of the situation....the way in which pleasure and fun function as reasons is logically dependent on the nature of the activity we are enjoying" (Dancy, 1993, p. 61). The claim is that pleasure really does give us a reason to do an action in which pleasure is taken in something which is morally good or neutral, but it does not give us reason to do an action in which pleasure is taken in something morally bad. Therefore, there is no true

contributory principle of the form 'there is a moral reason in favour of all actions which cause pleasure.'

The common anti-holist response to such examples is to argue that they can be perfectly well understood along non-holist lines. They claim that by considering cases in which pleasure is taken in evil, we discover that something we thought was a universal moral reason giving feature (such as pleasure) actually does not give us a moral reason at all. Perhaps, it is only "permissible pleasure" which gives us a moral reason in favour of an action and this does so universally.

Dancy is unimpressed by such suggestions, because he thinks that they are only convincing to somebody who is already under sway of the idea that contributory principles must be adequate to our particular judgments. He claims that, "*pre-theoretically* particularism is a perfectly natural view with obvious attractions. It is adherence to theory of some sort that upsets the apple cart, and causes us to abandon our initial particularist intuitions" (Dancy, 1993, p. 62, his emphasis). In Dancy's view, an evaluation of the moral evidence given by our contributory judgments supports holism just so long as we do not start out with the belief that moral judgments must accommodate principles.

In contrast, pro-principle philosophers suppose that the imperative to fit particular judgments into generalizations "underlies not only rationality in ethics, but...rationality, and its exercise in enquiry, as a whole" (Crisp, 2000, p. 34). Case Evidentialism agrees with this imperative towards generalization, claiming that a well-supported inductive and explanatory generalization provides evidence which allows us to improve our particular moral judgments through rational revision. If we accept that there is a rational imperative

towards generalization, then it is often possible to offer a theoretical explanation which makes Dancy's holistic analysis of cases appear mistaken.

As another example, Dancy offers us the case of Contraband, a family game in which some players must lie to another player as part of the fun of the game. Dancy claims, "[t]hat an action is a lie is commonly a reason not to do it; here [in playing Contraband] it is a reason in favour" (Dancy, 1993, p. 61). This analysis of Contraband supports holism by showing that although lying usually is a moral reason against an action, in this particular case it is a moral reason in favour.

However, pro-principle philosophers would claim that Contraband does not establish holism but merely shows that the principle 'lying is always a moral reason against' is incorrect. Stratton-Lake argues that "[t]he fact that our act is one of lying gives us moral reason not to do it only because (and if) in lying we break an implicit promise to tell the truth, or harm someone" (Stratton-Lake, 2002, p. xxxv). Although the Contraband example counts against the moral principle 'lying is always a moral reason against,' it does not count against the moral principles 'promise-breaking is always a moral reason against' and 'causing harm is always a moral reason against.' The latter two moral principles explain why lying usually contributes to the wrongness of actions, because lying usually involves breaking a promise or causing harm. Furthermore, the two principles explain our judgment that lying in the case of the Contraband example is not a moral reason against playing the game, because in this case, the lying neither breaks a promise nor causes harm (and the pleasure the lying brings to the game is a moral reason in favour).



The problem Dancy sees with the sort of analysis offered by Stratton-Lake is that it employs moral principles to revise our contributory judgments about lying. But this use of moral principles to revise our judgments is illegitimate if principle scepticism is true. Dancy claims that our pre-theoretical contributory judgments tell us that lying is a moral reason against ordinary cases of deception such as lying to your boss so that you can sleep in an extra hour. Against these pre-theoretical judgments, Stratton-Lake's principled explanation of Contraband requires us to accept the claim that lying itself is *never* a moral reason at all, but only appears to be because of its strong correlation with promise breaking and harm. The theoretical explanation preserves the contributory judgment that lying is not a moral reason against playing Contraband by asking us to abandon our contributory judgments that lying is a moral reason against ordinary acts of deception. If we already feel that it is important to interpret our judgments so that they are able to be captured under a parsimonious set of general principles, then this revision to our contributory judgments about lying will be thought to illuminate more clearly the behaviour of moral reasons by distinguishing genuine reasons from merely apparent ones. However, precisely because Dancy denies that there is an interpretive obligation to fit judgments to principles, it seems to him that such revisions to our pre-theoretical judgments conceal and distort the true behaviour of moral reasons.

Examples such as Contraband both support holism and do not support holism, depending on one's ideological starting point. When we entertain the thought that principle scepticism is true, we have greater deference to our pre-theoretical judgments and, thereby, find examples which seem to support holism. Equally, when we entertain the thought that principle scepticism is not true, we are more inclined to revise judgments

in light of well supported moral principles and thereby, explain away the apparently holistic implications of examples such as Contraband.

The fact that both holists and anti-holists can legitimately claim that their position is supported by the evidence, makes the debate over holism look like a dialectical standoff. However, if the examples provided as evidence for moral reason holism are only convincing to one who already accepts the very principle scepticism that reason holism is meant to support, then Dancy is not in a very strong position. He seems to need to presuppose his conclusion in support of his main premise.

We can make the awkwardness of Dancy's position explicit by writing out his argument for principle scepticism as follows:

**Premise 1)** Moral reason holism: No matter how many particular situations there are in which a feature of an action functions as a moral reason in a certain way, we may not infer that this feature will always function as a reason in that way.

**Premise 2)** Contributory principles assert that a feature always functions as a reason in a certain way.

**Conclusion 1)** Therefore, we may never infer contributory principles from particular moral judgments.

**Premise 3)** Only particular moral judgments are supported non-inferentially. (They are supported because they are observations of moral reality).

**Conclusion 2)** Contributory principle scepticism: we have neither non-inferential nor inferential evidence for contributory principles.

Moral reason holism in combination with the claim that only particular moral judgments are non-inferentially supported gives us a deductive argument for scepticism about contributory principles. Yet, as we have seen above, Premise 1 of the argument can only be supported by the evidence of examples such as Contraband, if we have already assumed the truth of Conclusion 2. Therefore, we might charge Dancy with arguing in a circle. Dancy would likely respond to this charge by pointing out that the pro-principle philosopher is forced to argue in just as much of a circle. The pro-principle philosopher must assume that Conclusion 2 is false in order to claim that the evidence of particular judgments does not support Premise 1.

