

**THE METAPHYSICS OF MORAL JUDGEMENTS OF
OBLIGATION**

by

MARK PATRICK SHEEHAN

A dissertation submitted to the Graduate Faculty in Philosophy
in partial fulfilment of the requirements for the degree of
Doctor of Philosophy, The City University of New York

2004

UMI Number: 3144138

Copyright 2004 by
Sheehan, Mark Patrick

All rights reserved.

INFORMATION TO USERS

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleed-through, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

UMI[®]

UMI Microform 3144138

Copyright 2004 by ProQuest Information and Learning Company.

All rights reserved. This microform edition is protected against
unauthorized copying under Title 17, United States Code.


ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346


2004

MARK PATRICK SHEEHAN

All Rights Reserved

This manuscript has been read and accepted for the Graduate Faculty in Philosophy in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

June 2, 2004 
Date Chair of Examining Committee

June 2, 2004 
Date Executive Officer

Professor Stefan Baumrin

Professor Jonathan Adler

Professor Simon Blackburn

Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

THE METAPHYSICS OF MORAL JUDGEMENTS OF OBLIGATION

by

Mark Patrick Sheehan

Adviser: Professor Stefan Baumrin

This thesis is an essay in philosophical moral psychology. It seeks to develop an account of the nature of moral judgements of obligation, where 'nature' is understood as the necessary preconditions for these judgements and where a judgement is taken to be a psychological act — the act of judging.

This aim is achieved by beginning with a key feature of our experience of moral judgements of obligation. This feature can be briefly characterised as a pressure to act which seems to originate independently of the agent. The two most important aspects of this experience are its capacity to motivate and its 'external' or 'objective' character. The argument then proceeds through a series of regressive steps. At each step we reach a conclusion that helps to unpack the nature of these judgements.

The account that is developed at the end of this series of argumentative steps makes claims about the structure of an agent's 'frame of mind'. That is, it provides a structural account of the configuration of the agent's mental states. The key feature of this structure is a set of desires (or more generally, attitudes) towards ways of representing the world. The simplest examples of desires of this kind are desires to notice.

The account developed in the thesis is strictly neither a cognitivist nor a non-cognitivist one. There are two reasons for this. First, the thesis is not concerned with whether moral judgements of obligation are capable of having truth-values but with the psychological structures involved. Second, the general structural account of the agent's 'frame of mind' given in the thesis does not privilege either belief-like or desire-like states but interlocks them in the key feature of the account.

However, the key structural element of the frame of mind is a type of attitude or desire-like state in quite non-cognitivist-friendly way. The arguments presented here and the account given of the nature of moral judgements of obligation are not intended to settle the debate between cognitivist non-cognitivist. Instead, it goes some way towards closing off one part of the dispute.

ACKNOWLEDGMENTS

There are very many people who have helped to make this dissertation possible and who have shaped its form and content. I am deeply grateful to my friends and colleagues in New York — at the Graduate Centre and at City College, and, more recently, at Keele. Without each of them this dissertation would be poorer.

Stefan Baumrin, Brad Armour-Garb, my father, Peter Sheehan, and Martin Harvey have undoubtedly had the greatest effect on the completion and shape of this dissertation. Stefan continually prodded and pushed what I had written both to make it more precise and to enable me to make broader historical connections. Brad and I thrashed out almost all of the key points in hours and hours of conversation. His relentless pursuit of the arguments has shaped the dissertation in profound ways. Time and again, Dad's experience and insight helped me to get over the psychological and structural hurdles that dissertation writing presents. His clarity and breadth of perspective invariably got me to see past and over whatever happened to be the current problem. Marty's laconic style and gentle argument always helped the issues and the beers go down more easily.

I dedicate this to my wife Debbie and to my mother and father. Deb and Mum, thanks for being patient when you were, for being impatient when you weren't and for supporting me always. Dad, thanks for giving me philosophy and for giving me such an example of how to use it to be a better person. I hope that this represents a small step in a right direction.

TABLE OF CONTENTS.

	Page
Chapter 1: Preliminaries.	1
A. Goals, Limits and Method.	
B. The Experiential Feature of Judgements of Obligation.	
Chapter 2: Obligation and the Origin of the Pressure.	39
A. The Origin of the Pressure.	
B. Obligation and Ignorance of Fact.	
C. 'Ought' and Judgements of Obligation.	
Chapter 3: The Particular and the General.	73
A. Explaining and Agent's 'take'.	
B. Rules.	
C. Aristotle's Minor Premise.	
Chapter 4: Explaining the Experience (Part I).	107
A. Aristotle's Practical Syllogism.	
B. Explaining Attention.	
C. Cognitivism and Sensibility.	
D. Two Points of Contention.	
Chapter 5: Explaining the Experience (Part II).	164
A. Shyness and Sensitivity.	
B. Non-Cognitivism and Projection.	
C. A Return to 'the Pressure'.	
Chapter 6: The Desire to Notice.	206
A. An Initial Formulation.	
B. Extensions.	
Chapter 7: Cognitivism and an Agent's Concerns.	255
A. The Argument in 'Virtue and Reason'.	
B. Three Issues.	
Chapter 8: The Frame of Mind.	271
Bibliography of works cited.	282

CHAPTER 1: PRELIMINARIES.

“There are two ways in which the subject of morals may be treated. One begins from inquiring into the abstract relations of things: the other from a matter of fact, namely, what the particular nature of man is, its several parts, their economy or constitution; ...”
(Bishop Butler, Preface, *Fifteen Sermons*, §12)

In this thesis I will develop and argue for an account of the nature of moral judgements of obligation. In doing this I intend to give an account of the conditions for the possibility of these judgements.¹ I will not seek to catalogue and classify the judgements of obligation that we make but to find and explain features that those judgements share and which make them judgements of that sort. This is a descriptive project — it is essentially an endeavour to give a particular kind of explanation of a particular kind of judgement that, it is assumed, we ordinarily make. It will be my contention that all moral judgements of obligation have the components that the account specifies and that these components explain our making of these judgements, again, because these components are the conditions under which a judgement is a moral judgement of obligation.

This is an exercise in philosophical moral psychology. It is psychological in part because it deals with the act of judging rather than the content of what is judged. It is

¹ There are three broad senses of ‘nature’ that might be intended given an investigation into moral judgements: necessary and sufficient conditions for moral judgement, causal/empirical conditions of moral judgement and conditions for the possibility of moral judgement. I am concerned with the third of these senses.

philosophical rather than empirical because it seeks 'conditions for the possibility of' rather than causal underpinnings of these judgements.

Having said this what follows is very much about giving a particular kind of psychological explanation. By 'psychological explanation' here I do not intend anything too controversial. In our ordinary practice of explaining and understanding our own behaviour and the behaviour of others, we use the full resources of language to attribute an entire range of mental states to the agent in question. How precisely this goes and whether it is, in the end, acceptable and consistent with science is not of concern here. Ordinarily, the juxtaposition of these various states is taken to explain what people do and say by reconstructing, in the terms and according to the rules of the practice, the (perhaps partial) configuration of the agent's mental states which combined (or would combine) to issue in the behaviour being explained. In all of this the primary evidence on which the explanation is constructed is what the agent says or does — expressions or reports of psychological states such as beliefs, desires or emotions, observed behaviour in particular contexts and reasons and justifications given in support for judgements made and acted upon are all used to develop the kind of psychological explanation of interest here.

In these terms, the end point of the project will involve reference to a particular configuration of an agent's mental states — the agent's 'frame of mind'. This configuration will be largely schematic — not so much giving the details of particular mental states (it is hardly likely that this will be possible for all moral judgements of obligation) but providing a structural account of the relation between the various kinds of

states. It is this structural configuration that will be the end point of the project and, it will be argued, amount to the (necessary) preconditions for moral judgements of obligation.

This first chapter has two objectives. First, in section A, the project will be outlined. This will include a discussion of the objects of study, the form that the account will take and the method that will be used. It will also outline the structure and argument of the thesis as a whole. Section B of this chapter involves a discussion of an experiential feature of moral judgements of obligation. As will become clear, this is the crucial first step in the development of the argument of the thesis.

A. Goals, Limits and Method.

Rather straightforwardly, there are three issues to be addressed in what follows:

- (1) What are the objects of study?
- (2) What method will be used to achieve the stated goals?
- (3) What is the structure of the study and its arguments?

This section is divided into three parts corresponding to each of these issues.

1. The Objects of Study.

At the heart of this project is a distinction between two questions that we may ask when doing moral philosophy: 'What is the nature of moral judgement?' and 'Which moral judgements are right and which are wrong?'. The first of these questions presumes that we make, or at least think that we make, judgements of a particular sort. It is thus, essentially descriptive. The second question takes moral judgements as its object and

seeks to apply a standard to them, to distinguish those that are right from those that are wrong. My concern is limited to the first of these questions.² This stipulation gives rise to what we might call, the 'non-normative requirement'. It is beyond the scope of the thesis to say which of judgements are right, preferable, justifiable or correct — the account is in this sense non-normative.³

The focus here is more specific than moral judgement generally. We are concerned primarily with moral judgements of obligation. The thought is that moral judgement in general can be divided into judgements of obligation and judgements of value. This is a relatively standard distinction and is sometimes construed as a distinction between practical and evaluative moral judgements.⁴ By and large this distinction is

² Such a concern is warranted by the fact that moral judgements themselves form a large part of subject matter of moral philosophy and as such, an investigation into their nature will be or, at least, could be instructive. It may be claimed that such an investigation ought to be conducted before normative questions can properly be broached. I will have something to say about some normative questions insofar as they are of importance in explaining the account, however I will not maintain or argue that an answer to the former question is required before we can address the latter. It seems reasonable to think that the question 'What is the nature of moral judgement?' is an interesting philosophical question independently of whether an answer to it is needed in order to answer the question 'Which moral judgements are right and which are wrong?'. Given this, an account that explains what it is to make a moral judgement, constituting an answer to the former question, is also of interest.

³ While the account itself is intended to be as non-normative as possible, what it is an account of is essentially normative. Indeed it is hard to imagine that any account of the nature of moral judgement could fail to contribute to an explanation of the normativity that these judgements possess.

⁴ For instance: Frankena, W.K. *Ethics*. Englewood Cliffs, NJ: Prentice hall, 1963: 9; Wiggins, David. 'Truth, Invention and the Meaning of Life' in *Needs, Values, Truth*. Oxford: Clarendon Press, 1998: 95; and, Darwall, S, Gibbard, A and Railton, P. 'Toward *Fin de siècle* Ethics' *The Philosophical Review*, Vol.101, no.1 (January, 1992): 115-189. Reprinted in *Moral Discourse and Practice: Some Philosophical Approaches*. New York: Oxford University Press, 1996. (Further references to this article will refer to the reprinted version.)

between judgements involving what ought to be done (by myself or others) and what is good, preferable, and so on (evaluative judgements tend to involve the predication of a value term to some object or event). In spite of this distinction it is natural to suspect that an account of judgements of obligation will, at some point, run into an agent's values. A good deal more will be said both about this distinction and about the connection between judgements of obligation and value.⁵

There are a number of ways of understanding 'moral judgement'. First, in what follows I will understand 'moral' in the sense that is to be distinguished from 'non-moral'. 'Moral judgements of obligation' thus picks out the set of judgements of obligation that involve 'moral considerations' as opposed to those that do not.

Second, I will assume that all judgements are judgements that such and such and I will construe what counts as a judgement quite broadly. In addition to inferring and concluding, say, I include such things as seeing that such and such, discerning, recognising, interpreting, appreciating, realizing and becoming aware. I will take it that judgement can be expressed, thought or reflected in action. There are two senses of judgement: what is judged (or the content of the judgement) and the act of judging. In the former sense, 'judgement' refers to the content of the judgement — on the content of particular thoughts or utterances involving obligation: that I ought to assist this confused student or that one should give to charity. In the latter sense, the judgement is a particular kind of act, a performance or the exercise of an ability. I will mean this latter sense of judgement, the act of judging, when I speak of a moral judgement. Thus we are concerned with my *judging* itself rather than what it is *that I judge*.

⁵ See particularly section B of this chapter and section C of chapter two.

It is important to isolate those ‘occurrences’ that are the focus of the investigation — moral judgements of obligation. We must be careful to ensure (i) that all judgements of this kind are included and (ii) that the non-normative requirement is satisfied. The most pressing danger here is the way in which we draw the distinction between what is moral and what is non-moral⁶. The trouble is that if we restrict the *content* of moral judgements by saying what constitutes a ‘moral matter’, we would be limiting what someone could make a moral judgement about. But to dismiss someone’s purportedly moral judgement as not being concerned, really, with any moral matter would be to judge that they were wrong in so judging. This would be to make a normative appraisal and for this project, something to be avoided.

This is an important point for what follows. A few examples are in order. We can readily imagine someone for whom matters of etiquette take on much more significance; for whom what is ‘proper’ and ‘improper’ is not distinguished from moral rightness and wrongness. For such a person, those judgements whose content involves matters of ‘mere’ etiquette are moral judgements. If we suggested that issues of etiquette are not a moral matter, this person might respond by citing our claim as further evidence of the moral decay of modern society. We may judge that this person is wrong to confuse what is improper with what is morally wrong. But however strongly we disagree with or deride this set of values; it remains a set of moral values. The more we become involved in such disagreements, the more we become embroiled in normative issues.⁷ On the other hand, consider someone who thinks that morality only involves how one treats members of

⁶ Cf. Blackburn, S. *Ruling Passions*. Oxford: Clarendon Press, 1998: 1-13.

⁷ Crucially while these normative issues will remain non-moral for us they will conversely remain moral for our opponent.

their own group, be it sex, race, religion or community. This 'thin' conception of morality means that judgements that concern people of other groups — how one should treat other races or the opposite sex, say — are not moral judgements. In short, saying what counts as a moral matter is, very often at least, a moral matter.⁸

We need then to ensure that the nature of moral judgements is not presupposed and that their contents are not limited by our own conception of what is a moral matter. Failure on either of these fronts will render the account trivial or simply the reflection of our antecedently held moral convictions.⁹

The search for a way to isolate the set of judgements under consideration has run into problems. Perhaps the best we can do is the following:

For some agent A, the content of a judgement made by A is moral when A thinks or says (or would think or say) that the judgement was a moral one.

A, presumably, will think that his judgement was a moral one when he thinks that the content of his judging was moral — that is, when he had a moral issue to judge about. To say any more than this would be to give content to A's distinction between moral and non-moral issues. Of course I may disagree with A on his assessment of his judgement. This would mean that I disagree with A's understanding of what constitutes a moral

⁸ In order to distinguish moral from non-moral judgements we might alternatively try limiting the *acts of judging* that are to be counted as moral judgements but this would require an understanding of their nature — the very investigation in which we are engaged.

⁹ Maurice Mandelbaum in *The Phenomenology of Moral Experience* (Baltimore: Johns Hopkins Press, 1955) writes: "Unless it can be demonstrated that the moral judgements found among people of other cultures, among socially ostracized individuals, and among those whose standards we abhor, are essentially different in character, and not merely in content, from the moral judgements we ourselves make, we have no right to exclude them from consideration." (p.34) Unless otherwise stated all references to Mandelbaum's work will be to this book.

judgement. Again, I might disagree with A about what is or is not a 'moral matter' or I might deny that the kind of 'judging' involved in this case was of the moral kind. Within this investigation, I have no grounds for thinking that my understanding of what constitutes a moral content is to be preferred — this would be a normative claim. Similarly, even if I had access to A's actual judgement, I could not, *at this point in the investigation*, have grounds for deciding that it was or was not a judgement of the relevant kind.

I do not intend that this way of isolating moral judgements is in any way substantive. Indeed the criterion is fairly unhelpful. At this stage, the 'criterion' is meant to give a general sense of the objects of the study. It will include all moral judgements but it may, when we know more about the nature of these judgements, include some non-moral judgements. I will have a little more to say in this regard at the beginning of the next section.

2. Method.

There are two things to decide upon in considering the appropriate method to use in a project such as this: where to begin and how to proceed. Before deciding upon these however we should say a little more about what is involved in judgements more generally. I intend this picture to be fairly minimal and certainly uncontroversial. In short, a moral judgement will have an expressed or expressible content, will be made in a context and will be accompanied or preceded by particular psychological features.

The judgement itself is presumably a psychological phenomenon and can be understood as the ‘culmination’¹⁰ of, what might be called, initial psychological states. So the psychological component of the judgement can be separated into two: the judging itself and the initial psychological states. The latter are those psychological states that combine or are brought together in some way in the making of a moral judgement. The end point of the project will involve crucial reference to precisely these states. In the terms introduced at the outset — it is structure of the particular configuration of the agent’s mental states, the ‘frame of mind’, that will provide us with the necessary precondition for moral judgements of obligation.

The former, the judging itself, is precisely what is to be explained in the account. Our experience of such judging is of some importance here. One of the reasons for our interest in the act of judging might be to better understand and explain our experience of those judgements. A satisfactory account of moral judgement will be one that does justice to our experience of such judgements. Any story that departs from this experience must provide strong arguments or explanations for such a departure. Even still, such departures can only be, as it were, the reluctant results of thorough investigation. Our experience of the moral judgements we make and the content of those judgements are the grounds against which an account can be measured. However when the focus of the account is on the act of judging rather than directly on the content, as it is here, its convergence with this experience is of great importance. It may be objected that we do not need to worry about the experience of moral judgement since what we need is a model that is psychologically plausible — one that we *could*, rather than do in fact use. This misses the

¹⁰ Whether this ‘culmination’ constitutes a ‘concluding’ or a perception, etc. will depend on the judgement in question.

point just made by failing to stipulate what the account models. Were the account an attempt to provide a model of the content of our moral judgements, convergence with our actual experience would be of secondary importance. Rightly, such a model would need to converge with the contents of our moral judgements. Since this is to be an account of our moral 'judgings', it will need to converge with our experience of those 'judgings'.

There are three, more obvious, possible starting points¹¹: an analysis of moral language, examination of initial psychological states and an investigation of the experience of moral judgements of obligation. All three of these candidates bring with them a favoured strategy.

In the first case, our inquiry would, at the outset, be concerned with conceptual analysis. We would gather together behavioural and linguistic evidence — reports of moral judgements correlated with situations; accounts of situations, reasons for action, etc. We would then examine our moral language in an effort to find commonalities and patterns and in an effort to develop an account of our moral concepts. By developing an understanding of concepts like 'obligation' we get at the nature of obligation and so, eventually, the nature of judgements of obligation.

The second candidate is an examination of the initial psychological states that 'culminate' in the moral judgement. Whatever is common to all moral judgements of obligation will be revealed through what is common to the psychological states prior to the judgement. Naturally, beginning with the psychological preconditions would incline

¹¹ If we count the initial psychological states and the experience as separate components then there are five components of moral judgements to be considered. The two that are immediately ruled out are the context and the content. Both seem equally hopeless as starting points simply because of the massive diversity in what we judge and in the situations in which we do so.

us to use a psychological method. The psychological method appropriate with the initial psychological states as a starting point takes its leave from empirical psychology. We would be interested primarily in uncovering the causal relationships that underpin our moral judgements. Perhaps more specifically we would wish to discover which components of our psychological makeup, our psychological natures, are responsible for our moral judgements and in what way. Proceeding in this manner, we would have at our disposal the full complement of psychological studies along with theoretical approaches taken by both psychologists and cognitive scientists. In all this, we would rely most heavily on behavioural evidence (both real and hypothetical) — including expressions/reports of psychological states such as beliefs, desires and emotions as well as expressions of moral judgements and reasons both for action and judgement. Quite clearly, all of these could be brought to bear on moral judgement in a way that has all the non-normativity and generalisation that we require.

Given the aims of the project however both approaches have limitations. Recall that (1) our interest is to understand the nature of moral *judging* as distinct from the content of the judgement, and (2), we are interested in those conditions that make such judging possible — that is, conditions necessary for a judgement to be a moral judgement. (1) makes the linguistic method unappealing because it does not target our primary interest, the judging, and it seems as though we may need a bridging strategy that takes us from the conceptual issues to the psychology. (2) makes things difficult for the psychological approach. In developing an account that systematically details the causes of moral judgements we will again be wide of our mark. It is not clear that in giving the causes of moral judgements that we will be giving the nature of moral judgements. A

causal/empirical account yielded by this approach would need to be supplemented with a story that detailed why the causes of moral judgements exhausted the nature of moral judgements. But to supply this would be to give an account of their nature — and this is the task at hand.

This leaves the third option — our experience of moral judgements. In starting here we would not be looking for common features within the content of our moral judgements or amongst the psychological precursors of such judgements, but for general features of our experience of the judging itself. The advantage of this strategy is that it addresses, immediately, that aspect of moral judgement against which the account is to be assessed. As we observed earlier, part of the interest in the nature of moral judgement is to better understand how and why we experience moral judgement as we do. An account that fails to cohere with the ordinary phenomena of moral judgements will undoubtedly fail to satisfy us. By beginning with experience we are seeking to include at the outset, the standard of acceptability to be applied to the entire account. The very serious disadvantage of beginning with our experience of moral judgement is that it looks to be a purely introspective matter. It is not obvious what would constitute an adequate strategy for investigating, as it were, pure experience at the required level of generality. I might claim that my experience is such and such on the basis of inspection but it is hard to see how I could compare this to someone else's experience. In the cases of the linguistic and psychological starting points there was a natural way to proceed. But beginning with experience and proceeding by a continued concentration on experience is in severe danger of rendering the account merely one of personal testimony. Dwelling too much on

the experience of moral judgements, since our best access to it is introspection, will not provide us with an account that is sufficiently general to be of interest.

The mistake here is to think that where we start dictates how we should proceed. If we begin by focussing on our experience of moral judgements we have not prevented ourselves from utilizing precisely the considerations adopted in the linguistic and psychological strategies, namely, careful attention to the behavioural and linguistic evidence — reports of moral judgements correlated with situations, expressions/reports of psychological states, reasons for action and descriptions of action. Thus instead of using the ‘data’ to explore the meaning of moral words or to uncover the initial psychological states antecedent to moral judgement, we will use it in the attempt to uncover and investigate general features of our experience of moral judgements. Introspection plays a purely negative role in this. It is a check against which any suggestion is to be measured. Introspection is only a problem for those accounts that place too much justificatory burden on it. By initially focussing on our experience of moral judgements and then attending to the behavioural and linguistic evidence, introspection is involved but only in a ‘supervisory’ capacity.

Maurice Mandelbaum’s project in *The Phenomenology of Moral Experience*, has several important similarities with this one. The one thing that Mandelbaum’s discussion of method, I think, lacks is an explicit description of how to proceed from the starting point through the course of the discussion to the resultant account — what I have been calling ‘a strategy’. This lack is brought to the fore by some remarks made by P. H. Nowell-Smith in his review of the book. These remarks also serve to highlight what I think is a general discomfort on the part of ‘analytic philosophers’ when it comes to

approaches grounded on a 'phenomenological method'. Nowell-Smith, it seems, is unsure of what Mandelbaum means by the 'phenomenological method'. At one point in the review he suspects that "the linguistic and the phenomenological approaches are identical". Several paragraphs later when speaking of phenomenology more generally he claims that it "reduces to psychology"¹². Although these remarks appear contradictory, they are both somewhat understandable. Mandelbaum proceeds through the book in a way very reminiscent of the kind of conceptual analysis to which Nowell-Smith was attracted. However, the discussion is centred on primarily psychological notions rather than conceptual ones. As I hope can be seen from the discussion above, this confusion is the result of the fact that the phenomenological approach which Mandelbaum espouses does not bring with it an explicit strategy. The final step here is to provide this strategy.

The form of argument that is most often and most naturally associated with the combination of the kind of starting point and end point that we have here is a transcendental argument. And, with a few important exceptions such is the form that I will adopt. Given that we will begin by examining our experience of moral judgement in the attempt to elicit general structural features and that we hope to provide an account of the nature of moral judgement, for the relevant understanding of 'nature', a transcendental-style argument seems natural.

A transcendental argument, generally, can be construed as one that endeavours to establish that some condition or conditions, A, are necessary for the possibility of some state of affairs (pertaining to the world or to our experience), B. Both Charles Taylor and Karl Ameriks argue that transcendental arguments make their moves by a regressive

¹² P. H. Nowell-Smith, Review of '*The Phenomenology of Moral Experience*', *Philosophy*, 32 (1957): pp.170-3. Both quotations are from page 172.

argument.¹³ Thus, a transcendental argument might be said to be an argument with an “unchallengeable” feature of experience as a starting point, which proceeds in a regressive manner by a string of necessary, ‘indispensable’ claims and whose conclusion reveals conditions for the possibility of some state of affairs pertaining to the world or to our experience. The regressive style and indispensability of the claims reflect the idea that these arguments provide analyses of the initial claim about experience — as it were, uncovering ‘deeper’ structural features of the experience. We can I think generalise these conditions and describe the form of these arguments in the following way:

The state of affairs, B, occurs. In order for B to occur we must have an experience E. This can be alternatively stated: E is part of what it is for B to occur (E is constitutive of B). But in order to have E we must have A.

Hence, A is a necessary condition for the possibility of B.

This generalised argument has only one step — that from experience E to condition A.

However, in more complex transcendental arguments there would be more intermediate steps.¹⁴

¹³ Charles Taylor, ‘The Validity of Transcendental Arguments’, *Proceedings of the Aristotelian Society*, Vol. 79, 1978-9, pp.151-65. Karl Ameriks, ‘Kant’s Transcendental Deduction as a Regressive Argument’, *Kant Studien*, vol. 69, 1978, pp.273-287. Taylor characterizes these arguments by three features:

(1) “They consist of a string of what one could call indispensability claims.”
 (2) These claims “are not merely probable, but apodeictic.”
 (3) These claims concern experience which “is held to give us an unchallengeable starting point.” (pp.159-60).

¹⁴ The example that is given by Taylor concerns the relation between playing chess and the possession of the rule governing the movement of the queen. It may be paraphrased in the following way:

We play the game of chess. Part of what it is to be able to play a game is to grasp the point of the game. That is, grasping the point of the game is a necessary feature of our experience of playing games. But grasping the point of the game is knowing what the game is about. Chess is a game

It is notable that if the necessity of the first, experiential step is denied the necessity of the conclusion fails to be secured. That is, whether A is a necessary condition for the possibility of B depends upon the experience E being constitutive of, or a necessary feature of, B. I suspect that it is this step that is the most problematic in transcendental arguments precisely because it makes such a strong necessity claim about experience. In the argument that I will provide in the ensuing chapters I will weaken this initial step. Thus, what makes the argument that I will present *not* strictly an example of a transcendental argument is that I will not argue, and hence not claim, that the experiential feature of moral judgements upon which the argument is built, is a necessary feature of moral judgement. I do think that this experiential feature is a very good candidate for being a necessary feature of moral experience but I do not see any way of proving that it is so¹⁵.

Even though this obviously central characteristic of transcendental arguments will be absent from my argument, enough other characteristics will be retained to make it transcendental in style. I will begin with a general claim about our experience of moral

about trying, by moving the pieces around the board, to achieve a particular final state. Given that the queen is one of the pieces, knowing the rule governing the way this piece can move is a necessary part of knowing how to move the pieces around the board. Hence, knowing (or possessing) the queen rule is a necessary condition for being able to play chess. (Taylor, *op.cit.*, pp.161-3)

This is by no means a complete argument but it does serve to highlight the kind of structure in question. The first step in this argument is one concerning an essential feature of our experience of playing games and since chess is a game, our experience of playing chess. The steps following this are the products of an analysis of this feature — what it is to grasp the point of activities such as games, what, specifically, chess is about, etc. An analysis such as this has as its disposal facts about our behaviour, our language and our experience.

¹⁵ G. E. Moore, in *Principia Ethica* ((1903) Cambridge, C.U.P., 1962), uses ‘self-evident’ to describe a similar kind of situation to this one. See particularly pp.142-5.

judgements of obligation. The feature that I will isolate and discuss is, I think, a good candidate for an essential, experiential feature of these judgements. The argument will then proceed by analysing this experience in the light of the behavioural and linguistic evidence in order to characterise the structural features of these judgements of obligation. After several steps in a 'regressive' chain of premises we will have an account of the nature of moral judgement. Of course given that I will not claim that the initial experiential feature is necessary the concluding account will not yield *necessary* conditions for the possibility of moral judgement. Instead the resulting conditions, and so the account, will be conditional upon the acceptance of the initial step.

3. Structure and Argument

Overall the thesis can be divided into two unequal parts. The first part, comprising chapters one through five, contains the development of the quasi-transcendental argument mentioned above. At the end of chapter five we have arrived at the key feature of the account. Chapters six and seven develop this feature and explore its relationship to both cognitivist and non-cognitivist positions.

The specific steps in the argument are as follows:

1. We make moral judgements of obligation. (Assumption)
2. When we make moral judgements of obligation, it is consequent to feeling a pressure which seems to originate independently of us. [There are two features that are characteristic of the pressure and which must be accounted for in the explanation of moral judgements of obligation: the motivational character (the

conative force) of the pressure; and, the 'external' or 'objective' character of the pressure.]

3. Whenever we feel this pressure (which seems to originate independently of us), it is a result of our 'frame of mind' being configured in a certain manner.
 4. Whenever our 'frame of mind' is configured in this manner (i.e. such as to result in the judgement and the experienced pressure), we have attended to particular features of the context in a certain way.
 5. The tools that are available to explain both the judgement and the two characteristics of the pressure are states that are broadly characterized as belief-like or desire-like. (Assumption)
 6. An account which privileges belief-like states in the explanation of moral judgements of obligation (typically, cognitivist accounts) cannot accommodate the motivational character of the pressure.
 7. Desire-like states can have three kinds of objects: (1) representations (things, states of affairs, etc.), (2) other desire-like states or (3) ways or modes of representing. (Assumption)
 8. An account which relies on desire-like states with objects of kind (1) or (2) cannot accommodate the second, 'objective' characteristic of the pressure.
- C. An account which centres on desire-like states which have 'ways of representing' as their object is well placed to account for both characteristics of the pressure that need to be accommodated.

Each of the first five chapters corresponds to steps in the argument. So, in the remainder of this chapter we consider premise two. Chapter two considers and defends

premise three by examining the arguments in H. A. Prichard's 'Duty and Ignorance of Fact'. Chapter three primarily examines premise four. Among other related things, this chapter looks at the way in which discussions of rule following impinge on the argument here. Chapter four involves discussion of Aristotle as well as an investigation of John McDowell and David Wiggins' cognitivist accounts of moral judgement. It is concerned with the development of premise six. Premise eight is the subject of chapter five. There we examine some non-cognitivist strategies for handling the experienced pressure. At the end of chapter five we are left with the conclusion of the argument

Chapters six and seven contain an exploration and development of the concluding step of the argument. Importantly the methodology of these two chapters is significantly different from the preceding chapters. Rather than developing the steps of the argument, these chapters, and particularly chapter six, are concerned to flesh out the position. The initial formulation of the position is refined and made more precise. The focus here is on how we should understand a 'desire to notice' and more generally desires that have 'ways of representing' as their object. The development and extension of the position is done in large part through a series of examples. Each of the examples is intended to illustrate an important feature of the central idea. Overall in chapter six the discussion centres on the 'structure' of these desires and the way in which they explain our moral judgements of obligation. Chapter seven examines the influence and role that the agent's beliefs play alongside these desires. The resulting picture is a complex blend of desire and belief-like states that are united and concentrated at the point of judgement by a desire-like state exemplified by a desire to notice.

B. The Experiential Feature of Judgements of Obligation.

In this section we turn to the experiential feature associated with moral judgements of obligation. As we have just seen, this feature is the starting point of the account — the initial, assumed premise of the argument. In the first part of the section I will give an account of the phenomenon in question whilst in the second, I will draw a useful and hopefully simplifying distinction between two kinds of judgement of obligation. This distinction will help focus attention more clearly on the paradigm case judgements.

In order to maintain the plausibility of the claim that all of our moral judgements of obligation are accompanied by this experience it will be important not to over-specify or over-describe it. There is a danger, the danger associated with introspection, that in characterizing this experience too finely it will lose its generality. A finely grained description of this experience runs the risk of simply being a characterization of features of my experience rather than experience generally.

1. A 'Pressure' or Force.

The experiential feature that accompanies moral judgements is a pressure or force of a certain kind. It is a pressure on us to act, to feel or to think. It influences our attitudes, our emotions, our thoughts and the things that we do. Feeling drawn to pick up the child who has tripped or becoming indignant when a helpful gesture goes unacknowledged are normal kinds of instances that exhibit this pressure or force. We might feel pulled towards the child, as though the child, simply by its predicament, demanded assistance. The rudeness exhibited in ignoring a helpful gesture makes us feel

indignant. We might feel an urge to say something sarcastic like “You’re welcome”. When it occurs to us that Fred ought not to treat his wife like that, it is Fred’s treatment of his wife that is offensive, inappropriate or un-called-for. The pressure that Fred’s actions exert on us is expressed in terms of their inappropriateness or that they are un-called-for. We recoil from the offensiveness.