However, the situation really is better for the pro-principle philosopher because the pro-principle philosopher has several strong arguments against holism which do not involve the evidence of particular cases. First, the pro-principle philosopher can argue that the truth of moral reason holism undermines the very inferences which were used to support it, making the claim self-defeating. Moral reason holism is, itself, a general claim about moral reasons which Dancy attempts to establish from the evidence of our particular moral judgments. Yet, if moral reason holism is true, then no general claim about moral reasons can be inferred from particular moral judgments.

Second, the pro-principle philosopher can make use of Hare's objection from moral disagreement presented in Chapter 3. Hare objected that we cannot rationally resolve moral disagreements by relying on our moral judgments. This objection was

avoided by Case Evidentialism with the help of moral generalizations. Since moral particularism denies that such generalizations should be used in moral argument, we are left with no account of how disagreements between different people's particular moral judgments can be rationally resolved. Dancy gestures at a particularist response to this objection by claiming that the consideration of one particular judgment can be relevant to the consideration of another:

Particularists claim that the move from what we say about one case to what to say about another is never automatic. But that does not show that appeal to one case can do nothing whatever to reinforce one's tentative judgement about another (Dancy, 2004, p. 157).

It is one thing to assert that one particular judgment can reinforce another, but quite another thing to give an epistemological account of how this is possible. Dancy leaves it entirely mysterious *why* considering a new moral judgment sheds light on and helps us to revise old moral judgments. Underlying this mystery is the fact that there does not seem to be an inferential relationship between one particular judgment and another. We must first infer moral generalizations and then from these supported moral generalizations we may inferentially support or undermine particular moral claims.

Think back to Chapter 1 and the example of hating green peppers. There is no reason why my liking of red peppers should incline me towards revising my opinion of green peppers because there is no inference directly from one judgment to the other and because there is no reason to generalize from such judgments. If moral judgments are disconnected from generalizations like judgments of food preference, then moral disagreements over particular moral judgments would be rationally intractable. By denying the possibility of using particular judgments to support moral principles, Dancy

has left himself with no way to account for our ability to rationally criticize and revise our moral judgments.

These arguments give us strong reason to doubt the truth of holism and principle scepticism. There may nevertheless be an important insight in Dancy's work which does not imply principle scepticism. This insight is that our best account of the behaviour of moral reasons may not involve the use of contributory principles. In the next section, I develop from Case Evidentialism an account of moral principles which allows for the possibility that contributory principles are inadequate to capture the behaviour of moral reasons.

#### **4.5 Beyond Contributory Principles**

We will begin by investigating why Dancy singles out contributory principles as a target for particularist criticism. Recall that a contributory principle is one that takes the form: a feature F is always a moral reason for/against an F action. For example, 'that an action involves telling lies is always a moral reason against doing it.' The only contrasting type of principle that we have so far come across are overall principles, such as 'murder is always morally forbidden.' Dancy thinks that it is much less controversial that overall principles fail to capture the full nature of our moral thought, than that contributory principles fail. Therefore, he focuses on arguing against contributory principles.

Overall principles are thought to be inadequate because we need to make sense of the idea that moral duties can conflict in a particular case, and the notion that when they do conflict, we are subject to genuine moral regret for having been unable to fulfil one of our duties. Overall moral claims cannot easily make sense of our judgements of

conflicting duty or of moral regret, whereas contributory claims can. We can say that when there are moral reasons against an action which we are nevertheless morally required to perform, we ought to feel moral regret and that this regret has implications for future actions.

For example, Francis makes a solemn promise to meet his love, Jennifer, at the airport to say goodbye before Jennifer leaves on a long journey. On the way to the airport, however, Francis sees a woman who is hurt by the side of the road. Francis's overall moral judgment tells him that he must stop to help the woman, even though he knows that doing so will make him miss his rendezvous with Jennifer. However, Francis also judges that his promise to Jennifer gives him a moral reason not to stop to help the woman. Further, after helping the woman and missing his meeting with Jennifer, Francis judges that he has a moral obligation to do something to make it up to Jennifer for breaking the promise. Overall principles are simply inadequate to capture the moral complexity of such situations. As David Ross argued, to do them justice we need to give an account of the various conflicting considerations that are involved in complex cases (Ross, 1930).

Therefore, Dancy thinks, in arguing against moral principles, overall principles are an easy target. He takes it that Ross largely accomplished the task of demolishing the hopes for a successful moral theory featuring overall moral principles. Ross claimed that we can only ever hope to articulate contributory principles, and that overall moral judgments must be left to our evaluations of particular cases. This scepticism about overall moral principles might be supported by a weaker form of holism than the moral reason holism Dancy advocates. We might call this weaker holism 'moral verdict holism'

and it would be expressed by the claim that we cannot infer generalizations like 'murder is always morally forbidden' from our overall moral judgments about particular cases. Dancy simply takes Ross one step further, by arguing that moral reasons themselves are holistic, and cannot, therefore, be subject to contributory principles. As we have seen reason to reject moral reason holism, we should also reject moral verdict holism. However, rejecting moral verdict holism is compatible with the moral theoretical claim that our moral evidence counts against moral theories which contain only overall principles. It may be that this moral theoretical rejection of overall principles was all that Ross intended in the first place.

Seeing moral reason holism as one step beyond moral verdict holism leads us to ask whether there is a next step which would allow us to construct moral theories which give a principled account of moral reasons but which do not employ contributory principles. We begin by noting that Dancy himself endorses something called a default reason: "A default reason is a consideration which is reason-giving unless something prevents it from being so" (Dancy, 2004, p. 112). The existence of default reasons is consistent with the denial of contributory principles because it does not imply that a feature always functions as a reason in the same way. Default reasons can be expressed by a moral principle such as 'undefeated, cruelty is always a moral reason against an action.' Such principles, which we may call default principles, are like contributory principles except they allow that sometimes the reason giving function of a feature may be defeated by some other feature of the situation.

We must not make the mistake of thinking that a default principle is the same as a rule of thumb such as 'cruelty is a reason against an action most of the time.' The first

difference between such a rule of thumb and a default principle is that the rule of thumb says nothing about why cruelty sometimes fails to be a reason against an action. Rules of thumb are not sufficiently rigorous to play a role in moral theorizing because they give us no guidance in explaining or predicting the cases that are exceptions to the general rule. In contrast, a default principle states that a feature always acts as a reason in a given way unless some other feature is present in the situation which alters its reason giving function.