This pressure or force is also reflected in a good deal of our moral language. “You ought to do such and such” is intended to bring with it some pressure or to provide some compellingness. The experience of the so-called ‘moral ought’ is precisely what is at issue. In cases where we are obliged or duty-bound we feel constrained. In cases where we think someone else has an obligation we feel a similar pressure but it is a pressure on them and their actions.

The philosophical literature is scattered with allusions to this kind of pressure or force¹⁶. In the first chapter of *Ethics: Inventing Right and Wrong*, Mackie is concerned with what he calls, ‘the claim to objectivity’ in moral language. He writes that “the ordinary user of moral language means to say something about whatever it is that he characterizes morally, for example a possible action, as it is in itself, or would be if it were realized ... But the something he wants to say ... involves a call for action or for the refraining from action”¹⁷. Here, the pressure is the call for action or for the refraining

¹⁶ It is an interesting fact that for the most part those who seem sympathetic to an investigation of this feature have so-called ‘cognitivist’ or deontological leanings. A concentration on this experience is found in such authors as Butler (particularly when he introduces conscience early in his *Fifteen Sermons*) and Ross (*Foundations of Ethics*). More recently, it has been given some significance by David McNaughton (in *Moral Vision*) and Mark Timmons (in *Morality Without Foundations*). Timmons, whilst not falling in the above categories, thinks this phenomenon needs to be accounted for.

¹⁷ Mackie, J. L. *Ethics: Inventing Right and Wrong*. (1977) London: Penguin, 1990: p.33.

from action. However, the concern here is not the force or the pressure that I intend my utterances to carry but that demand which, when they are 'successful', they do carry for those for whom my utterance is intended. We might be tempted, in this kind of case, to think that the intention that my utterance place a particular kind of demand on another person makes sense because I have reason to believe that this person might be affected by such demand. I may, for instance, have evidence that this person responds to demands of this, perhaps moral, kind. Alternatively, I may simply think that people ought to be affected by demands of this kind. In both of these cases, the intended effect of my utterances seems to presuppose something about the way in which the hearer might be affected by my utterance.

In a passage in chapter three of his book, Mackie captures the idea of the pressure nicely by reflecting on etymology. He writes:

'Ought' is a past tense of 'owe', but used as a present: if *A* ought to *G*, *A* somehow owes that he/she/it should *G*. But what is it for Jones, say, to owe Smith five dollars? Is it as if five dollars in Jones's pocket or bank account were trying to fly across into Smith's? Rather, since Jones's pocket and bank account may both be empty, it is as if there were an invisible hook reaching out from Smith and fishing for the money in Jones's pocket. Or an immaterial suction-pipe. The owing is itself something like a demand for payment. Similarly if *A* ought to *G*, there is something about the situation that sets up an expectation or presumption of *A*'s *G*-ing. A similar metaphor is buried in the term 'obligation'. If *A* is obliged to *G*, it is as if it were tied down to *G*-ing: an obligation is an

invisible cord. If something is offered without obligation, there are no strings attached.¹⁸

Another way of characterizing the pressure associated with judgements of this sort is through an analogy with aesthetics. This idea has a long and storied history¹⁹. As Ross writes in *Foundations of Ethics*,

There seems to be something not altogether different in the way in which a situation calls for a certain act, and the way in which one part of a beautiful whole calls for the other parts. ... The harmony is not of the same kind — rightness is not beauty; but there seems to be a genuine affinity, which justified the Greeks in their application of the word καλόν [beautiful, fine, noble] to both.²⁰

In a perhaps surprising section of *The Critique of Aesthetic Judgement*, entitled ‘Beauty as the symbol of morality’, Kant writes that

Even common understanding is wont to pay regard to this analogy; and we frequently apply to beautiful objects of nature or of art names that seem to rely upon the basis of a moral estimate. We call buildings or trees majestic and stately, or plains laughing and gay; even colours are called innocent, modest, soft, because they excite sensations containing something

¹⁸ *ibid.*, p.74.

¹⁹ Among those that might be mentioned here are Plato, Clarke, Shaftesbury (particularly *British Moralists*, §201, p.172) and Price.

²⁰ Ross, W. D. *Foundations of Ethics*. Oxford: Clarendon Press, 1939: 54.

analogous to the consciousness of the state of mind produced by moral judgements.²¹

Mandelbaum, following Ross (as well as Broad and Ewing²²), exploits this analogy with aesthetics in explaining moral experience²³. His notion of fittingness is intended to pick out this pressure. An action seems to fit or complete a situation and we are drawn by this need for completion to act or judge. Here, the comparison with aesthetics is intended to be a variant or instance of the pressure. However, it might also be used to point to the kind of experience that is particularly associated with moral judgement more generally — concerning matters broader than action and obligation.

In all of these examples of our moral language, it is important, again, to stress that the experiential feature I am attempting to depict is not the force or pressure on action which our utterances *are intended to exert* but the pressure or force that inclines us to make such utterances at all. This difference might best be exhibited in cases where we are tempted to make such an utterance but do not for whatever reason. We might think to

²¹ Kant, Immanuel. *The Critique of Judgement* (trans. J. C. Meredith). Oxford: Clarendon Press, 1952: 225 (in the original pagination: 354). Paul Guyer explores the connection between aesthetics and morality in Kant in his book, *Kant and the Experience of Freedom* (Cambridge: C. U. P., 1996).

²² Broad, C. D. *Five Types of Ethical Theory*. (New York: Harcourt, Brace & co., 1930):164-5. Ewing, A. C. *The Definition of Good*. (New York: Macmillan, 1947.)

²³ Mandelbaum, p.60ff. I have some reservations about characterising all occurrences of the experience in question using quasi-aesthetic terms like ‘fitting’ and ‘unfitting’ as Mandelbaum does (for the most part in line with the tradition). This is mainly due to the fact that some of our moral experiences affect us more ‘violently’ than could be conveyed by these terms. To find examples of this one need look no further than novels and films recounting experiences in war or testimony relating the atrocities of the Holocaust (see for instance, Raimond Gaita, *Good and Evil: An Absolute Conception*. London: Macmillan, 1991: 1.) Such examples are ‘gut-wrenching’. They can sap the spirit in a way that seems to exhaust the adequacy of the aesthetic analogy. This, of course, is not meant as a substantive criticism but a quibble about the picture that the analogy gives us.

ourselves how appalling some behaviour was but not say anything because of 'present company'. We are still affected by the behaviour and we still make the judgement (we still think it is appalling) even though it is not expressed. Whether or not we end up expressing our judgement, given that the judgement is made I want to claim that we experience this normative pressure, we are affected in a certain way by the situation. This is not to suggest that when we do give voice to our moral judgements that we do so without the knowledge that they will have an effect or with the intention that they do have that effect. Instead, I only wish to highlight the presence of the experienced force accompanying the moral judgement.

I have said that this experience accompanies moral judgement but to this point the sense of this accompaniment has been obscure. In attempting to place this experience we might be tempted to think that the experience is part of what it is to have judged that such and such. This would make the experience something analogous to the qualitative experience associated with belief. In turn this would mean that the moral experience is 'what it is like' to judge that such and such is wrong. This is not the relation between moral experience and moral judgement that I intend. Rather, the relation I intend may be illustrated through the following examples. I am confronted by a situation in which Gilbert Harman's cat is being mistreated and judge that I ought to do something about it. I witness Fred's conduct with respect to his wife and judge that he ought not to act in that way. In each of these instances and, I claim, in all other instances of moral judgements of obligation we experience a certain kind of pressure which goes into and is thus explanatorily prior to the judgement. Somehow, in a way which needs to be investigated,

the children's mistreatment of the cat and Fred's conduct put us under a certain kind of pressure.

I realise that this is a far from adequate characterisation of the relationship between the judgement and the phenomena. However there is no real need for me to be more precise since it will be recalled that it is this pressure that is being used to understand the nature of moral judgements of obligation. It is by investigating this moral experience and, as it were, 'unpacking' it that we can understand and explain what 'goes on' when we make a moral judgement — that is, in accord with to the method outlined in the first section. In any case, the precise sense in which the pressure accompanies the judgement *is* obscure. It has been suggested that rather than separating the judgement of obligation from the experience, the two should be thought of as identical.²⁴ That is, it might be thought that the feeling/registering/recognition of a pressure simply *is* the judging that one has an obligation. This would certainly explain why the accompaniment relation is difficult to characterise. To properly investigate the worth of this much stronger claim is itself a major undertaking and certainly not one to be attempted here. However it should be pointed out that it is not crucial for the project that the pressure be thought of as separate from the judging. Recall that I did not wish to claim that the pressure is a necessary feature of moral judgements and that this meant that the argument would only be a 'quasi-transcendental' one. If it could be shown that the experienced pressure amounted to a judgement of obligation then the necessary connection would be more than adequately met. The remainder of the argument of the thesis would remain the same except that any forthcoming explanation of the pressure would be a direct

²⁴ In different ways and in different contexts by Stephen Toulmin and David McNaughton.

explanation of the judgement. For the purposes here I will rest with the weaker and more obscure claim.

Now quite clearly things may not proceed in the way that the examples above suggest. I may see the children's treatment of the cat and not be moved or I may not have noticed anything about the way Fred treated his wife. In both of these cases there is no pressure, but there is also no moral judgement. I may of course be judged by someone else for not being moved or not noticing the way Fred behaved. The third party, my judge, would then have been affected, perhaps insulted, by my lack of judgement. Alternatively, I may think that Fred's wife 'had it coming', that she deserved to be treated in that way, in which case the pressure would be different and I might judge that Fred ought to have treated her as he did. A final, but important, variation is where one notices but does not judge²⁵. Here, I might see the mistreatment and recognise it as such, but feel that it is a private matter and none of my business. Similarly, I might be familiar with Fred's exchanges with his wife and simply not have the energy to bother about it. Perhaps they are keeping me from watching an interesting basketball game. In both of these scenarios the mistreatment is acknowledged but no moral judgement ensues. In the former, my response that it is none of my business might be the result of a prudential judgement. That is, I might think that in cases like this it is best (for me) to remain uninvolved. The pressure is felt but the resulting judgement is not something that I take to be a moral judgement. It is prudential. In the latter version, again I feel the (all too

²⁵ This variant is intended to have parallels with an example considered Simon Blackburn and John McDowell in their debate. The example, concerning recognising that someone is shy and sensitive, is initially given by McDowell in 'Are Moral Requirements Hypothetical Imperatives?' (*Proceedings of the Aristotelian Society*, supp. vol. 52 (1978): 21-2.) and is taken up by Blackburn in *Ruling Passions* (Oxford: Clarendon Press, 1998: 100).

familiar) pressure but lethargy or my desire to concentrate on the game prevents my judgement.

An instructive point may be taken from these final variations. My claim with respect to the experience of pressure and its relation to moral judgements of obligation is that when we make such a judgement we experience this pressure, *not* that when we experience the pressure we judge. There are, it seems, two kinds of cases in which one might experience the pressure and not make the relevant judgement. First, the pressure may be taken to concern a non-moral matter and second, the pressure may be, as it were, overcome by a pressure or force of another kind. In short, the experience 'pre-empts' the moral judgement — it is not consequent to the judgement and it is not a product of the judgement.

There are several important features of this experience that should be emphasized if we are to proceed with its investigation. First, it is useful to contrast this kind of experience with that of preference. In the normal case, a decision or judgement based on preference does not possess the same kind of force. We may strongly feel like coffee rather than tea and we may dislike tea so much as to be repelled from the idea of drinking it. However, such tendencies are experienced as pressures emanating from 'within' the agent. A crucial feature of the experience accompanying moral judgements is that the pressure or force to act, judge, feel or think is experienced as independent of the agent. I must be clear not to overstate this point. My claim is only that the force is *experienced* as independent, not that it is independent. How it comes to be experienced in this way and whether the force in fact emanates from outside the agent will be the subject of much consideration below. Here the distinction is between the experience of the force of a

preferential judgement and the experience of the force of a moral judgement. Understanding 'objective' as independent of the agent and 'subjective' as dependent on the agent, the claim is that the experience accompanying a moral judgement is experienced as objective²⁶.

A second feature of the force or pressure that I have been attempting to describe is that it might well be called simply a normative force. It might appear as though the discussion has drifted from one concerning a particularly moral experience to the broader region of normativity and a normative experience. Such a complaint would be reasonable, although a reminder is needed. The normative force in question is indeed, I would claim, a feature accompanying normative judgements generally and not just moral ones. We might say that the experience of normativity just is this pressure. Hence, the normative force is experienced when we judge Van Gogh's *Irises* to be beautiful or when we think we ought to exercise regularly. In both of these cases, as in the examples above, the pressure is not experienced in the same way as simple preference. It feels as though the painting itself is, as it were, doing all the work. It feels as though beauty is an attribute of the object, there for all to see²⁷. Similarly, although the last thing I want to do is go for a run, I might feel as though I really should. It is certainly the case that I have a general desire to remain healthy by exercising regularly but in this instance it is not what I want to do. My general commitment to health through exercise presents itself not as a

²⁶ See Mandelbaum, pp.55-8.

²⁷ The position in aesthetics with respect to the 'experience of beauty' and judgements of beauty seems to me to parallel the moral case. (As we saw earlier parallels of this sort have been noticed and discussed widely.) Whether or not beauty is a property of the object is not the issue. It is the experience of the normative pressure as originating outside the agent that is of concern.

‘subjective’ demand but as one imposed on me from without. Mandelbaum sums this up nicely:

In order to feel an obligation to satisfy one’s desire, the fulfillment of that desire must be seen as a means to some further end which itself appears as “objectively good,” that is, as good independently of one’s preference for it. (p.57)

In short there is nothing in the experience of the pressure that distinguishes a moral force from a non-moral, but normative force. What makes the present discussion one concerning moral judgements and the experience associated with them is only the very tenuous criterion from the previous section. That is, the pressure that I have been describing is not to be thought of as isolating or picking out *moral* judgements of obligation. As we saw in the first section, to separate the moral judgement from the non-moral judgement requires distinguishing between contents or some account of the nature of moral judgement. The former runs the risk of a breach of the non-normativity requirement, whilst the latter is the current task.

One final, more complex kind of case should be mentioned²⁸. Sometimes we are confronted by two (or more) normative pressures that pull us to act in opposite ways. Sometimes we might judge that we have conflicting moral obligations²⁹. Following the

²⁸ This kind of case derives from Mandelbaum’s discussion and criticism of Ross (pp.74-81). His main point is that the initial claims that two conflicting *prima facie* duties make on us is not resolved by appealing to the relative strengths (“stringency”) of the two duties, but that the reconsideration of the situation, acknowledging the conflicting duties, enables us to judge which course of action is fitting. A consequence of this is that there is no need to analyse and rank the relative strengths of *prima facie* duties. Issues related to this will be discussed in the next chapter.

²⁹ As we have noticed the normative pressure need not be taken to involve a moral issue. In such a case the conflict would not be described as a moral conflict.

story to this point, we experience a pressure to act in way X and a pressure to act in way Y. The ensuing judgements that we are obliged to do X and also that we are obliged to do Y give us the conflict. In instances such as this, it is natural to introduce the notion of deliberation and decision — we have to think about what to. Now there are clearly many ways of thinking about or reconsidering a situation and here is not the place to consider them. If there is a upshot of this ‘deliberation’ and a judgement or decision is made then, I submit, there is also a ‘new’ pressure which precedes the judgement. Now, I do not mean anything too significant by suggesting this new pressure. It may be the reiteration of a previous pressure or it result from the recognition of a compromising solution. The motivation for calling a ‘new’ pressure is that it includes the prior recognition of the conflict and the results of whatever deliberation took place. All that I wish to register here is that this experiential pressure accompanying moral judgement does not preclude such important features of our moral judging as pondering, dwelling and reconsidering — subsumed under a general notion of deliberation.