The second difference is that the feature picked out by a default principle may not function as a reason most of the time. It could be that the reason giving function is usually defeated or that the reason giving function is dormant by default and requires a special feature to enable it. For example, consider the default principle, 'when enabled, that an action is romantic is always a moral reason in favour.' The claim is that being romantic is a reason in favour whenever it is activated by the appropriate circumstances, such as being in love with the person to whom the romance is directed. Yet, most of the time, such enabling conditions will not obtain. So even if the romance default principle were true, the rule of thumb that romance is usually a reason in favour would be false because we are not usually in love with people towards whom we are capable of performing romantic actions.

As we can see from the consideration of the romance default principle above, in order to evaluate the plausibility of a default principle against our moral judgments, we must know what the defeating and enabling conditions are that (de)activate the default reason. However, this shows that default principles alone are unable to constitute a moral theory. An unaccompanied default principle gives us no account of when the default



reason is activated or deactivated. Therefore, default principles are partial theoretical tools which stand in need of completion from a general account of defeaters and enablers. Let us call a feature which modifies a default reason, such as an enabler or a defeater, a reason modifier. In order for default principles to give us a moral theory, they need to be complimented by reason modification principles.

For example, the default principle 'undefeated, lying is a reason against an action' must be complimented by one or more reason modification principles that tell us when lying does not function normally. A reason modification principle that could plausibly serve this role would be the principle of informed consent. This principle states, 'when a feature gives a moral reason against an action because this feature constitutes a moral wrong to an individual, then this reason is defeated (or much diminished) if the individual gives informed consent to the taking of the action with the feature in question.'

So lying in the case of Contraband is not a reason against playing the game because the default reason given by lying is defeated by the fact that all players of the game consent to be lied to in the course of the playing. The difference between this theoretical explanation of the Contraband case, and the one we examined earlier by Stratton-Lake, is that this explanation does not try to find a feature which invariably acts as a reason in the same way. Rather, it offers a principled account of the way in which lying serves as a reason in some cases but not in others. If, as I am claiming, the cases which Dancy cites as evidence of reason holism actually support the moral theoretical claim that contributory principles are inadequate, then moral theorists should investigate alternative generalizations of these cases which use default principles in combination with reason modification principles.

#### 4.6 Moral Principles

From our consideration of overall, contributory, default, and reason modification principles we can advance a general conception of what it is to be a moral principle.

Throughout Chapter 2, I elaborated Case Evidentialism primarily in terms of moral generalizations, mentioning moral principles only in passing. I have said that some of our moral generalizations could be rightly called moral principles, but did not explain what it takes for a generalization to count as a principle.

The moral principles of a given moral theory are those moral generalizations that are considered by the moral theory to be ineliminable. When we claim that a moral generalization is a moral principle, we are claiming that our best and most parsimonious account of the moral evidence will feature that generalization. For example, in common sense morality, 'murder is wrong' is a candidate for a moral principle, while 'murder is wrong on Sunday' is merely a correct generalization. Since common sense morality holds 'murder is wrong' to be a correct generalization and since 'murder is wrong on a Sunday' can be derived from 'murder is wrong,' then 'murder is wrong on a Sunday' can be eliminated from the final statement of the common sense moral theory.

Given a hypothetical moral theory which only endorses principles such as 'murder on a Sunday is wrong,' 'murder on a Saturday is wrong,' etc., but which does not endorse the broader generalization that 'murder is wrong,' then 'murder is wrong on a Sunday' is among the moral principles of the theory. It may sound counterintuitive to claim that such a specific and idiosyncratic generalization could count as a moral principle of any theory, but this counterintuitiveness arises from the fact that we would be unlikely ever to

construct such a theory and we are certainly unlikely to seriously consider it for acceptance. It is actually a point in favour of this way of drawing the distinction between principles and non-principles that we do not have to presuppose that a generalization is a plausible candidate for a correct moral principle in order to call it a principle at all. Allowing implausibly narrow generalizations to count as moral principles of some theories ensures that we are not sneaking any normative substance into our mere definition of a moral principle.

Having distinguished moral principles from other moral generalizations, we need a name for these other generalizations. Let us call a generalization that is held by a moral theory to be correct, but which is ultimately eliminable, a mid-level generalization. Mid-level generalizations are general moral claims that are less than maximally specific but which are not as general as possible given our overall theoretical commitment. They fall between maximally specific generalizations and moral principles on the scale of specificity.

Although mid-level generalizations do not express our best explanation of moral judgments, they are essential way stations in the process of moral theorizing.<sup>9</sup> Our cognitive limits may not allow us to test our moral principles directly against particular moral judgments. Our moral principles are likely to encompass so many particular judgments that it is not within our capacity to evaluate the balance of evidence in favour of a principle from the particular judgments themselves. Instead, we use the evidence of

---

<sup>9</sup> In addition to playing an important role in moral theorizing, mid-level generalizations also play an important role in applied ethics where there is a practical value in possessing a set of more specific moral generalizations that are especially relevant to the area under consideration. For example, in biomedical ethics, the principle of beneficence as it features in moral theory is too broad to be of use. Here it helps to develop a mid-level generalization of medical beneficence which is specific to the doctor-patient relationship. Unfortunately, such area specific mid-level generalizations are called moral principles in the bioethics literature, which could lead to some confusion (see Daniels 1996).

particular moral judgments to support a manageably sized mid-level generalization, and several such generalizations, in turn, support broader generalizations. These broader generalizations might support still broader ones, and we continue this process as far as we can go. The generalizations from which we are unable to generalize any further make up our set of moral principles.

The conception of moral principles as ineliminable generalizations of a moral theory shows overall and contributory principles to be two varieties among many. In particular, this conception of moral principles allows for default principles and reason modification principles. Although Dancy's analysis of cases like Contraband do not support holism or principle scepticism, they may support the claim that contributory principles are inadequate. If so, default and reason modification principles might be better supported by our particular moral judgments.

A further argument against moral reason holism is generated by the finding that the evidential relationship between particular judgments and default or reason modification principles is logically no different than that between particular moral judgments and overall or contributory principles. In both cases, the principles stand as attempts to parsimoniously generalize from our particular judged contents. Therefore, we should not be committed to a pre-theoretical bias against overall and contributory principles. The fact that Dancy's moral reason holism expresses such a bias is another problem with his position. Once we see that contributory principles are related to the moral evidence in the same way as other types of principles which do not attribute invariant reason giving behaviour to features of a situation, we come to see that moral reason holism is epistemologically arbitrary.

#### 4.7 Infinite Holism and Scepticism

While the conclusion that moral reason holism is epistemologically arbitrary is another victory for pro-principle philosophers, Dancy is a principle sceptic to the end. Against the idea that the moral theorist could support default principles in combination with reason modification principles, Dancy (2004) says, "we should not yield to the pressure to turn our defeasible generalizations into much more detailed universally quantified ones" because "the latter would not...be able to survive exceptions" (p. 113). In other words, although we can hold default principles (what Dancy is calling defeasible generalizations) we should not try to combine them with reason modification principles because the combination will fall prey to the same kind of counterexamples that are said to support holism about moral reasons. Dancy confirms that his rejection of a principled explanation of reason modifiers is based on his commitment to holism when he writes, "[t]here neither is nor need be a detailed...account of which cases of cruelty are ones where the cruelty is wrong (given our commitment to holism)" (*ibid.*, p. 114).

Dancy would be mistaken here if he meant to claim that holism about moral reasons counts against our ability to support reason modification principles. Holism about moral reasons only counts against our ability to support contributory principles. In order to develop an argument against reason modification principles, Dancy needs to endorse a further holist claim about reason modifiers. It is only natural that Dancy should endorse reason modifier holism as well as reason holism. However, making such a claim would only push the debate between particularists and generalists one step further, towards the adequacy of reason modifier modification principles. Thus, Dancy endorses holism all

the way up, through infinite levels from reasons to reason modifiers, modifiers of modifiers, to modifiers of modifiers of modifiers, etc. Infinite holism claims that we may never infer the likelihood of any universal moral statement from any set of particular moral judgments. This infinite holism seems to prevent all attempts at moral theorizing by fiat. It is far from clear how one could possibly have evidence for such a grand claim.

The pro-principle philosopher need not dispute the truth of infinite holism because the scepticism which follows from infinite holism is a double edged sword which the moral particularist is ultimately unable to grasp. Infinite holism brings about a sceptical problem for our particular moral judgments which Dancy himself acknowledges, writing:

This claim [that particularism renders moral knowledge unattainable] derives from the holism of reasons and the ever-widening circle of enablers and disablers, enablers for enablers, enablers for disablers..., and so on.... No matter how much one knows, there will always be something more that one doesn't know, but whose truth or falsehood makes all the difference (Dancy, 2004, p. 158).

Dancy responds to this sceptical worry by arguing that it sets the bar for knowledge too high. He argues that if in order to know that *p* "one has to know the truth-value of everything whose status would or might affect the truth-value of *p*, one could never know anything at all" (Dancy, 2004, p. 158). To avoid this general sceptical conclusion, Dancy argues, we need to find an epistemic filter which allows us to claim moral knowledge despite our inevitably finite ability to rule out potentially relevant considerations. Dancy suggests that "all and only aspects that the agent is capable of recognizing and/or at fault for not recognizing or for not being able to recognize can pass the filter and so affect the morality of the action" (*ibid.*, p. 159). By endorsing this

epistemic filter, Dancy avoids the inference from infinite holism to scepticism about particular moral judgments.

What Dancy fails to acknowledge is that applying an equivalent epistemic filter to our evaluation of moral principles blocks the inference from infinite holism to principle scepticism. Such a filter would state: all and only the particular judgments we are capable of making and/or at fault for not making, or for not being able to make can pass the filter and so affect the correctness of moral principles. This filter tells us that in order to count as moral knowledge, our moral principles only need to be adequate to the moral verdicts, moral reasons, reason modifiers, etc., that are brought to our attention through the particular moral judgments that are accessible to us in moral inquiry. Another way of saying this is that our moral principles must be adequate to the moral evidence we have at our disposal. This is, of course, what Case Evidentialism has advocated all along. Dancy's position is that the epistemic bar must be lowered so that our particular judgments can count as knowledge despite the implications of holism, while also surreptitiously raising this same bar back up again in order to argue that infinite holism implies principle scepticism – This is blatant epistemological discrimination and moral principles everywhere won't stand for it! Infinite holism either supports global moral scepticism, or if limited by an epistemic filter such as the one Dancy describes, it supports scepticism about neither principles nor particular moral judgments.

In summary, we have seen that although some of the doctrines associated with moral particularism may be well grounded, the radical scepticism about the use of moral principles in moral thinking which has been advocated by some moral particularists is not

tenable. However, moral particularist arguments for scepticism about contributory principles may express an important moral theoretical discovery. This discovery was articulated in terms of the need for moral theories to capture greater complexity in the behaviour of moral reasons by developing default principles and reason modification principles in place of contributory principles.

From our consideration of such principles, I argued that we can understand moral principles (of whatever kind) as the moral generalizations required by our moral theories and supported by our particular judged contents. Since overall principles and contributory principles are also moral generalizations and, therefore, do not differ in their evidential relationship to particular moral judgments from other kinds of moral principles, it is arbitrary to assert that contributory principles cannot be inferred from particular judgments while allowing that other kinds of principles can be so inferred.

The particularist may respond to this arbitrariness by asserting an infinitely expansive holism, one which rules out all inferences from particulars to principles. However, using infinite holism to argue for general principle scepticism leads the particularist unavoidably to scepticism about particular judgments as well. The strategies which the particularist can use to avoid scepticism about particular judgments are equally effective in avoiding principle scepticism.

Therefore, thoroughgoing principle scepticism cannot be consistently supported by a moral particularist. Yet, neither can the particularist support selective scepticism about contributory principles because this requires drawing an arbitrary distinction between contributory principles and other kinds of moral generalization. Therefore, principle scepticism is false, and whatever truths there are in the particularist analyses of



cases such as Contraband can be understood as claims about what type of moral principles are supported by our moral evidence.

## **Chapter 5: Moral Evidence and Reflective Equilibrium**

In the third chapter I defended Case Evidentialism against challenges to the evidential use of moral judgments. In the fourth chapter I defended it against challenges to the evidential use of moral principles and mid-level generalizations. If these defences have been successful then we have made a strong case for accepting a method consistent with the bi-directional revision in our moral views that is advocated by Reflective Equilibrium. For this reason, I think that most philosophers who accept Reflective Equilibrium will find that Case Evidentialism provides a much needed epistemological framework to ground Reflective Equilibrium methodology.