It might appear as though there are two general kinds of exceptions to what I am proposing. Both actions performed out of habit and reflex actions might be taken as moral actions but do not involve an experienced moral pressure (or at least an awareness of a pressure). Let us take habit first. Wilbur gives his change to the homeless man on the corner so often that he no longer notices that he does so. It is true that there is no experienced pressure associated with a particular instance of such a habit. Wilbur, on a particular day, does not feel any pressure to give to the homeless man, but importantly, nor does he judge, say, that he ought to do so. He might think “There’s John” and reach into his pocket. We might say that his giving is a conditioned response. Notice however

that if Wilbur were to reflect on what he had just done or what he usually does he might think that he is right to give regularly. What had been, as it were, a 'mere' habitual response might then become something that he chooses to do and is aware of as being morally appropriate behaviour³⁰. His moral judgement was one that applied to his pattern of behaviour as a whole rather than to a particular instance of it.

A similar lack of judgement can be seen in cases of reflex or, perhaps, reaction. A clear example of this would be catching a child who seems to be lurching into on-coming traffic. Such a case seems akin to saving oneself, which in turn looks to be similar to, say, recoiling from intense heat. The key point here is that we might be disinclined to attribute a judgement in any of these instances and given this, they do not constitute counterexamples. Some of these cases however do not appear straightforward. Even if one is tempted to attribute a judgement in the first or second kind of case, it seems hard to also resist allowing for the experience of the pressure. Thus, in both cases of habit and reflex behaviour it is true that there is no experienced pressure but there is also no moral judgement made and these cases do not constitute exceptions to my claim that the experience of pressure accompanies moral judgement³¹. Of course we may step back and

³⁰ There is an interesting parallel here with Butler's discussion of the principle of benevolence and conscience. In Butler's terms, by reflecting on what he does and that, according to him, it is right, Wilbur subsumes his behaviour under the principle of benevolence and is approved of by his conscience. This behaviour is then governed by a "much more settled principle" (Sermon I, §8, p.39). Butler, Joseph. *Fifteen Sermons Preached at Rolls Chapel*, ed. Revd. D. Matthews. London: Bell & Sons, 1958.

³¹ Mandelbaum's concern in discussing these kinds of cases is to distinguish between actions that are 'willed' and those that are spontaneous. Actions are willed when "the agent, even while performing them, feels responsible for them" (p.47). He suggests that there is a third kind of spontaneous action that is neither reflex nor habit and gives three examples of such actions: "I sense the embarrassment of a person and turn the conversation aside; I see a child in danger and catch hold of its hand; I hear a crash and become alert to help." (p.48) I am not so convinced about this third kind particularly if

judge our own patterns of behaviour or reactions in certain situations, but these judgements are more akin to our judgements of others (as we shall see) and *are* accompanied by the pressure.

2. Two Types of Judgement of Obligation.

Describing the experience in the way that I have, as a force, provides us with a useful heuristic for understanding it more thoroughly. Forces can be described in terms of strength, direction and origin³². A block on an inclined plane is held in place by the force of friction. The strength of this force depends on the angle of the plane and the weight of the block. The direction of the force is up the plane and its origin is the contact between the surface of the block and the plane. Similarly, our experience of the normative force when we make moral judgements is an experience containing these three components. I shall postpone any discussion of the strength and origin of the normative force for the

we talk of reactions rather than reflex or spontaneity (although there are problems with this also). The first example does not seem to be obviously a 'non-willed' case. Changing the topic of conversation seems a very deliberate and conscious act, especially when it because of someone's embarrassment. The second example needs to be specified a little more. It would seem that if the situation required the helper to move towards the child to help (say a couple of steps), the action would be willed in Mandelbaum's sense. He clearly does not mean this. The example, I assume, is intended to be more of a reflex action, akin to sidestepping to save oneself. However, now there does not seem much difference between this and a genuine reflex. That someone should have this reaction rather than some other remains puzzling. For instance, someone might push the child out of the way rather than catching and pulling back. The basic point here is twofold. Catching the child looks to be a reflex action and so explainable as such — like avoiding an electric shock. But even in some of these spontaneous cases there appears to be more choice or deliberateness than Mandelbaum seems to allow. The third example, I think, bears this out. Upon hearing a crash, I certainly become alert, this much is reaction. But becoming alert *to help* seems to require more explanation. Someone might become alert to help if they knew or were expecting that something requiring help could happen.

³² The central idea here comes from Mandelbaum (p.54) although he exploits it slightly differently and not to as great an extent as I will.

moment, instead concentrating on variations in direction. The distinctions that I will make here correspond to and are derived from those made by Mandelbaum.

From his suggestions there are two types of moral judgements of obligation, each of which reflect differences in the direction of the experienced normative force: direct moral judgements and removed moral judgements³³. He writes

By the first I denote those judgements made by an agent who is directly confronted by a situation which he believes involves a moral choice on his part. By the second I denote those judgements made by an observer on the conduct of another person (or by an agent on his own past conduct). (p.45)

The most important amendment that I think needs to be made to this classification concerns direct moral judgements. Mandelbaum's account requires that (i) the agent be confronted by a choice and (ii) that the agent be aware that the choice confronting him is a moral choice. These two requirements are, it seems, a little strong. We might say that for the most part these two conditions are satisfied but there do seem to be occasions where one or both fail to be satisfied and we make a moral judgement. In regard to the first condition Mandelbaum writes:

Now, in every willed action the agent is in effect confronted by at least two alternatives of whose presence he may be said to be aware: at the very

³³ He actually suggests that there are three types of moral judgement. The third, judgements of moral worth, I take to include and indeed exhaust judgements of value. Given the link to action in the two types mentioned above and given our concern with judgements of obligation, I will not discuss this third category.

There is more to be said about these issues. Darwall, Gibbard and Railton (*op.cit.*, note 91, p.41) suggest that the phenomenological characteristics of the two kinds of judgement are different. I have some sympathy with this idea. As I remarked earlier it might be that the idea of 'fittingness' works better for judgements of value.

least he has the possibility of abstaining from that act which is aimed at the goal which he contemplates. (p.49-50)

The notion of a 'willed action' enters the picture for Mandelbaum because he wants to limit the kind of thing that can be the object of a direct moral judgement. At this point I am not so concerned with the object of these judgements as with the way in which we draw the distinction the two kinds of judgement. Thus, for my purposes, we can read "willed action" in the above passage as 'direct moral judgement'. Considering this reading, there do seem to be exceptions. It seems possible that an agent finds himself in a situation where, as far as he is concerned (as far as he judges), there is no possibility of not acting — he *must* act³⁴. Instances of this are (perhaps) many cases of telling someone the right time when asked or giving (or at least attempting to give) someone the best directions to their destination. It would rarely, if ever, even occur to us to give the wrong time or the wrong directions. In kinds of cases like these where we just act, we certainly feel compelled to act — we experience the pressure. However, it is as though there is no other choice — the situation just seems to present us with what we have to do, with what is required. What is peculiar here is the 'felt necessity' of the action. We might explain this by speaking in terms of the strength of the pressure to act. It is so great or so clear that other 'actual' alternatives do not, as it were, appear as alternatives³⁵.

³⁴ Bernard Williams ('Practical Necessity' in *Moral Luck*. Cambridge: C. U. P, 1981) discusses issues that are related to this one. He describes these and related cases as practically necessity. Other useful discussions of these issues can be found in Peter Winch, 'The Universalizability of Moral Judgements' (*Ethics and Action*, London: Routledge, 1972) and Raimond Gaita, *Good and Evil: An Absolute Conception* (London: Macmillan, 1991).

³⁵ This is the precursor of a point that will be discussed later in connection with some of John McDowell's suggestions. McDowell calls this suppression of alternatives "silencing" ('Virtue and Reason'. *Monist* 62 (1979): 331-50.). It is also noteworthy that

These examples are also of use with respect to the second of Mandelbaum's criteria for direct moral judgements — that the agent be aware that the choice confronting him is a moral choice. Quite trivially if the agent is not aware of there being a choice then he could not be aware of it being a *moral* choice. Another example is useful however. It seems quite plausible that two people (A and B) may agree that one of them (B) had an obligation to act in a certain way without agreeing that the obligation was a moral one.

A: "You ought to have worked late on that project because that was the best way to ingratiate yourself with your boss and staff cutbacks are coming".

B: "No, I had promised the boss that I would have the project completed by the next day."

A might go on to produce evidence in support of his claim and so convince B of his point. Having been convinced, B would no longer think that the obligation to work late was a moral obligation. This is an instance where the agent is confronted with a choice but thinks, at the time of choosing, that the choice is a moral one. Similar examples can be imagined where B's change of mind went the other way and also where B had not considered the kind of obligation at all. In an example of the latter kind the agent would only have been aware of the obligation to act one way rather than another. Consideration of what kind of obligation, specifically moral or other, is not necessary for an agent to think that there was an obligation. Of these cases, those of interest to us are clearly those in which the agent decides (or would decide) involve a moral issue and hence that the judgement was a moral judgement.

these cases satisfy Mandelbaum's account of willed action (p.47). Even though there is no choice, there is, it would seem an espoused goal and the experience of being responsible for the action.

There is an important distinction implicit in this discussion which highlights the current point and also introduces the distinction between direct and removed moral judgements. This distinction may be seen by considering the differences between (i) the judgement that I ought to do such and such, (ii) the judgement that I ought to have done such and such and (iii) the judgement that my obligation in case (i) was a moral obligation. This last judgement, (iii), is not a judgement of obligation. It is the product of reflecting on the situation in which I made the judgement in relation to the content of the judgement. The example in the paragraph above was intended to highlight the possibility of making a judgement of type (i) without a judgment of type (iii). That one can make a moral judgement without being aware that it is a moral judgement, is demonstrated by the fact that we can reflect on situations in which we judged and assess (or reassess) the kind of the judgement that we made.³⁶ Judgement (i) is intended to satisfy the requirements of direct moral judgements: I feel a pressure *on myself* to act and judge that I ought to do so. Judgement (ii) is intended to be a removed moral judgement. The judgement concerns my past action. I reflect on the situation with which I was confronted and feel a pressure, as it were, on my 'past self' (perhaps better: myself in that past situation). The judgement in a case like this is similar to the more standard, removed moral judgement where an agent judges the actions of another in a certain situation. The experienced pressure characteristic of these judgements is directed towards the agent in question. It is felt *by the judge* as a demand on the other person. It is important not to confuse the agent who is experiencing the force with the direction in which the force is acting. In the direct

³⁶ There is a good deal more that could be said at this point. For instance, it seems as though the relationship which holds between (i) and (iii) could also hold between (ii) and (iii). That is, I may not be aware that my removed judgement was a removed moral judgement.

judgement case, (i), the force is experienced as acting on the agent himself, thus the direction is towards the 'experiencer'. In the removed judgement case, (ii), the force is experienced as acting on someone else, thus the direction is towards some third party.

Having distinguished these two kinds of moral judgements of obligation and remarked on the direction of the pressure accompanying them, I will, for most of what follows, set aside removed moral judgements and concentrate on direct moral judgements. I certainly do not mean to suggest by this concentration that direct moral judgements are more basic or fundamental. I do think (and this is the main reason for the concentration) that in direct moral judgements we have the clearest and simplest access to the pressure. Removed moral judgements given that they involve others more closely are inherently more complex. It is worth observing that quite often a removed moral judgment leads to a direct moral judgement. Having judged that Fred ought not to treat his wife in that way in may be incumbent upon me to decide how I ought to act in the situation.

As a strategy then, it appears reasonable to focus, at least initially on those moral judgements where the variability and complexity is at a minimum. Quite clearly however, there is a danger in this strategy, namely that the considerations educed from the simpler cases fail to apply in the more complex ones. It will be necessary therefore to be aware of the consequences of the following investigation for removed moral judgements of obligation.

CHAPTER 2: OBLIGATION AND THE ORIGIN OF THE PRESSURE.

A. The Origin of the Pressure.

In this chapter we will begin the task of understanding the origin of the normative pressure that is experienced when we make moral judgements of obligation. As we have seen, using the analogy with physical forces, we can isolate variations in the experienced direction of the pressure. Corresponding to these 'directions', I have suggested a division within moral judgements of obligation into direct (where the pressure is experienced as being on me) and removed (where the pressure is experienced as being on another) moral judgements. As noted, in what follows I will concentrate primarily on direct moral judgements, as it is in these cases that the elements involved in the judgement are the least complicated. Whatever conclusions we arrive at regarding the pressure will need to be assimilated with judgements of obligation and our expression of such judgements. A large part of this chapter will be spent considering just these issues.

Perhaps the most intriguing and unique feature of the normative pressure is that it is experienced as 'objective'. That is, as we saw in the previous chapter, it is experienced as an external 'force' on the agent. It *feels as though* it is the situation itself that places me under this pressure, demands a response of a particular kind. I am compelled, urged or merely nudged towards acting, thinking or saying, seemingly, by the position in which I find myself. Understanding and explaining this feature of moral judgements is crucial if we are to proceed since, given the method outlined in the first chapter, it is this feature that separates normative judgements from non-normative ones. The first order of business

in this chapter, then, is to sort out whether the experience of the pressure as having an external origin reflects its actual origin. Do we experience the pressure as having an external origin because it, in fact does or is it somehow dependant upon our 'take' (version, view) on the situation?

A 'take' on a situation for our immediate purposes here includes those accounts that would explain the experience of the pressure and the judgement of obligation as, in some way dependent on the relationship between the agent and the situation. For the moment we will not be concerned with this alternative. Instead, I want only to determine whether the experience of the pressure is, in fact, objective (understood as independent of the agent) given that it is experienced as such.

Before continuing it is instructive to try to be clearer about 'a situation' and how to understand an agent's being in a particular situation. Mandelbaum suggests that

Among such [initial] conditions are to be included not merely the present conditions which we find ourselves confronting, but those past and future events which we recognise as being relevant to the choice which we are to make.³⁷

Not surprisingly, this characterization of 'a situation' has much in common with what was called 'the context' in chapter one. There I mentioned the time, place, the people involved, their psychological makeup/states and the immediate (or not so immediate) past and future. These features are included in the story because of their relevance to the choice. Here, however, we run into a problem. Part of the point of saying what amounts

³⁷ Mandelbaum, p.61. Here Mandelbaum is following Ross (*Foundations of Ethics*, p.81ff) against Broad (*Five Types of Ethical Theory*, pp.218-220) in that he wishes to allow more than utility into the account of the situation. Carritt is good on this also. See *The Theory of Morals*, p.91.

to 'a situation' is to help determine what counts as 'external to' the agent. However using Mandelbaum's account of a situation, with its dependence on those features that are relevant to the choice, will not allow us to isolate this external requirement adequately. On his view, two people could be in the same place, at the same time, witnessing the same events and yet be in different situations. Similarly, how we understand the situation confronting an agent, whom we observe, may be different from the situation which the agent understands himself to be in. That is, what we, as observers, take to be relevant to the agent's choice may differ from agent's view of his situation. What is relevant to the choice, therefore, will not allow us to distinguish that which is external to the judge since it is, at least partly, dependant upon the position of the judge.

Ross' account of a situation given at the beginning of chapter seven of the *Foundations of Ethics*³⁸ is helpful here. He suggests that we understand the situation in which an agent finds himself as containing an objective (or external) and a subjective (internal) element.

The objective element consists of the facts about the various persons and things involved in the situation. ... [Whilst,] the subjective element consists of the agent's thoughts about the situation. (p.147)

For the purposes of distinguishing between these two elements that comprise what it is to be in a situation, let us call the objective element, i.e. the facts of the situation, the context and the subjective element, the agent's thought about the context. Now perhaps the most important point about the facts that make up a context is that we have no way of limiting them. As we have just seen with Mandelbaum's use of 'relevance', we cannot give, with

³⁸ Ross, W. D. *Foundations of Ethics*. Oxford: Clarendon Press, 1939.

any precision, the set of facts that define a context. We will be able to say, of particular facts, that they are a part of a context but it will be much more difficult to say when a fact is not part of one. When considering, say, whether to purchase a pair of socks, the fact that the socks were made in a sweat shop in Mexico by underpaid labourers³⁹ may be relevant to some but not to others. A point similar to this one is sometimes made in connection with versions of utilitarianism. Sometimes buying a pair of socks is supporting a particular economic arrangement — this might depend on other, seemingly even more removed facts such as the effect that one customer refusing to buy an item has had on the store's decision to re-order that item. A perhaps more mundane example might be one of breaking a past commitment. The nature of an agent's situation is markedly changed if we include the fact that one possible course of action would preclude the fulfilment of a past promise.