However, there are two worries which adherents of Reflective Equilibrium might have with Case Evidentialism. First, one might argue that Case Evidentialism privileges particular moral judgments over moral principles and claim that such epistemological privilege is antithetical to Reflective Equilibrium. Second, one might invoke Rawls' distinction between narrow and wide reflective equilibrium and claim that Case Evidentialism's focus on the evidence of particular judgments and general principles only captures the methodological directives of Narrow Reflective Equilibrium. The more sophisticated method of Wide Reflective Equilibrium also requires that we balance both judgments and principles against background theories. In this concluding chapter, I will argue that neither of these worries are well founded.

## 5.1 Epistemological Privilege

Norman Daniels, in his now classic defence of Reflective Equilibrium, explicitly rejects the claim that particular moral judgments should be seen to "carry the epistemological privilege" in moral theorizing (Daniels, 1979, p. 264). Daniels thinks that in order to properly defend Reflective Equilibrium we must claim that no moral belief is epistemologically privileged, but that particular judgments, accepted principles, and background theories are all brought together on equal terms. Daniels seems to have two worries about assigning special epistemological status to particular moral judgments. First, he thinks that there is no compelling account of "why we should pay homage ultimately to those judgments and [only] indirectly to the principles that systematize them" (*ibid.*, p. 265). Second, he thinks that allowing special epistemological status to particular moral judgments implies that these judgments cannot be extensively revised in light of theory, whereas "[R]eflective [E]quilibrium does not merely systematize some determinate set of judgments" (*ibid.*, p. 266).

The best response to Daniels claim that Reflective Equilibrium should dispense with "fairy tales about epistemic priority" is to show that our best accounts of the nature of morality support an epistemology which grants special status to particular judgments (Daniels, 1979, p. 265). In Chapter 2, I argue that three of the best meta-ethical positions today imply that the contents of particular moral judgements are supported non-inferentially by our moral judgments. In contrast, moral principles can only be supported by inference from particular moral judgments or by inference from mid-level generalizations which are themselves supported by a chain of inferences terminating in particular moral judgments.

Each meta-ethical position gives a different account of why particular moral judgments are non-inferentially supported. According to observational realists, they are supported in the same way that empirical observations are supported. Wiggins' subjectivism holds that particular judgments are supported by our partial grasp of complex subjective predicates which is revealed to us through our application of moral predicates to particular cases. Finally, quasi-realism holds that particular judgments are supported because they express attitudes which we currently endorse but which we may stop endorsing in light of experience and moral argument. None of these meta-ethical accounts allow an assignment of non-inferential evidence to moral principles or other general moral beliefs. Therefore, contrary to Daniels' worries about fairy tales, we do have a good story to tell about why we should grant a special epistemological status to particular moral judgments. If the methodological tenets of Reflective Equilibrium can be preserved by an epistemological account that is backed by several plausible meta-ethical accounts, then the adherents of Reflective Equilibrium should gratefully accept a heterogeneous distribution of non-inferential evidence.

Daniels' second worry is a more serious one. He thinks that on any account which epistemologically privileges a subset of our moral beliefs, this privileged set must be held more or less fixed and unrevisable over the course of our moral theorizing. If Reflective Equilibrium is to offer us a method to develop moral theories which might provide genuine social criticism, it needs to be able to claim that most or even all of our pre-theoretical particular moral judgments could be abandoned in the course of theorizing. Can Case Evidentialism allow for such radical revision in light of theory, or does the fact

that it assigns non-inferential evidence only to particular moral judgments commit it to moral theoretic conservatism?

Several considerations speak against the conclusion that Case Evidentialism does not allow for radical revision of our pre-theoretical particular moral judgments. First, it holds that a particular moral judgment which is isolated from any theory whatsoever is on very shaky evidential ground. This means that in order to count as a plausible claim, a single judged content always requires inferential support from moral generalizations in addition to the non-inferential support it receives from its moral judging. Therefore, it is possible that only a small number of our judged contents are able to garner sufficient inferential support from moral principles to be confidently endorsed. If there is a major shift in the theoretical framework we accept as most evidentially supported, there could be a corresponding radical shift in the judged contents that we are able to confidently accept.

The second way in which Case Evidentialism can allow for radical shifts in our particular moral judgments in light of changes in our accepted principles is through the claim that our particular moral judgments are themselves partially theory dependent. We internalize moral generalizations both unreflectively through living with a set of cultural norms and also reflectively, by training ourselves in a well supported moral theoretical framework. These internalized moral generalizations affect how we conceptualize moral situations and, thereby, affect the particular moral judgments we make. If we internalize new moral generalizations which are better supported by evidence than the old internalizations, then the resulting changes in our particular moral judgments should be seen as overall improvements over the old judgments. Further, if there is a major change

to the moral generalizations we have internalized, it may bring about a radical change in our judged contents. Therefore, by allowing that particular moral judgments are strongly but not entirely dependent on internalized general moral beliefs, Case Evidentialism preserves the claim that particular judged contents are non-inferentially supported while also allowing for our set of accepted moral judgments to change radically as a result of moral theorizing.

It is difficult to offer a compelling example of a radical change in our moral judgments which has occurred as a result of moral theorizing because such radical changes tend to cover their own tracks. Once the new theory and set of judgments has been established, it is easy to forget the system of moral thought that had supported the old set of judgments. What would be needed to properly demonstrate the power of theory to transform our way of morally judging would be a careful philosophical analysis of the history of moral thought akin to Thomas Kuhn's work on the history of scientific revolutions (Kuhn, 1970).

Obviously, such a task cannot be undertaken here. However, I can point to several great transformations in the moral thinking of the West that have occurred in the 20<sup>th</sup> Century: the civil rights movement, feminism, and gay liberation. Although each of these changes in the West's moral thinking was brought about by political action rather than simple argument, the success of these political movements cannot be completely understood without recognizing the cogency of the moral arguments that these political movements presented against racism, patriarchy, and homophobia, respectively. How else could such disempowered groups have won political and legal support for their causes? It is my hope that a Case Evidentialist analysis of the moral arguments employed

during these historic shifts in moral thinking will offer a better understanding of the potential usefulness of moral theorizing in challenging entrenched cultural values.

Having responded to both of Daniels' objections to epistemological priority, I would like to point out a serious problem with Daniels' own view that Reflective Equilibrium should attribute a homogeneous epistemic status to all our moral claims. The problem is that in telling us to adjust moral principles sometimes and to adjust moral judgments other times, Reflective Equilibrium gives us no guidance about when to do the former and when to do the latter. If we are just able to revise whichever side strikes our fancy, with no further epistemological constraints, then Reflective Equilibrium would allow for somebody to systematically reinforce a prejudiced moral position simply by always choosing to revise principles or judgments which disagree with their prejudices.<sup>10</sup>

Case Evidentialism avoids this problem by providing clear guidance in cases of conflict between a general moral claim and a particular moral judgment. If the particular judgment together with other judgments supports an alternative moral generalization, then the conflicting general moral claim should be revised. On the other hand, if the general moral claim is well supported by other particular moral claims, and the conflicting particular judgment is not inferentially well supported by other moral generalizations, then the particular judgment is to be revised. This ability to give methodological guidance in cases of conflict between judgments and principles makes Case Evidentialism an essential addition to Reflective Equilibrium methodology. Without such guidance, Reflective Equilibrium is subject to the charge that it extremely

---

<sup>10</sup> This line of argument against reflective equilibrium has been developed in Haslett (1987). It is argued in Ebertz (1993) that reflective equilibrium might avoid this problem by endorsing a modest foundationalist epistemology that assigns defeasible direct justification to some of our moral claims.

underdetermines the moral claims which are to be held in equilibrium, even when there is initial agreement about judgments and principles.

The claim that advocates of Reflective Equilibrium should deny special epistemological status to particular moral judgments can now be seen to be unfounded. First, as we saw in Chapter 2, there are good meta-ethical reasons to assign non-inferential evidence to particular moral judgments and not to general moral statements. Second, giving particular moral judgments this sort of special status does not imply that they are unrevisable or mostly unrevisable. Finally, a homogeneous approach to moral evidence leaves us with no guidelines for making adjustment decisions when particular judgments and general principles conflict. By allowing that particular judged contents are epistemologically special in being non-inferentially supported, Reflective Equilibrium gains meta-ethical support and avoids vicious underdetermination, while maintaining its capacity for radical moral criticism.

## **5.2 Wide Reflective Equilibrium**

Several years after proposing Reflective Equilibrium in *A Theory of Justice*, Rawls drew a distinction between Narrow Reflective Equilibrium and Wide Reflective Equilibrium (Rawls, 1974-5). In seeking a state of narrow reflective equilibrium, we bring our moral principles and moral judgments into harmony by sometimes modifying principles, sometimes modifying judgments, until we achieve maximum coherence between the two. In contrast, achieving a wide reflective equilibrium involves balancing not only our moral judgments and moral principles, but also balancing both of these against background theories. As examples of such background theories, Daniels gives "a



theory of the person, a theory of procedural justice, general social theory, and a theory of the role of morality in society" (Daniels, 1979, p. 260). Using such background theories ensures that our moral principles are better grounded in a broad understanding of humanity and society.

The focus of my discussion of Case Evidentialism has been on the evidential relationship between particular moral judgments and moral principles. This may give the impression that Case Evidentialism is committed to a version of Narrow Reflective Equilibrium, and so to the irrelevance of background theories in moral inquiry. In this section, I show that Case Evidentialism can offer an account of the role that background theories play in moral inquiry.

In her discussion of the relative merits of Wide and Narrow Reflective Equilibrium, Margaret Holmgren notes that the concept of a background theory used by Rawls and Daniels is not sufficiently clear. "Some background theories – for example, a theory of procedural justice – appear to be straightforward moral theories.... Other background theories appear to be straightforward nonmoral theories" such as general social theories like economics and sociology (Holmgren, 1989, p. 50). Daniels seems happy with this ambiguity, claiming only that at least some of our background theories are themselves supported by moral judgments, but allowing that some may be non-moral theories. I think it would be helpful to deal separately with the relevance of non-moral background theories to moral inquiry, before turning to examine the relevance of background theories with an irreducibly moral component.

Case Evidentialism accounts for the importance of non-moral theories in moral debate in two ways. First, particular cases are characterized for us through non-moral

facts. We come to understand the non-moral facts of a situation, and it is about these facts that we make our moral judgment.<sup>11</sup> An improvement in our non-moral theories may bring features of a case to our awareness which we previously missed or show that certain features we thought were present are actually absent. The better our non-moral understanding of the case, the more likely we are to be cognizant of all and only the morally relevant features which are truly present in the case.

The second way in which Case Evidentialism can make sense of the importance of non-moral background theories to moral theorizing is that non-moral background theories play a crucial role in the evaluation of moral evidence. In order to tell which moral generalizations are supported by particular moral judgments, we must be able to see which non-moral features different cases hold in common. For example, in order to evaluate the evidence in favour of the generalization, 'pain is always a reason against,' we must be able to identify which cases are ones which involve pain and then what our judgments about these cases tell us. For example, in order to evaluate whether the case of a consensual sadomasochistic relationship conflicts with the generalization 'pain is always a reason against' we need a non-moral theory to tell us whether the sadomasochistic relationships in question really involve inflicting pain.

For example, the neurological theory that pain is caused by activity in a region of the brain, would enable us to determine through empirical study whether the masochist feels pain during sadomasochistic activity. If we determine that the masochist does feel pain, we may judge that this pain is morally neutral. This judgment would then count as

---

<sup>11</sup> Here I am only claiming that some of our particular moral judgments must be understood as judgments about non-moral facts. We may also make higher order particular moral judgments which are about combinations of non-moral facts and particular moral claims. For example, an overall moral judgment may be about both non-moral facts and contributory moral claims.

evidence against the generalization 'pain is always a reason against' and vice versa. This judgment would also provide support for the reason modification principle of informed consent from Chapter 4, because this principle can explain the neutrality of pain in sadomasochistic relationships as well as the neutrality of lying in the Contraband case. On the other hand, if it were discovered through brain research that the masochist does not feel pain during consensual sadomasochistic sex, then our judgment about the case could not be used as evidence against the generalization that pain is always a reason against an action.