Returning to the point at hand, two counterexamples will suffice to demonstrate that although the pressure is experienced as having its source external to the agent such experience cannot be understood simply in terms of features of the situation which are external to the agent, i.e. the context. First, suppose that Betty comes home to find her roommate asleep on the couch. She carefully puts a blanket over her roommate and goes to her room. As it turns out (and as Betty later discovers) the roommate is not asleep but has passed out because she has taken the wrong pills. The roommate is in grave danger and needs to be taken to the hospital immediately. Since Betty does not know that her friend has taken the wrong pills she cannot feel moved to help. Second, suppose that Betty happens to look in bathroom and sees the empty pill bottle on the sink. She realizes

³⁹ For the purposes of the example we can ignore the obvious evaluation involved in this 'fact'.

what her roommate has done. Suppose additionally that, unbeknownst to her, Betty is a brain in a vat and the events that occur before her are the horrifying products of her programmer's imagination. It seems clear in this case that, given the kind of person she is, helpful, caring and sympathetic, Betty will not only feel the pressure to assist but will endeavour to do so. Both of these cases straightforwardly show that the experience of the pressure, although it feels as though its origin is external to the agent, does not originate there. In both, it was Betty's 'take' on the situation that determined whether she experienced the pressure or not rather than the situation itself. In the first, she did not experience the pressure to help because she did not know about the pills and in the second, the pressure that she experienced was not a result of her context but of her programmed thoughts about her context (which, as it happened, were radically mistaken).

We can quite easily cast these examples in terms of obligation. In the first case it did not occur to Betty that she ought to help her roommate because she did not know that help was needed. Indeed, Betty may well have thought that she ought to be as quiet as possible. In the second case, Betty definitely thought that she ought to help even though her situation was radically different from the way in which she took it to be. Notice however that when construing the examples in this way we have moved from considering the experience of the pressure to considering the judgements that Betty makes. Indeed we could also go further and consider the judgements that we would make about Betty's predicament and her judgements.

The examples above show that the source of the pressure is not external. That is, although it feels as though some force is being applied on us from without, it is not and instead, it is a result of how we are 'taking' the context to be. Now as things stand, this

conclusion may be accepted without thinking that it tells us anything interesting about our obligations. All this shows, it might be suggested, is that the experienced pressure is not a good guide to our real obligations — the origin of the pressure might be ‘internal’ but this is not to say anything of the origin (ground) of obligation. In the terms of the examples above, it might be agreed that Betty didn’t experience the pressure and indeed that consequently she did not judge herself to have certain obligations. However, this is not to say that there is no sense in which she did in fact have those obligations — i.e. to save her roommate. If this were true it would spell serious trouble for the project as a whole. It would mean that there is a sense of obligation and a correspondingly, a whole slew of judgements of obligation, that are not related or only incidentally related to the experienced pressure. Given this worry, we will be in a much stronger position if we can say something about the origins of not just the pressure, but of our obligations. Recall the options, now in terms of obligation: the source of our obligations is either ‘internal’ or ‘external’.

It is worth noticing that these two positions do not quite match up with the ‘objectivist’ and ‘subjectivist’ traditions in meta-ethics. Importantly, the intuitionist tradition would count as both objectivist and internal. This mismatch limits the kinds of arguments that we can use to address the present problem. The subjectivist tradition has a long history of argument that would yield just the right kind of conclusion but it does so at the expense of certain important versions of the objectivist tradition. For instance, J. L. Mackie in *Ethics: Inventing Right and Wrong* gives us the argument from queerness in which he assumes that intuitionism is implausible and then proceeds to show that all

forms of objectivism reduce to intuitionism.⁴⁰ Since we are concerned with the internal/external issue as characterised above and have made no assumptions about the objectivity or subjectivity of moral judgement, this kind of argument is unavailable to us.

In what follows immediately below then I will introduce a distinction between two senses of 'ought' discussed by W. D. Ross. In the next section (B), we will examine the relationship between these senses and the origin or ground of obligation. In the final section (C) we will return to the sense of 'ought' in an attempt to reconcile the conclusions of B with Ross' two senses.

At the point in the *Foundations of Ethics* where Ross introduces his characterization of a situation, he also introduces a corresponding distinction between two senses of 'ought'. Let us call them respectively 'ought_o' and 'ought_s'. If an agent ought_o to perform an act then, on the basis of the objective element of the situation, that act "would in fact be the best possible fulfilment of the various *prima facie* obligations resting on the agent" (p.147)⁴¹. Correspondingly, the act which an agent ought_s to perform depends upon the agent's thoughts about the situation (the subjective element). That is, "the act which is morally suitable to them, i.e. the act which the agent, in view of his opinion about the situation, thinks will be the maximum fulfilment of obligation, will be in that respect right" (p.147). If the two elements of a situation coincide — that is, if

⁴⁰ "A special sort of intuition' is a lame answer, but it is the one to which the clear headed objectivist is compelled to resort." J. L. Mackie, *Ethics: Inventing Right and Wrong*. London: Penguin, 1977: 39.

⁴¹ It should be noted that this account includes Ross' own normative theory. The important point for us is the connection between the kind of obligation (or the sense of 'ought') and the objective element, not how the content of the obligation is determined. We shall return to consider *prima facie* obligations below (section C).

the agent's thoughts about the situation correspond to the situation as it is — then the act will be “completely right”.⁴²

Given that these two senses can, in the same situation, attach to two different acts — that is, Fred ought_o to do X and, at the same time, ought_s to do Y — Ross stresses that it is important to determine which is “ethically the more important”. The following example demonstrates why this is an issue which needs to be resolved. Suppose that Fred boards a train and takes the last seat just as a pregnant woman, loaded with grocery bags, steps into the same car. She looks tired presumably because she has been on her feet for the duration of her shopping expedition. Fred thinks that he really ought to give up his seat (‘it would be the right and proper thing to do’) but he does not because he was there first and he wants to continue reading his book. Suppose further that the woman is not tired but stressed. All day she has been treated as though she was incapable of doing anything of her own accord, now that she is ‘in this condition’. She has just scolded a worker at the supermarket for suggesting that she could not manage with her bags. In this case, Fred ought_s to have offered the woman his seat but, since she would have been offended, he ought_o not to have offered. In not doing what he ought_s to have done, he inadvertently did what he ought_o to have done. What makes this case problematic is that what Fred did was, by his own lights, wrong, however his doing so accidentally meant that the best result was achieved.

⁴² A second account of the various senses of ‘ought’ is given by W. K. Frankena in ‘Obligation and Ability’ (in Black, M. *Philosophical Analysis: A Collection of Essays*. Englewood Cliffs, NJ: Prentice Hall, 1963: pp.148-65.) I discuss these senses and their relation to Ross’ in the final section of this chapter.

In *The Right and the Good*⁴³ Ross solved this problem by saying that the act which an agent ought_o to perform is the right act, whereas the act which ought_s to be performed is the morally better act. So, as it happened, Fred did the right act even though it was “morally inferior”. We may think, as Ross came to think, that there is something unsatisfactory about this solution since it drives an awkward wedge between what is morally preferable or morally superior and what is right. Ross was uncomfortable with the possibility that an agent could do what, by his own lights, was wrong but end up doing the right thing. Indeed, in the *Foundations of Ethics* he gave up the solution he gave in *The Right and the Good*. The cited reason for his dissatisfaction and change of heart is H. A. Prichard’s paper, ‘Duty and Ignorance of Fact’.⁴⁴ It is to this paper that we now turn.

B. Obligation and Ignorance of Fact.

In ‘Duty and Ignorance of Fact’, Prichard is concerned with an issue very similar to that which we faced at the beginning of the chapter. Precisely how similar and in what ways will become clear as we proceed. The question that he considers is:

If a man has an obligation, i.e. a duty, to do some action, does the obligation depend on certain characteristics of the situation in which he is, or on certain characteristics of his thought about the situation? (p.18)⁴⁵

⁴³ Ross, W. D. *The Right and the Good* (1930). Indianapolis: Hackett, 1988.

⁴⁴ H. A. Prichard, ‘Duty and Ignorance of Fact’ in *Moral Obligation*. Oxford: Clarendon Press, 1949: 18-39.

⁴⁵ Prichard gives several alternative formulations of this question. Of particular note is the one given on p.28: “Does the obligation depend on the fact that the action would possess

As I remarked with respect to the occurrence of 'take' in our question, so Prichard calls the reader's attention to the vagueness of the word 'thought' in his and indicates that it will not influence proceedings at the outset⁴⁶. The two options presented in this question look to be identical to Ross' distinction between the objective and the subjective senses of obligation. That is, if we answer the question with the former alternative we might be saying that the man ought_o to do the action (and vice versa). Immediately, then, there is an issue. What does Prichard want to decide here? Surely these are just two different senses of 'ought' and whether the man's obligation depends on the context or his thoughts about the context will depend on which sense is intended. So, the man's obligation depends on the sense of 'obligation' being used.⁴⁷

The resolution of this problem, I think, lies in the fact that Ross and (presumably) Prichard take this paper to address, what we might call, 'Ross' problem': which of these senses is "ethically more important". However casting the question in this way looks to make the question a normative one. Indeed, more than this, it makes it a moral question. It does not appear that this is what concerns Prichard either: his arguments are not moral arguments. Moreover, if the paper were about answering a moral question then, given our project, the paper would not be of any use to us since we are not in the business of settling normative moral questions.

this character in a greater degree than any other, or does it depend on the fact that the man, after full consideration, thinks it likely that it would?"

⁴⁶ Prichard suggests this, so far as I can tell, for similar reasons to mine.

⁴⁷ These remarks, at least partly, anticipate objections made by Frankena and Ewing. See below.

What Prichard does wish to consider is the ground of obligation. This is suggested by the way in which the argument of the paper is presented and by the presence of 'depend' in the initial question. He wants to know what the obligation *depends upon*. There are various ways of stating this quest for 'the grounds of obligation': what makes it the case that the man has an obligation, what is the source of the obligation or what is the basis of the obligation. In each of them the grounds play some kind of justificatory role. Perhaps the best strategy here is to examine Prichard's arguments and to determine from this what kind of justification he sought. In any case, we should resist the temptation to align the alternatives in Prichard's question with Ross' senses of 'ought' or understand his question to be one which requires a moral judgement. It seems possible to decide the question of the ground of obligation without denying that there are different senses of 'ought'. That is, one might think that 'ought' could *mean* different things but think that there was only one kind of ground. (This is discussed further at the end of this section.)

Being at least partially clear about what Prichard is attempting in the paper gives us a clue as to the method that he will use. The two competing positions in the question are labelled the objective view and the subjective view and are taken to be the only two options (p.25).⁴⁸ By examining cases in which an agent, in a given situation, is actually taken to be under an obligation, we will, it is hoped, be able to decide in favour of one of

⁴⁸ A word about the use of 'objective' and 'subjective' here: It is advisable, I think, to take 'objective' and 'subjective' to be more akin to 'external' and 'internal' (i.e. as I used it with respect to the experienced pressure). There is a sense of this distinction that is related to truth-aptness that we would do well to avoid, partly because the issue of truth and moral judgements will come up later but perhaps mostly because it is possible to hold the subjective view and also hold that moral judgements are capable of being true or false. David Wiggins (see his *Need, Values, Truth* (3rd Ed.). Oxford, Clarendon Press, 1998.) and Michael Smith (*The Moral Problem*. Oxford: Blackwell, 1994.) are perhaps two good examples.

the views. Thus, given that there are only two alternatives both of which locate the ground of obligation differently, by finding cases where we would ordinarily think “without question” that there was an obligation to perform a certain action, we can check which of the views is satisfied.

The role of Prichard’s discussion of the general characterization of moral rules is designed to focus our attention on those characteristics of acts “which have led us to think them duties” (p.18). By developing a highly abstracted account of a moral rule we will have some general grasp of the kind of characteristics of acts and contexts that seem to issue in obligations. That is, we will be in a position to fill out, in a general way, the “certain characteristics” of the situation and perhaps of the agent’s thought about the situation. After some discussion he arrives at the following version:⁴⁹

⁴⁹ The initial form of the generalized moral rule is: “A man ought to do such an act or acts, i.e. bring about such a thing or things, as will cause a thing of kind *A* to assume a state of kind *X*.” (p.20) But since we are considering cases where we think there is an actual obligation this needs to be qualified. To use Prichard’s example, “in asserting that a man ought to support his indigent parents, we clearly do not mean that a man ought at *any* time to support his indigent parents.” (p.20) If we were considering, as it were, an ideal or ‘theoretical’ obligation (something which is generally or *in principle* obligatory) then the above statement might suffice — respecting one’s parents, for instance. When someone has an obligation we take it that the context is such that it contains a thing of kind *A* capable of assuming a state of kind *X* and that something that the agent can do (bring about) will cause *A* to assume state *X*. [The involvement of ‘can’ or ‘capable of bringing about’ in this latter condition is somewhat problematic since it could have different meanings. It would appear reasonable to suggest that how we understand ‘can’ here is connected to how we understand ‘ought’. Thus, since we are considering actual obligations we should understand ‘can’ to be something like ‘actually can’. These issues are discussed by Frankena in ‘Obligation and Ability’ (*op.cit.*.)] Also, since we cannot always be sure that these two conditions are satisfied in particular cases (for the latter condition, we can never be sure), Prichard further modifies the rule to one having a conditional form. [This last modification is the result of the intention that the general characterization express what we know of those features of acts which have led us to think them duties. But if we, at least sometimes, are in doubt then what we know will need to be expressed as a conditional.]

If the [context] in which a man is contains a thing of the kind *A* capable of having a state of the kind *X* produced in it, and if also it be such that one of the things which he can bring about directly will, if he brings it about, cause *A* to assume such a state, then he ought to bring about that thing. (p.23)

This statement of the generalized form of a moral rule is taken by Prichard to express the objective view — it gives the characteristics of the context mentioned in the antecedent as the grounds of obligation. However, it leads to a number of problems. “On this view that an obligation, if there be one, depends on certain features of the situation, we are driven to the extreme conclusion that, although we may have duties, we cannot know but can only believe that we have.” (p.24) Three further problems follow from this one: First, we can never do a duty because it is a duty — that is, because we know it to be what we ought to do. Second, we may have done things in the past that were our duty even though we did them believing them to be things we ought not to have done. Finally, we “may do some act which is in fact a duty, although in doing it [we] do not even suspect that it will have the effect which renders it a duty.” (p.24)

These problems, whilst not necessarily decisive, lead Prichard to turn to the subjective view since they all stem from the dependence of obligation on the facts of the context. After an extensive discussion of the immediate problems and responding adjustments to this view, Prichard concludes that the “fundamental difficulty with the subjective view “is that it represents the duty of doing some action as depending not on the fact that the action would have a certain character if we were to do it, but on our

thinking it likely that it would.” (p.31) Interestingly, this is simply a refusal to accept the subjective view.

After thus stating and clarifying the general form of a moral rule, and setting up initial problems with both views, we get to the heart of the paper. Prichard writes that “the first thing to do appears to consist in ascertaining which of the alternative views better corresponds with the thought of our ordinary life.” (p.28)

According to Prichard, there are several reasons why we might be tempted to answer the question in favour of the objective view. First, “we frequently think without question both that the situation contains some thing in a certain state, and also that some action which we could do would produce a change in it of a certain kind, and then think without question that we ought to do the action.”(p.28) That is, we seem to think that we have an obligation because of the facts of the context and the changes that our action would effect on that context. Second, “we often seek to change the mind of someone else about a duty by trying to convince him that he is mistaken about the facts; and, in doing so, we seem to imply that the question whether he has the duty depends on the nature of the facts.”(p.29)

Against this, Prichard gives a number of examples. One of these, suggested to him by R.G. Collingwood, is as follows:

Imagine that we are watching a car approaching along a road which we know forks, and of which one fork has, we know, just suffered a landslide, and that we have no idea which road the driver is intending to take, or whether he knows about the landslide. (p.29)

We would, Prichard rightly contends, say that we ought to warn the driver of the landslide. We are in this example supposing that the driver does actually intend to take one of the forks in particular and that he does either know about the landslide or not. Thus, there are facts about the situation that we do not know but which, if we did, might change our view of what ought to be done. So, for instance, if we knew definitively that the driver intended to take the safe fork we would not be obliged to warn the driver. Importantly, we would think that we ought to warn the driver in spite of our ignorance of the facts and indeed, irrespective of them. Thus, from this and his other examples Prichard generalizes that

the objective view is in direct conflict with all the numerous cases in which we think without question that we ought to do something which we are thinking of as of the nature of *an insurance* in the interest of someone else. (p.30)

In addition to this, Prichard suggests that the much of our ordinary thought *seems* to involve the objective view because of our tendency to think that something's being likely or probable is a "fact in nature". Thus, we tend to think that 'He has probably fainted' and 'Shouting will probably revive him' are statements representing certain facts called probabilities and that "an obligation to shout then arises from these probabilities". Prichard suggests however, that "it takes but a little reflection to realize that there are no such things as probabilities in nature. There cannot, e.g., be such a thing as the probability that someone has fainted, since either he has or he has not." (p.30) So, in the case above, even though we may think it likely that the driver does not know about the landslide either he does or he does not. We may think it probable to a certain degree that

the driver intends to take the dangerous fork, but again either he does or he does not.