This understanding of the role that non-moral theories play in moral inquiry is at odds with the view sometimes implied by moral philosophers that we can directly infer the plausibility of a moral principle from a non-moral background theory. For example, Derek Parfit claims that his theory of personal identity supports utilitarianism by showing us that the divisions between one person and the next are "less deep" than common sense tells us (Parfit, 1973 as cited in Holmgren, 1989). However, Holmgren criticizes this direct inference from the deepness of a non-moral fact to a moral conclusion. She argues, "nonmoral facts, whether deep or shallow, have no *independent* moral significance. (E.g. the fact that there are four basic forces in the universe...has no independent moral significance, although it seems a good candidate for a deep fact.)" (Holmgren, 1989, p. 56). Our non-moral understanding of the world only has relevance for moral theory after it has passed through our moral judgment.

Nevertheless, Parfit's claim that a theory of personal identity might support a moral theory can be understood in a way that does not try to infer moral claims from non-moral facts. A Parfitian could argue instead that the moral judgments most people make

about the cases which are put forward as counter-examples against utilitarianism are contaminated by the faulty commonsense understanding of personal identity. They would claim that once we have properly internalized Parfit's superior account of personal identity, then the cases which were previously judged to be at odds with utilitarianism would be judged to be unproblematic.

I now move on to discuss moral background theories. Moral background theories can be distinguished from our theory of moral principles in the following way. Moral principles are general statements about the moral status of actions. For this reason they are the general moral claims that are most closely related to practical decision making. What I have been calling 'moral theory' is the attempt to find the best set of such action guiding principles. Therefore, what I have been calling moral theory might be more properly called the theory of moral action. However, I will stay with common usage and say that a moral theory is a collection of moral principles which tell us about the moral status of actions. In contrast, moral background theories are theories which do not tell us directly about the moral status of actions, but which nevertheless rely, in part, on the evidence of moral judgments.

For example, a value theory is a *moral* background theory because it is supported by our particular judgments about what is morally good and bad. Value theories can be developed from particular value judgments in the same way that moral principles are developed from particular overall and contributory judgments. However, value theory is a moral *background* theory rather than a moral theory, because we cannot infer the moral status of actions from a knowledge of their good and bad consequences alone. In addition, we need a moral principle such as 'the action which does the most good is

morally obligatory.' Other moral background theories include a theory of moral virtue, a theory of evil, a theory of justice, a theory of the moral status of persons, etc. Some of these background theories may involve conceptual analysis and/or empirical claims as well as moral claims.

Our moral background theories do not determine our moral principles because they do not speak directly to the moral status of actions. However, moral background theories do help us develop moral principles because they enable us to formulate principles which assign a moral status to actions based on a moral description of the action rather than a non-moral description of it. For example, the consequentialist principle mentioned earlier, 'the action which does the most good is morally obligatory,' can only be formulated and tested in light of a theory of moral goodness. Similarly, the generalization that cruelty is always a moral reason against, can only be tested against particular judgments given a moral understanding of which actions are cruel. Such an understanding might be improved by developing a moral background theory about cruelty.

From these examples we can see that the role of moral background theories in evaluating moral principles is very similar to the role of non-moral background theories in evaluating them. In both cases, our judgments about particular actions are informed by our understanding of the background considerations. Second, our ability to evaluate the evidence provided by moral judgments for moral principles depends on our background understanding of the non-moral and moral action descriptions which the principles employ. In this way, moral and non-moral background theories function very similarly to auxiliary assumptions in science. In order to test a hypothesis through an experiment,

scientists must make many assumptions that cannot be tested by that experiment. These assumptions are grounded in other theories, which are not currently under investigation. For example, an archaeologist uses carbon dating along with other indicators to determine the age of a burial mound. However, in doing so, the archaeologist must presuppose the theories on which the carbon dating testing is based.

However, non-moral background theories and moral background theories are not related in quite the same way to moral theory. The difference comes when we consider the claim of Wide Reflective Equilibrium, that our background theories themselves can be revised in light of our moral principles and moral judgments. To be plausible, this claim must be restricted to moral background theories because we do not have epistemic reasons to revise a non-moral theory from moral considerations (although, of course, we might have moral reasons to do so).

Sometimes, scientific theories make claims that do not sit well with our moral commitments but we should not discount the plausibility of non-moral claims on moral grounds. For example, suppose that a scientist claims to discover that a certain ethnic group is genetically predisposed to low intelligence. Such a finding is morally problematic because it could encourage discrimination against the ethnic group in question. However, the rational response to such a problem is not to claim that we should ignore or discount scientific evidence in order to avoid moral wrongdoing. Neither need we be complacent in the face of those who try to use the scientific finding to reach moral conclusions we find unacceptable. We may argue that the moral conclusions of those who support discrimination are not supported by the scientific finding. Alternatively, we could

argue that intelligence is not a morally neutral concept, or that the researcher's methods for measuring intelligence were biased by discriminatory assumptions.

In contrast, we can legitimately revise our *moral* background theories in light of moral judgments and well supported moral generalizations. For example, the principle 'justice is always a moral reason in favour of an action' seems to be very well supported by our particular moral judgments about justice and about just acts. If we encounter an act with certain non-moral features that make the act just according to our background theory of justice, but these same features are also judged to be reasons against the action, then our evidence may count more strongly against the claim that the act is just, than against the principle that 'justice is always a moral reason in favour.' In other words, we modify our background theory of justice to bring it into coherence with our well supported moral contributory principle that justice is always a reason in favour of an action.

I have tried to show that Case Evidentialism can make sense of Wide Reflective Equilibrium's central claim that background theories must be taken into account in moral theorizing. The conception of the role of background theories which I have offered is more complex than Daniels' injunction to balance judgments, principles, and background theories. However, this complexity should be viewed as a virtue of Case Evidentialism because it yields more specific methodological guidelines which avoid the problem of vicious underdetermination that arises with more epistemologically nebulous accounts of Reflective Equilibrium.

## Conclusion

Our attempt to find a meta-ethically grounded epistemological basis for the key tenets of Reflective Equilibrium has given us a complex picture of moral theorizing. I have attempted to show how many different kinds of judgments and many different kinds of theories function in moral inquiry according to familiar rules of evidence. When one is engaged in actual moral debate, the complexity of this evidence sometimes gives one the feeling of futility. It sometimes feels as if all moral theorists have to work with are fleeting intuitions, intuitions which seem to be hopelessly biased or intractably at odds with one another.