Prichard continues:

No doubt it is extremely difficult to formulate the precise nature of the fact which we express, for instance, by the statement: 'X has probably fainted.' But at least we must allow that, whatever its precise nature may be, the fact must consist in our mind's being in a certain state or condition. And, once this is realized, it becomes obvious that most of our ordinary thought involves the subjective view. (p.30)

Finally, (referring to the second reason for thinking that our ordinary thought implies the objective view) Prichard points out that when we try to change someone's mind about a duty, we think that once we have made the attempt, whether or not the person has the duty will still depend on that person's opinion of the facts. Thus, suppose that two people, A and B, agree that they ought to go to war if their country has been economically compromised by another but A thinks that their county has been so compromised and B does not. Even after B has attempted to convince A of the facts, whether or not A will be obliged to go to war will depend on his (A's) opinion of the facts (p.30-1).⁵⁰

Now, although this is not the end of the account that Prichard gives, this is enough for us. Ross, in his account of Prichard's argument in the *Foundations of Ethics*, agrees:

⁵⁰ It should be noted that this example and the one given by Prichard are not particularly good ones. This is partly because one might insist, against them, that, for instance, A ought not to go to war irrespective of whether B has convinced him of the facts or not. My reason for including this point and an example will become apparent when we consider the second objection to Prichard's argument. There the insistence against Prichard will involve a return to the senses of 'ought', which, it seems, is part of the problem underlying the current example.

“This last part of Professor Prichard’s theory, while both true and important, is not necessary for the saving of the subjective view.” (p.156) Admittedly, Ross continues to give his own defence of the subjective view but it, like Prichard’s is not immediately relevant to our discussion. The part of Prichard’s discussion that I have omitted is a closer investigation of what an obligation is an obligation to do. There he introduces and defends the idea that the obligation is to “set oneself” — an idea closely related to willing.⁵¹ He goes on to conclude that we should understand obligations to be possessed by agents rather than understanding ‘obligatoriness’ or “ought-to-be-doneness” as a feature possessed by acts.

Let us return to the most important part of the above argument to be clear about what it shows. The example of the car approaching a fork in the road gives us an instance in which we have an obligation but where that obligation does not depend on the facts of the case. We have an obligation even though there are crucial facts about the context which we do not know and, moreover, realise that we do not know. Again, according to the example, we think that we have the obligation *without* knowing these facts. Thus, these facts could not be *grounding* the obligation.

There may be a question about why we would think that it is *these* facts which would ground the obligation in the first place. The short answer is that on the objective view if we did know them then we would think that they did ground the obligation (or, in the case of the driver knowing of the landslide, the lack of obligation). But more fully, we can see here the point of Prichard’s discussion of the general characterisation of a moral

⁵¹ Frankena (*op.cit.*, p.163) has an objection to this final step. But since it is related to Frankena’s senses of ‘ought’ and is not strictly relevant, I will not discuss it here. In the end however, it misses Prichard’s argument for the same reason as his other objection which is discussed at the end of this section and in section C.

rule. The rule in its final form is intended to give us the kinds of facts upon which the obligation depends. That is, on the objective view, the two kinds of facts on which the obligation depends are (i) that the context contains a thing of kind *A* which is capable of assuming a state *X* and (ii) that the agent is able to bring about the state *X* in *A*. In the example before us, where there is an obligation, these would become: (i) the driver does not know of the landslide and intends to take the unsafe fork rather than the safe one and (ii) that I can stop the car and inform the driver of the landslide so that he may take the safe fork. The problem, of course, arises because we do not know the first of these conditions and yet have the obligation. Thus, by first saying something in general about the kinds of facts which ground obligation on the objective view — for Prichard, those facts which we would ordinarily think grounded obligation — and then giving a number of cases where we understand there to be an obligation without knowing precisely those facts, Prichard's argument undermines the objective view.⁵²

Bringing all of the pieces together, we might put the major thread of argument in the following way (where 'O' stands for 'an agent's having an obligation in a certain context', 'F' for 'certain characteristics of the context' (i.e. facts), 'T' for 'certain characteristics of the man's thought about the context', 'F₁' for 'the context's containing a thing of kind *A* capable of assuming a state *X*' and 'F₂' for 'the agent's being able to bring about the state *X* in *A*'):

Either O depends on F or O depends on T (Assumption).

⁵² It might be objected at some point here that the relevant fact about this situation on which the obligation depends is that (i) is not known to be false — i.e. the fact that it is not known that the driver knows about the landslide and intends to take the safe fork. But this is a fact about my (our) thoughts about the context (namely what is known) and if our obligations depended on these facts they would seem to depend on one's thoughts about the context.

If O depends on F then O depends on F_1 and F_2 (Generalized Moral Rule Argument).

O does not depend on F_1 and F_2 (From the example).

Therefore, O does not depend on F.

Therefore, O depends on T.

The most natural response by the proponents of the objective view would be, I think, to challenge the kinds of facts included in the generalized account of a moral rule. Such a response would endeavour to show that Prichard's account of those characteristics of the context upon which the obligation depends, fails to do justice to the objective view. Now while there is no argument to show that such an attempt could not be successful, Prichard's position does seem to be a strong one. The benefits of his position are that it is general, it is restricted to actual obligations and it is intended to capture what we 'ordinarily think'. Any version of the objective view that was intended to challenge the one given would also need to satisfy these conditions. In addition, Prichard's remarks about probability seem to rule 'future facts' out of contention. We shall have more to say about this alternative below when examining criticisms that have been made of the argument.

Having said all this Prichard's argument is, I think, a little harder to follow than it need be. This is because he does not distinguish between the possible perspectives from which the judgement of obligation is made. Perspectives might seem to matter in the following way. We might think, in some cases, that an agent has an obligation of which he is unaware because there are certain things he does not know or that the agent thinks he has an obligation because he is unaware of certain things. Thus from an observer's

perspective we might know certain facts about the agent's context to which the agent is not privy. As a result, we might judge the agent's obligations in way that differs from that of the agent. The way in which Prichard's arguments make room for this kind of difference of perspective and judgement is not immediately obvious. Although this matter does not affect the argument, it does make it a little harder to appreciate. The amendment that will help to remove this opacity is the introduction of the distinction between direct and removed judgements into the equation. This distinction helps by specifying the perspective from which the agent is taken to have an obligation. Prichard's initial question is thus resolved into two:

- 1*. Does the obligation on the agent (as judged by the agent) depend on certain characteristics of the context in which he is, or on certain characteristics of his thought about the context? And,
- 2*. Does the obligation on the agent (as judged by a third party) depend on certain characteristics of the context in which agent is, or on certain characteristics of the third party's thought about the context?

This amendment allows us to take possible differences of perspective out of the issue. There are two things to notice. First, question 1* amounts to 'Does the obligation that the agent thinks he has depend on the context, or on his thought about the context?', and would seem to require the subjective view as its answer. It is quite plain that my thinking that I have an obligation is dependent upon my thoughts about my context and not on the context itself. That is, if I believe that I ought to do *X* this is because I believe that the context is such as to require my doing *X*. Second, this split does not privilege either the objective or subjective views. The separation of the two questions is a separation of the

perspective of the judge so that we might say it is a distinction within the subjective view. Thus, given both of these points, the question of interest is the second, whether the obligation depends on the judge's thoughts about the context or on the context itself.

This distinction is of most use, I think, in making sense of the subjective alternative although a few adjustments are required to Prichard's example to enable us to see how the point now goes. We have now to imagine that we are witnesses both to the car approaching the fork in the road and to the agent who is in a position to alter its course. Additionally we should suppose that we know that the agent is able to do something.⁵³ Finally we should suppose that we know that the agent knows of the landslide and that, like us, he does not know whether the driver is aware of it or which fork the driver's intends to take. In spite of this, again, we would judge the agent to have an obligation to stop the car.

However, this only goes part of the way to responding to the issue of perspectives. One may still object that there is a perspective from which a judge, who knows enough about the context, would be able to determine the agent's obligation on the basis of the facts. If someone knew, for instance, that the driver had heard about the landslide on the car radio and had already decided to take the safe fork then they would be under no obligation to stop the car. However, suppositions such as these rapidly approach a claim about obligations from a 'God's eye' perspective. If we knew all that there was to know then surely what we are obliged to do is based on what we know. But now we have missed the point of Prichard's project. When we have *actual* obligations, such as in the

⁵³ This, of course, will be harder to know than in the first person case since more things could go wrong in deciding whether someone else is capable of performing an action than in, for example, my knowing that I can perform some action.

initial example, we do not know all that there is to know. We must decide what it is our obligation to do on the basis of what we *do* know, that is, on the basis of our thoughts about the context. That there is or could be a perspective from which our obligations are different cannot effect our actual obligations in this context precisely because we lack the knowledge that we would have from that perspective.

We can now return to the unresolved question about the nature of 'grounds'. It seems relatively clear from the discussion to this point that the kind of grounds of obligation in which Prichard is interested are reflected, in his initial question, by the idea that the obligation *depends on* one of the two alternatives (say, A depends on B). It is natural to interpret this as, 'without B there is no A' — that is to say, B is a necessary condition for A. Thus, in terms of Prichard's account of the objective view: F_1 and F_2 are necessary conditions for an agent's having an obligation. The grounds of obligation look to be the conditions necessary for obligation. This however is not all. The final statement of the generalized moral rule expresses the two conditions of the objective view (F_1 and F_2) as being sufficient for the agent's having an obligation: If F_1 and F_2 then O. Hence, it would seem as though the targeted account also includes those facts, satisfying F_1 and F_2 , as sufficient conditions for obligation. We have reason then, to think that Prichard's concern is with both necessary and sufficient conditions for obligation.

But does the argument gel with this reading of 'grounds'? That is, given our understanding of the argument, should we understand what Prichard means by 'the grounds of obligation' to be 'necessary and sufficient conditions for obligation'? I believe so. Prichard's example is one where the agent has an obligation (according to us) to warn the driver of the car *irrespective* of whether the driver, in fact, knows of the landslide or

not and *irrespective* of which fork he, in fact, intends to take. If the facts about the driver were necessary for the agent's being obliged then we would not be in a position to determine if there was an obligation until these facts had been ascertained. Clearly this is not the case. If the facts about the driver were sufficient for the agent's being obliged then these facts would be enough to guarantee the agent's obligation. This too is problematic since we would still require that there was someone who was in a position to access those facts. Such an individual is, at least in the example, only hypothetical and yet we think that the agent's obligation is an actual obligation and not merely hypothetical. Hence, if the facts (namely, F_1 and F_2) were either necessary or sufficient for the obligation then the judge would need to have the right kind of (epistemic) relation to the facts either in order to judge there to be an obligation at all or in order for these facts to be enough to judge that there is an actual obligation. Prichard's example demonstrates the problem with these requirements by providing a case where we unequivocally judge there to be an actual obligation but where we do not have any (epistemic) relation to the facts, let alone the right one (i.e., knowledge). Thus, Prichard's denial that the objective view can provide the grounds for obligation amounts to a denial that the facts of a context are either necessary or sufficient for obligations that we actually have.

What is noticeable about the above discussion, and indeed the outline of the main argument of the paper presented above, is the absence of an account of the subjective view. Unlike the objective view, for which we are given an argued account, we are not given much clue about how the subjective view would develop. Prichard does, in the discussion of the problems of the subjective view, give us a little but he does not present

anything like a final attempt at the grounds of obligation as the subjective view would have them. In short, we are given an account of what the objective view entails, namely F_1 and F_2 , but comparatively little regarding the subjective view. We are led to the subjective view not through positive argument, as it were, but via the denial of one of the disjuncts of an exclusive disjunction. This, as Prichard points out, is a little discomfoting since the objective view seemed to have much appeal. It all means that a good deal of work is required to make sense of the subjective view. One way of seeing what follows in the next chapters is as the, undoubtedly small, beginnings of such a project.

Before moving on however, there are a number of objections that have been made to the arguments in 'Duty and Ignorance of Fact' that need to be addressed. Each of them, I believe, rests on a misunderstanding of the nature of the task which Prichard set himself. First, A. C. Ewing, in *The Definition of Good*, suggests that Prichard's difficulties, in part, arise because he "assumes" that probability cannot be objective⁵⁴. Against this Ewing holds that probability is an objective relation between certain data, "namely those available at the time the judgement is made" (p.128). The suggestion is that, given a particular set of data, the probability of X is determined relative to that set. The set of data available to the agent (and upon which the probability is determined) is simply the agent's view of the context. If what I know about the context is limited and if my obligations depend on probabilities based on what I know, then my obligations will depend on what I know about the context, i.e. on my view of the context. Of course, an observer, with a different set of data (i.e. different view of the context), might arrive at a different probabilities (or, indeed may be certain) and so may have a different opinion of

⁵⁴ Ewing, *The Definition of Good*. New York: Macmillan, 1947: footnote, p.125.

my obligations. The probability is calculated on the basis of a data set that is incomplete. We determine that a given event is likely, probable or improbable on the basis of what we know but, quite clearly, we do not know all that there is to know — that is why we think it probable or likely. This does not make probability ‘subjective’ — that is, for a given set of facts, there may be ‘real’ probabilities — but it does, as Prichard notes make the probability dependent on the agent’s view of the context.⁵⁵

Second, an objection made by both Ewing and Frankena, involves Prichard’s failure to consider other senses of ‘ought’. More specifically, Frankena’s objection is the following:

Prichard’s arguments ... refute the objective *view* of obligation but do not show that ‘ought’ is not actually and defensibly used in an objective *sense*. They seem to show this only because Prichard implicitly assumes ‘ought’ to have but one sense; as soon as it is seen that it has more it becomes plain that they show nothing of the sort. (p.156-7)

However on the interpretation above, Prichard is not concerned with the various senses of ‘ought’ but with the ground of obligation. As we have seen the objective and subjective views are views of the grounds of obligation and are to be distinguished from the senses of ‘obligation’. Recall that Ross’ problem, that which, in part, defined Prichard’s project, concerned which of the senses, objective or subjective, is more fundamental or “ethically the more important”. Consistent with Prichard’s comments, we have understood the search for the grounds of obligation to be the search for the more fundamental of the two senses. What does not follow from this is that once we have decided the question of the

⁵⁵ This issue is considered further in the next section.

grounds we have also denied one of the senses of 'obligation'. One of the senses could, it seems, be more fundamental than the other without the latter sense being obsolete.

The part of Frankena's objection that does have some force is the first: that Prichard's arguments do not show that 'ought' is not actually and defensibly used in an objective sense. On the interpretation I have given, this is an omission rather than a problem for Prichard. Part of the strength of the objective view was that it allowed us to explain the objective sense of 'ought' — 'we have obligations in a real sense *because* our actual obligations depend on the facts'. Once this grounding is taken away it is no longer clear what gives sense to our thinking that we have obligations in the objective sense. Prichard does not give us any help in this regard but then it is not strictly a part of the paper. Thus, some reconstruction is needed to make out a plausible position concerning the objective sense of 'ought'. Since it would be madness to deny that 'ought' is not actually used in the objective sense, the main issue in this reconstruction is whether this use is defensible.

C. 'Ought' and Judgements of Obligation.

To this point in the chapter we have examined the 'ground' of the experienced pressure and of obligation and found in both cases that the objective view is inadequate. As we have just seen however, there is a sense of 'ought' — i.e. a sense in which we might be said to have obligations — which can be understood as an objective sense. In the final section of this chapter we will consider this sense of 'ought' more closely to see whether it is a problem for Prichard and for the current project. This will involve examining a more detailed account of the senses of 'ought' given by W. K. Frankena in

his paper 'Obligation and Ability'.⁵⁶ There he maintains that there are three senses of 'ought' that are to be regarded as objective. We will examine each of these in turn.