In this work, I attempted to penetrate the evidential complexity of real moral theorizing by working from relatively simple forms of moral theoretical argument, towards more complex forms. In Chapter's 2 and 3, I focused exclusively on the evidential support between overall moral judgments and overall moral principles. In doing so I may have given the impression of endorsing an overly simplistic conception of moral theorizing.

In the fourth chapter, I expanded my account to consider the additional theoretical complexities introduced by moral reasons. Moral reasons allow us to make sense of much more complicated moral arguments which involve contributory judgments and principles. Yet, even contributory principles may not be sufficient to capture the full complexity of moral reasons, so I developed the idea of default principles and reason modification principles.



Finally, in Chapter 5, I attempted to extend my account of moral evidence to include all legitimate forms of moral argument by accounting for the role of background theories in the evaluation of moral principles. Necessarily, this final discussion was more programmatic as the scope of potentially relevant background theories seems limitless. Therefore, I attempted only to show the general strategy by which Case Evidentialism would make sense of the role of background theories. Moral theorists themselves will have to work out the details of the evidential relationships between, for example, a theory of justice and a theory of moral action.

By working from simple forms of moral theorizing towards an account of moral theorizing in all its messy reality, I hope to have shown that the messy reality should not lead us towards epistemological despair. The difficulties in contemporary moral theory are real, but they need not be seen as arising from an inherent flaw or limitation in the approach which has become popular since the advent of Reflective Equilibrium. Rather, these difficulties should be understood as arising from the magnitude of the task with which moral theory faces us. Given our numerous moral judgments, the need for these judgments to be inferentially supported, the vast number and kinds of moral principles which can be proposed, and the relevance of non-moral and moral background theories, it should come as no surprise that finding a consensus in moral theory is a difficult task. We must also remember that contemporary moral theory is a young discipline, and one which is still beset by the growing pains of methodological disagreement. If the start of contemporary moral theory is marked by the publication of Rawls' *Theory of Justice*, then moral theorists should commend themselves for the progress they have made in so short a time, rather than despairing at the continuing disagreement within the field. If an

evidential account of moral theorizing like the one I have advocated here is correct, then we have every reason to believe that progress in moral theory will continue steadily.

## References

- Blackburn, S. (1993). *Essays in Quasi-Realism*. Oxford: Oxford University Press.
- Blackburn, S. (1999). Is Objective Moral Justification Possible on a Quasi-realist Foundation? *Inquiry*, 42, 213-228.
- Blackburn, S. (2002). Précis of Ruling Passions. *Philosophy and Phenomenological Research* 65(1), 122-135.
- Crisp, R. (2000). Particularizing Particularism. In B. Hooker & M. O. Little (Eds.), *Moral Particularism*. Oxford: Oxford University Press.
- Dancy, J. (1993). *Moral Reasons*. Oxford: Blackwell.
- Dancy, J. (2000). The Particularist's Progress. In B. Hooker & M. O. Little (Eds.), *Moral Particularism*. Oxford: Oxford University Press.
- Dancy, J. (2004). *Ethics Without Principles*. Oxford: Oxford University Press.
- Daniels, N. (1979). Wide Reflective Equilibrium and Theory Acceptance in Ethics. *The Journal of Philosophy*, 76(5), 256-282.
- Daniels, N. (1996). Wide Reflective Equilibrium in Practice. In N. Daniels. *Justice and Justification*. Cambridge: Cambridge University Press.
- Ebertz, R. (1993). Is Reflective Equilibrium a Coherentist Model? *The Canadian Journal of Philosophy*, 23(2), 193-214.
- Haack, S. (1993). *Evidence and Inquiry: Towards Reconstruction in Epistemology*. Oxford: Blackwell.
- Haack, S. (2003). *Defending Science Within Reason*. Amherst, NY: Prometheus Books.
- Hanson, N. (1958). *Patterns of Discovery: An Inquiry into the Conceptual Foundations*

- of Science*. Cambridge: Cambridge University Press
- Hare, R. (1973). Rawls' Theory of Justice. *Philosophical Quarterly*, 23(91), 144-155.
- Hare, R. (1976). Ethical Theory and Utilitarianism. In H. D. Lewis (Ed.), *Contemporary British Philosophy IV*. London: Allen & Unwin.
- Hare, R. (1981). *Moral Thinking*. Oxford: Oxford University Press.
- Hare, R. (1997). *Sorting Out Ethics*. Oxford: Oxford University Press.
- Harman, G. (1977). *The nature of morality : an introduction to ethics*. New York: Oxford University Press.
- Harman, G. (1999). *Reasoning, meaning, and mind*. Oxford: Oxford University Press.
- Haslett, D. W. (1987). What is Wrong with Reflective Equilibria. *Philosophical Quarterly*, 37(148), 305-311.
- Holmgren, M. (1989). The Wide and the Narrow of Reflective Equilibrium. *Canadian Journal of Philosophy*, 19(1), 43-60.
- Korsgaard, C. (1996). *Sources of Normativity*. Cambridge: Cambridge University Press.
- Kuhn, T. (1970). *The Structure of Scientific Revolutions* (2<sup>nd</sup> ed.). Chicago: University of Chicago Press.
- McNaughton, D. (1988). *Moral Vision*. Oxford: Basil Blackwell.
- Monro, D. H. (1967). *Empiricism and Ethics*. Cambridge: Cambridge University Press.
- Parfit, D. (1973). Later Selves and Moral Principles. In A. Montefiore (Ed.). *Philosophy and Personal Relations*. London: Routledge & Kegan Paul.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge, Mass: Harvard University Press.
- Rawls, J. (1974-5). The Independence of Moral Theory. *Proceedings and Addresses of the American Philosophical Association* 48, 5-22.

Ross, D. (1930). *The Right and the Good*. Oxford: Oxford University Press.

Stratton-Lake, P. (2002). Introduction. In P. Stratton-Lake (Ed.) *The Right and the Good*.  
Oxford: Oxford University Press.

Wiggins, D. (1987). A Sensible Subjectivism. In D. Wiggins *Needs, Values, Truth*.  
Oxford: Blackwell.

Wilson, C. (2004). *Moral Animals: Ideals and Constraints in Moral Theory*. Oxford:  
Oxford University Press.