These four (relevant⁵⁷) senses of 'ought' can be briefly described in the following way: the first sense, 'ought(1)', is that which "occurs in general statements or 'moral rules'"(p.150); the act which I ought(2) to do is the act which is objectively right; the act which I ought(3) to do is that which is objectively right relative to a certain body of evidence; and, the act which I ought(4) to do is that which I think is objectively right. Of these, Frankena thinks that senses (1), (2) and (3) are objective and that only sense (4) is subjective.⁵⁸

'Ought(1)' is introduced in the context of *prima facie* obligations, where he understands statements like 'Promises ought to be kept' and 'Pleasure ought to be promoted' to be expressing the idea that certain characteristics possess ought-implying tendencies. It is not that someone actually ought to do such and such but that generally speaking or ideally, these kinds of things ought to be done. Reflecting this, Frankena

⁵⁶ Frankena, W. K. 'Obligation and Ability' in Black, M. *Philosophical Analysis: A Collection of Essays*. Englewood Cliffs, NJ: Prentice Hall, 1963: pp.148-65. There are other discussions of the senses of 'ought' that we could examine, for instance, those of G. E. Moore, Bertrand Russell, C. D. Broad, A.C. Ewing and E. F. Carritt. However Frankena considers all of these authors in his typically thorough manner. It should be noted that, unlike Frankena, we are concerned only with the senses of 'ought' and not with the relation between 'ought' and 'can'.

⁵⁷ Frankena does briefly mention a fifth sense but suggests that it is "explicable in terms of other senses". (p.162)

⁵⁸ Frankena, *op.cit.*, p.159-60. The four senses are identified and related by four statements:

- (a) I ought(1) to do what I ought(2) to do, if I can, and if other things are equal.
- (b) I ought(2) to do what will most fully fulfil my obligations(1) of the actions which I can do if I choose.
- (c) I ought(3) to do what, relative to evidence *E*, I probably ought(2) to do.
- (d) I ought(4) to [try to] do what I think I ought(2) or ought(3) to do. (note, p.164)

takes “Virtue ought to be rewarded” to imply “If anyone can reward virtue, he ought to do so” (p.153), not that there is someone in particular who can right here and now. This idea follows from a distinction made by Moore between “rules of duty” and “ideal rules” where Frankena seems to understand rules of duty to be actual duties and ideal rules to be something akin to *prima facie* obligations. Frankena admits that there is some gloss on both Moore’s ideas and Ross’ account of *prima facie* obligation, however what he is aiming at is, at least partially, captured by the relationship between a *prima facie* and an actual obligation. An agent’s *prima facie* duty is only his actual duty “if no other moral features of the situation prevent its being an actual duty” (note, p.152). This means that talk of one’s general *prima facie* obligations will be talk in abstraction, lacking a context and, we might say, theoretical.

Now as we have seen Prichard is only concerned with our actual obligations, where, I take it, these are obligations that we in fact take ourselves to have — they are ‘actually judged to be obligations’.⁵⁹ More importantly, on this account ‘ought(1)’ judgements look to be judgements of value rather than judgements of obligation. In chapter one, judgements of value were characterised as those that tend to involve the predication of a value term.⁶⁰ Frankena’s examples all seem to be of this type, as does the

⁵⁹ Understanding ‘actual’ in this way does not prejudice either view. The aim, it will be remembered, was to determine whether the obligations that we take ourselves ordinarily to have depend on the objective or the subjective view — i.e. to understand the grounds of obligation in cases where we would ordinarily think “without question” that there was an obligation to perform a certain action.

⁶⁰ So judgments like ‘Virtue ought to be rewarded’ (and by implication, ‘If anyone can reward virtue, he ought do so’) are judgments of value in much the same way as ‘It would be good if virtue is rewarded’.

idea of ascribing ought-implying tendencies to certain characteristics.⁶¹ The upshot of this is that Frankena's 'ought(1)' may well be an objective sense of 'ought' but it is one that features in judgements of value and as such it is best understood as an expression of value. Given this it does not trouble either the account here or Prichard's argument.

In the case of Frankena's second sense of ought, the act which I 'ought(2)' to perform is the absolutely or objectively right act in precisely the same sense as Ross'. It is, as Ewing puts it, "the action which an omniscient and perfectly wise being would advise us to perform."⁶² Recalling Prichard, the peculiar feature of this sense of 'ought' is that it is impossible for anyone to know what they ought(2) to do. As a result no one can ever do as they ought(2) knowingly. Indeed, as E. F. Carritt points out⁶³, my doing as I ought(2) could only be an accident.⁶⁴

⁶¹ Ross' analogy between physical laws and *prima facie* obligations suggests this also. As Mandelbaum points out:

The analogy which Ross draws between physical laws and *prima facie* obligations (both being tendencies which will manifest themselves if nothing intervenes) shows that he is primarily interested in normative generalisations, and has not put forward his theory of *prima facie* obligations in an attempt to give an adequate phenomenological account of our moral judgements. (Note 41, p.318)

⁶² Ewing, A. C. *The Definition of Good*. New York: Macmillan, 1947: p118.

⁶³ Carritt, E. F. *The Theory of Morals*. London: Oxford University Press, 1928: 92. Carritt is, admittedly, discussing the absolutely right act. However, when considering the absolutely right act it would appear as though the distinction between 'the right act' and 'the act that I ought to perform', discussed by Ross (*The Right and the Good*. Oxford: Clarendon Press, 1930: 3-4) collapses.

⁶⁴ With these interpretations of Frankena's senses (1) and (2) we can understand the plausibility of (a) and (b) in note #22 above. Considering (a), it seems true that, generally speaking or 'in the abstract', I ought to do what I really or absolutely ought to do. In (b), if we understand ourselves to have ideal obligations or 'obligations in theory' then it seems plain that our real or absolute obligation, in *any* given context, is to do what will most fulfil them.

That this sense of 'ought' is a long way from our actual obligations (in Prichard's sense) is clear and so that it does not trouble Prichard's argument should also be clear. The fact that we cannot make judgements of this sort (accept accidentally) matters for the current project also. It will be recalled that the task was to explain a particular kind of judgement that we ordinarily make. If we cannot make these judgements — i.e. those involving this sense of ought — they cannot be judgements that we would ordinarily make.

However part of the task of this section is to determine whether these senses of ought are used defensibly and so more needs to be said. The important point here is that 'occurrences' of this sense of 'ought' are heavily loaded with theoretical baggage. We need to know everything there is to know and be able to morally assess the situation correctly in the light of this knowledge — as Ewing notes our advising supreme being needs to be omniscient *and* perfectly wise. I could only say what my absolute obligations are if I had decided about the one, true ordering of values (or was a perfect judge). What separates judgements involving this sense of 'ought' from the others is the justificatory machinery in the background. They do not involve simply having a better or stronger justification than any competitors but a way of securing the metaphysical reality of the normative claim. Given all of this it is hard to see uses of 'ought(2)' as strictly defensible. We do often say that we really ought to have done such and such but this is sometimes just an overstatement of the case or, more often, from the perspective of knowing more that we originally did.⁶⁵ As such this seems more like an occurrence of Frankena's 'ought(3)'.

⁶⁵ Cf. Blackburn, *Ruling Passions. op.cit.*, p.294ff.

Some work is required to become clear about 'ought(3)'. This can be achieved by clarifying the difference between 'ought(2)' and 'ought(3)'. 'Ought(2)' is not relative to the available evidence and so includes facts about the future. Thus we can modify (b) to read: I ought(2) to do what will, in fact, most fully fulfil my obligations(1). Presumably the point of distinguishing between 'ought(2)' and 'ought(3)' is to reflect the fact that we cannot know the future and hence cannot know what we ought(2) to do. Given that the facts which determine what I ought(3) to do are, in some sense, available — that is, not in principle unknowable —, it is possible for me to know what I ought(3) to do. That we use 'ought' in this sense seems clear. To use one of Prichard's examples, we might not expect someone to know that there were cars coming along a main road before we entered but we would think that they ought to slow down on the chance that there were. Here, the available evidence is that the road to be entered is a main road and, as such, there is a high probability that there are cars approaching the intersection.⁶⁶ Thus, the driver has an obligation(3) to slow down, on the available evidence. That is, the driver probably ought(2) to slow down.

There remains a problem about how we are to understand to what the evidence is relative. Because Frankena takes only 'ought(4)' to be a subjective sense of 'ought' (p.159-160), 'ought(3)' must be an objective sense. But then, we must understand the evidence *E* to be relative to what is in fact available and not relative to the agent or to any observers. If *E* were relative to me as the agent then, of course, 'ought(3)' would be subjective. Indeed, if it were relative to any observer whatsoever (except an omniscient one) then it also would also be subjective. Hence 'ought(3)' is relative to all possible

⁶⁶ 'High probability' here can be taken to involve past frequencies — i.e. in the past it has often been the case that there have been cars approaching the intersection.

evidence if we are to take it as an objective sense of 'ought'. But if this is the case, then ought(3) (virtually) reduces to ought(2).⁶⁷

We can see how this works if we complicate the traffic example above. If I am the driver contemplating entering the main road without being able to see whether there is traffic on it, then the evidence available to me is limited to the time of day, and hence the likely amount of traffic on the roads as well as my past experiences of the traffic volume on this particular road. Relative to this evidence there is a certain probability of there being a car approaching the intersection. Let us suppose that Fred is standing on an overpass that spans the main road. He can see that there is a car approaching the intersection in question but does not know whether it will turn off before reaching it. Relative to Fred's evidence the probability of the car being at the intersection when I enter is high. Alternatively, the driver of the car on the main road knows whether he intends to turn off but does not know whether his car is functioning well enough either to turn or to continue on. For each individual the probability of the car being at the intersection when I enter is different and each probability is dependent on the evidence available from that perspective. The only way to avoid this subjectivity is to require that

⁶⁷ My hesitation about the reduction concerns the possibility that there is some residual indeterminacy about knowing what will happen even if one knew all that there was to know about the present and past. The distinction between senses (2) and (3) rests entirely on this indeterminacy issue, if 'ought(3)' is to be taken as an objective sense. If I can know all that there is to know about the past and present and yet the future remains probable then we have reason to distinguish between senses (2) and (3). Either way it does not affect the argument. The barriers to my knowing all that could be known about the past and present are substantial enough to make it, at the very least, practically impossible for someone to know either what they ought(2) or ought(3) to do.

This also applies to the earlier discussion of whether probabilities are subjective or objective (or better, the extent to which probabilities depend on the agent's thoughts about the context). It might be that knowing all that there is to know does not eliminate all of the indeterminacy.

the available evidence be all possible evidence. Requiring all possible evidence ensures (something like) objectivity because it is no longer relative to a limited set of evidence and hence is not relative to a particular perspective. But now we can see the problem. If I possessed all the possible evidence then I would know that there was a car on the main road and that it would reach the intersection as I was entering. I might even know that my causing an accident here and now would initiate a sequence of events preferable to the alternative.⁶⁸ I would then know what I ought(2) to do.

The conclusion then, is that if ought(3) is to be taken as an objective sense of 'ought' it suffers from the same drawbacks as 'ought(2)' — it is impossible (if not in principle, then practically) for me to know what I ought(3) to do. If, on the other hand, we take 'ought(3)' to be a subjective sense, then whilst there are many versions of what I ought(3) to do — corresponding to perspectives on the context — they will all be subjective because they are relative to the evidence. On this view — taking ought(3) to be a subjective sense —, 'ought(4)' is a quite specific version of 'ought(3)', namely, my obligation(4) is what I ought(3) to do relative to the evidence which I have about my context. Understanding 'ought(3)' in this latter, subjective way is to be preferred because it better captures the ordinary sense it was designed to reflect, namely the perspective which an observer might have on an agent's context, especially when it differs from that of the agent himself. Taking 'ought(3)' in this way has the additional advantage of allowing the reintroduction of our distinction between direct and removed judgements — when I make a direct judgement of obligation I express it using 'ought(4)' and when

⁶⁸ These last two claims ignore the indeterminacy issues mentioned above.

someone else judges that I have an obligation (a removed judgement) they express it using 'ought(3)'.

In sum, 'ought(1)' looks to be an objective sense of 'ought' but judgements in which it occurs are better thought of as judgements of value and should be distinguished from our actual obligations. Second, it is unclear that we can actually or legitimately make judgements involving 'ought(2)'. This is largely because of the steep requirements associated with something being objectively or absolutely right. Finally, Frankena's third objective sense of 'ought', 'ought(3)' looks not to be an objective sense but a sense relative to a particular set of evidence. Judgements in which it occurs are perhaps good examples of removed judgements of obligation.

This chapter provides the first step in the argument of the thesis. We have considered here the objective view and the subjective view (suitably understood) as possible sources of the experienced pressure and of our obligations generally. We found quite quickly that the 'pressure' did not originate independently of us but noticed that in spite of this our obligations still might. Further argument was required to show that the ground of our obligations did not lie with the objective view. This was provided by Prichard in 'Duty and Ignorance of Fact'. Overall, the important focus arising out of the argument of this chapter is on the subjective view — one way of seeing what is to follow is as an exploration of this view, subject of course to the constraints and aims outlined in chapter one.

CHAPTER 3: THE PARTICULAR AND THE GENERAL.

“But up to what point and to what extent a man must deviate before he becomes blameworthy it is not easy to determine by reasoning, any more than anything else that is perceived by the senses; such things depend on particular facts, and the decision rests with perception.”

(Aristotle, *Nicomachean Ethics*, 1109b19-24)

A. Explaining an Agent's 'Take'.

The conclusion from the last chapter which has the most immediate significance for the project is that in order to understand the origins of the experienced pressure accompanying moral judgements of obligation we must look to the judge's thoughts about the context — the agent's 'take'. That is, there is something internal to the agent's 'take' which is responsible for the experienced pressure and hence, for the judgement of obligation. As we shall see in this chapter, there are two ways of approaching the explanation of this responsibility — of saying what, in the agent's 'take' is responsible for the pressure and hence the judgement. One is by appealing to general features of the agent's 'take' which operate in each case and the other is by appealing to particular features of the 'take' and the agent's ability to 'construct' such 'takes'. Thus, in this chapter, after becoming clearer about what constitutes a 'take', we will attempt to decide between these two routes of explanation.

1. Preliminaries.

There are three matters regarding an agent's 'take' that are usefully noted before we proceed.

(1) The main motivation for calling the agent's view of the context a 'take' is so as to not privilege a position in the dispute between cognitivist and non-cognitivists. The issue between cognitivists and non-cognitivists is most broadly characterized as a dispute about the possibility of moral knowledge. Cognitivists claim that moral judgements⁶⁹ like 'You acted wrongly in stealing the money' can be known to be true in the same way that 'You stole the money' can be known to be true. There is, however, the possibility of confusion with this characterisation: Is the issue an epistemological or a metaphysical one? That is, are we considering the possibility of our *knowing* the truth of a moral statement or are we concerned with there being a truth-value for a moral judgement. Clearly, I can believe that moral judgements are either true or false but think that we have no way of knowing which, in particular instances. The dispute is, I think, generally characterized as the metaphysical one. That is, non-cognitivism is the thesis that moral judgements do not have truth-values. Thus, Kant and Moore are cognitivists whilst Hume, Hare and Ayer are non-cognitivists.

The reason that it would be inappropriate for the account to privilege either of these positions at this point is simple. We are investigating the very subject of this dispute, the metaphysics of judgements of obligation. Whether and how it is *possible* for moral judgements to have truth values *might* turn out to be a consequence of the account

⁶⁹ It should be noted that, for the purposes of describing this dispute, I have reverted to the 'object' way of understanding moral judgement. For the next few paragraphs *only* 'moral judgement' should be understood not as the act of judging but 'what is judged'.

but, since we have yet to reach the end, the arguments thus far cannot privilege this issue. Still, this is not enough. The target, as it was set up in chapter one, was not primarily one concerning truth but what goes on when we make moral judgements. Thus, in searching for what made the act of moral judging possible we were looking for the configuration of psychological states that were responsible for these judgements. In this regard we might take the distinction between cognitivism and non-cognitivism to be one concerning, as we had it in chapter one, the psychological states that are responsible for the judging. Thus the cognitivist would explain moral judging as involving primarily 'cognitive' states like belief, knowledge and perception whereas the non-cognitivist would argue that such judgements depend primarily upon 'non-cognitive' or conative states like desires and attitudes. Of course putting the matter like this keeps truth close by, but it does mean that our immediate attention need not be the possibility of truth for moral judgements⁷⁰. Thus, we can, I think, proceed with the project allowing the more general distinction between so-called cognitive and non-cognitive states with the understanding that the possibility of the truth of moral judgements will, in part, resolve itself once the account is developed.

(2) It is important not to ignore the role that the context plays in the agent's 'take'. The arguments of the previous chapter do not show that the context is irrelevant, only that it alone does not explain either the experience of the pressure or the obligation. It is the

⁷⁰ Typically, if the account of what goes on when we make a moral judgement relied heavily on having certain beliefs or knowledge we might expect it to come down in favour of the claim that moral judgements were capable of being true or false. The reverse expectation (that if the making of moral judgements are largely explained by reference to attitudes and desires then moral judgements are not capable of being true or false) might also be typical although there do seem to be clear exceptions. Michael Smith's account (*The Moral Problem*. Oxford: Blackwell, 1994.) is one such exception.

agent's *view* of the context, the agent's 'take', that matters. Indeed it is the *view* of the context that is necessary if we are to explain the experience of the pressure.

(3) By investigating briefly the ways in which an agent's 'take' might differ from the context we will illuminate the point above and also observe some of the complexities surrounding the distinction between direct and removed judgements. There would appear to be two main ways in which an agent's understanding of the situation confronting him differs from his actual situation (his context): (a) as a result of ignorance and (b) as a result of interpretation. Correspondingly we could describe these divergences in terms of the relationship between the agent and particular features of the context. There might be features of a context:

- (i) of which one could not have been aware;
- (ii) of which one could have been aware but was not (for various kinds of reasons);
- and (iii), of which an agent was aware but did not figure, or figure prominently, in her 'take'.

Clearly there is room for slippage between these categories. We might think in some cases that although it was not strictly impossible for an agent to have been aware of feature F it was practically impossible.

Perhaps the most interesting point about these distinctions is how they become involved in removed moral judgements. Notice that in direct judgements none of these distinctions apply. When we make a removed moral judgement (about what someone else ought to do or about what we ought to have done) we, either implicitly or explicitly, consider these features and make judgments about them. Judgements of responsibility,

culpability and the like will all turn on how we, as observers, understand not only the situation, but the agent's relationship to the situation. Sometimes not knowing is an excuse and sometimes it is not — we might judge that although the agent did not know, he ought to have. Similarly, with respect to the third kind, we might judge that an agent ought to have paid more attention to a particular feature or that the feature did not figure as strongly in the agent's thought about the situation as it ought to have. At other times features might be given too much emphasis, they ought to be downplayed or ignored — skin colour might often be one such feature.

There are, I think, two related lessons in all of this. First, there seems a clear sense in which our removed moral judgements are parasitic on these 'sub-judgements'. Given this, an analysis of removed judgements would, quite quickly, become very complex. Second, the complexity of removed moral judgements stems from our 'take' on the context, from our view of the agent's take on the context and from our (often) moral judgements about the agent's relationship to the context as we see it. Thus the interplay of the features of the context and the stances of the various judges is similarly involved in removed judgements as it was in direct judgements.

2. The Agent's 'take'.

We have shown, in our first step (the discussions of chapter two), that the experience of the pressure is to be explained in terms of the agent's 'take' on the context rather than simply in terms of the context itself (the 'subjective view'). Our next step is to endeavour to understand what 'goes on' with the agent's 'take' such that the pressure is experienced. That is, the question to which we now turn is 'What are the features of the

agent's 'take' that will explain (i) the fact that the pressure is experienced at all and (ii), particularly, that it is experienced as objective (arising externally to the agent)?'

In order to proceed it is important to focus on direct moral judgements and more specifically, general features internal to an agent's making of a direct judgement of obligation. This is a consequence of the method outlined in the first chapter, the conclusions of the last chapter and the remarks above concerning ignorance and interpretation. The first of the lessons above should give us enough reason to focus on direct rather than removed judgements. If we become embroiled in removed judgements we will constantly need to disentangle the perspectives of third parties from that of the agent.

It should not come as a surprise that the focus will also be on the general features *internal* to the agent's direct judging. Recall that the kind of account we are seeking is one that explains what it is that makes judgements of moral obligation possible. Now that we have seen that the experience accompanying these judgements originates from the agent's 'take' rather than from the context, it is to the features of the agent's 'take' that we must look for further explanation. Thus, in the terms of the first chapter, we must look to the components and structure of the agent's initial psychological states. Our 'materials' are the psychological precursors to the judgement. The kind of explanation that we are seeking is also related to this focus. In looking to explain what makes judgements of moral obligation possible, the interest was in those conditions that make the *judging* possible. Since the judging is a psychological phenomenon it is to be expected that our explanatory route will take us, at least, through the initial psychological states.

If we were to concentrate primarily on the ‘external’ features surrounding the judgement we might be able to develop a correlation between the judging and these features but we would not have made progress in giving the kind of explanation with which we are interested. By ‘external’ here I mean those facts about the context and about the agent’s behaviour accompanying the judgement. These kinds of facts could not be the primary focus here and for the project outlined in chapter one since they would only tell us about the relationship between the content of the judgement, the content and the agent’s behaviour, not about the nature of judging itself.⁷¹

One important consequence of all of this is that we can avail ourselves of a notion of right action when considering an agent’s direct moral judgements⁷². This, after all, is how whatever judgements the agent makes will appear, given that they are direct. The notion of right action here comes from within the agent’s own ‘normative theory’. What kinds of things *are* judged to be right and when a moral judgement is called for are beyond the scope of our concern — this is a consequence of the non-normativity requirement of chapter one. The current point is that we can, given our aims, consider any context/agent pair such that a moral judgement that the agent makes is one where the agent thinks that he morally ought to do such and such. That is, given that Fred understands certain things to be moral issues and also has views within this about what is

⁷¹ This is a similar point (at least in intent) to that made by Blackburn in his general objection to response-dependent accounts and the ‘Canberra Planners’. See *Ruling Passions*, pp.104-121.

⁷² An alternative to this is to consider where agents ‘go wrong’ — that is, make the wrong moral judgement or behave (judge) immorally. However, there are many ways in which an agent might go wrong (as the discussion about ignorance (above) suggests) both according to him and according to us and, in any case, focussing on such cases would require us to bring our own removed normative judgements to bear on the agent’s situations and judgements.

right and wrong, good and bad, we can look at cases where Fred judges that he ought to do *X*. There is no violation of non-normativity here since we have not stipulated which kinds of cases and what kinds of actions within these, are right. Again, for reasons of simplicity we have limited our considerations to *direct* judgements. It is not *our* (removed) judgement of right action but right action as understood in the agent's terms.⁷³

Our position, then, is as follows: An agent in a context (given, say, by properties $X_1, \dots, X_n$) constructs/develops a 'take' on that context which issues in (gives rise to) a pressure, experienced as objective and which accompanies the judgement that he ought to do *A*. It should be noted that this characterization of the context is quite problematic in the light of the discussion of 'a situation' in chapter three.⁷⁴ So long as this is clear it will not matter. However, more does need to be said about a 'take'. Although a crucial feature of an agent's 'take' is that it contains a subset of the properties of the context determined by, among other things, what the agent thinks is relevant and what the agent is able to detect, this is not all. It will also include what the agent thinks to be the relevant history, any assessments (presumably probabilistic) of the future and perhaps some weighting or emphasis on properties and relations between properties. This last feature is intended to capture the idea that agents tend to find some properties more salient or more 'in focus'

⁷³ It is interesting to contrast the notion of right action which the method of chapter one will allow with that utilised by Aristotle. The standard interpretation of Aristotle might hold that the right action is that which the man of practical wisdom would perform. Now if practical wisdom is given a significant amount of normative content then so too is right action. If, on the other hand, possessing practical wisdom is taken to be largely structural matter then my claim can be seen as close to the Aristotelian one. We will return to these issues in chapters five and six.

⁷⁴ Recalling that discussion, the problem here is the suggestion that the context is *given* by the set of properties ($X_1, \dots, X_n$). This, of course, is another 'take' — perhaps ours, perhaps one that many would agree about but a 'take' nevertheless.

than others. At one extreme, the bear ten yards away is likely to dominate my view of the context rather than the fly on my dinner. At the other, I may have 'seen' that the car was red and be able to describe it as such, even though its being red did not initially occur to me.⁷⁵

The first step that we will take in the attempt to determine the features of agent's 'take' that will explain the pressure is to distinguish particular features from general features. That is, presumably some of the features of an agent's 'take' will be related directly to the current situation while others will be applicable to or utilised across various situations and across judgements. So, the modified form of our question is: Are the features of the agent's 'take' that explain the pressure general or particular features? This question is posed, of course, in the context of the conditions laid down in chapter one and this dictates the kind of explanation that we require.

In what follows we will focus almost exclusively on rules as exemplars of the 'generalist' option. We will not consider other candidate general features and in particular general sentiments for a number of reasons. First, for the most part the discussions that will follow concerning rules and Aristotle's minor premise do I think apply more generally. That is, even if there is a generally characterizable sentiment involved in all judgements of obligation, some explanation of how agent's come to view contexts in ways that 'trigger' this sentiment will still matter. Just as we might think of the rule and the application of the rule in a particular case, so we might also think of the general sentiment and a situation in which it is 'activated'. Second, in case of general

⁷⁵ For the most part, my 'take' will consist of those features of which I am conscious. There may be some occasions where a feature influences how I see a context (and how I judge or act in that context) even though I am not aware of it doing so, but these, it seems are the exception rather than the rule.

sentiments, the main problem that is of consequence here is the explanation of the 'objective character' of the pressure. We will spend a good deal of time exploring this issue in chapters five and six.

B. Rules.

The agent knows that situations with, say, properties X_i and X_{i+1} are such that he ought to do A. That is, the agent possesses a rule or principle, the knowledge of which, when combined with the instantiation of the properties mentioned in the rule generates the pressure and thus, the judgement. The idea here is that such rules are the product of reason and are known (or deducible) by rational beings. Of course there are many variations and nuances that could be (and, indeed have been) suggested. These rules might be taken as 'dictates of rationality', constitutive of what it is for someone to be rational and laws of our nature. Alternatively, they might be a part of what it means to be a rational being who must interact and depend on others members of a community. There might be only one rule (monism), several rules or many rules (pluralism). Things are simplest in the case of monism. As Jonathan Dancy puts it, "if there is only one principle, then perhaps only one property needs to be thought of as morally relevant; the question whether an action is our duty will depend only on the degree of presence or absence of that property at the end of the day."⁷⁶ For pluralism there are several (or many) rules each of which can be understood to pair properties with obligations. Prichard's account of the general form of a moral rule will serve as an example:

⁷⁶ Jonathan Dancy, 'Ethical Particularism and Morally Relevant Properties', *Mind* 92 (1983): p.531.

If the situation in which a man is contains a thing of the kind A capable of having a state of the kind X produced in it, and if also it be such that one of the things which he can bring about directly will, if he brings it about, cause A to assume such a state, then he ought to bring about that thing.⁷⁷

Clearly, the number of ways of filling out A and X will correspond to the number of rules required for a complete theory.

How could an account of this kind explain the relation between the agent's 'take', the pressure and the judgement? Each rule is designed to isolate features of a context and by doing so dictate the agent's judgement or action. So long as the feature(s) of the context mentioned in the rule is(are) present in the agent's 'take', the rule will guide the agent in determining what ought to be done. For pluralism we can offer the following additional kind of case. Let us suppose that the context has as one of its properties X_j of which the agent is unaware — X_j does not feature in the agent's 'take'. Let us further suppose that the agent possesses another rule which states that when properties X_j , X_i and X_{i+1} are present he ought to do B . If the agent were aware of X_j he would judge that he ought to do B rather than A . Thus, the agent's 'take' matters because it determines which of the agent's rules or principles is applied in the particular case.

The pressure on this view is explained via the status that is given to the rules as rules of morality or of rationality. There are clearly many ways of cashing this out. In doing so the rules, as commands or imperatives, are given the authority to determine

⁷⁷ Prichard, 'Duty and Ignorance of Fact', *op.cit.*, p.23. Citing Prichard here is not meant to imply that he was a pluralist. Also, using 'situation' here instead of 'context' is intended to incorporate the arguments of section C of the last chapter. That is, A and X represent properties which feature in the agent's 'take' rather than (necessarily) features of the context.

judgements and action. For these kinds of accounts, the way in which the rules come to have the authority that they do plays the primary role of explaining the pressure. The idea being that it is the authority of the moral or rational requirements which is the source of the pressure. For this kind of account, it is the rule that is the source of the pressure rather than the agent's 'take'. The 'take' only determines which rule is 'activated'.

There are two points to make here. First, at the level of the phenomena, it was noticed that the pressure felt as though it originated 'externally' to the agent. However, by placing the emphasis on the rule as the source of the pressure we are not well placed to explain the 'externality' of the experience — the rule and its authority are features of me not the context. Let us suppose that the reason the rules have the authority to dictate my behaviour is connected to their being rules of rationality. That is, I am bound by them because I am a rational being. Now presumably my rationality is operating, as it were, in the background, giving the rules authority and making them compelling. This does not seem to explain how it is the situation which, in some way, is compelling. If the rule provides the pressure then I will need to be aware of the rule but not aware of it or its authority as something peculiar to or originating specifically from me. More is required if we are to explain how it is that my possession of the rule allows me to find a situation as demanding something of me.⁷⁸

Second, even if we can delineate a rule or a set of rules there will still be a question concerning how they are applied in particular contexts. This application of rules involves both an ability to recognise that certain features of a context fall under the purview of a certain rule (or that judging or acting in a certain way is judging or acting in

⁷⁸ There is a very strong parallel between this problem and that faced by the non-cognitivist in chapter five. The solution that is offered in chapter six will work here also.

accord with the rule) and an ability to recognise those features in the first place. It would not be much good if I knew the complete set of moral rules (suppose, say, they were complex formulae relating molecular structure) but could neither recognise the features mentioned in the rules nor determine that the rules applied even given my recognition of the relevant features⁷⁹. Alternatively, being able to apply the appropriate rule at the appropriate time (even according to the agent) will involve both having the right ‘take’ — where ‘right’ here means that which ensures that the appropriate rule is relevant — and recognizing that the rule applies to this ‘take’ or, from the agent’s perspective, to this situation⁸⁰.

In cases where the content of the rule is general or broadly specified it is easy to see that this extra step is required and, moreover, that it is crucial. How one follows the rule ‘Help those in need’ will depend on what counts as helping and when someone is in need. In cases where the content of the rule is highly specified then, although the range of interpretations might be much smaller (and the potential need for a greater number of rules much greater), there will still be a need to account for the ability to follow/apply the rule.⁸¹

⁷⁹ This latter point gets us immediately into Wittgenstein’s rule-following discussion. It seems hard to see how someone could have the rule and the property mentioned by the rule but not make the judgement. It seems automatic and compelling in just the same way as going on in the ‘add 2’ case. See below.

⁸⁰ How to understand ‘appropriate’ here is part of the issue at hand. We might easily be able to imagine a case where the rule/context pairing is obvious but even this might require an extra step.

⁸¹ I intend this discussion to cover the distinction between general and universal rules. This is nicely discussed both generally and with respect to Aristotle by Martha Nussbaum in ‘The Discernment of Perception: An Aristotelian Conception of Private and Public Rationality’ (reprinted in *Love’s Knowledge*. New York: Oxford University Press, 1990: 54-105.) See especially, pp.67-8. Moreover, the entirety of section two (The Priority of

At this point we are confronted by Wittgenstein's rule-following discussions. The discussion of Wittgenstein's remarks about following a rule that bears the closest relation to the matters before us, is McDowell's paper 'Non-Cognitivism and Rule-Following' and Blackburn's reply 'Reply: Rule-Following and Moral Realism'⁸². After outlining the general argument as it applies here, I shall briefly consider their discussion (again, primarily as it applies here).

One way of seeing Wittgenstein's project in the *Philosophical Investigations* is as an attempt to undermine philosophical theories by revealing their un-stated and unjustified commitments.⁸³ In the case we are interested in, these commitments are to a particular picture of the operation of the mind. Thus, the kind of theory which suggests that the meanings of words can be laid down in advance, once and for all, is challenged by examining the consequences of such a theory for how we are to think about the 'processes' of the mind. In particular, the idea is that if meanings can be given as something like rules, then when someone understands (knows, grasps) the meaning of a word, they have understood the rule for its meaning. In other words, if I know the meaning of a word then my knowledge of the rule will be manifest in my behaviour. What, then, explains my consistent and continual use of the term is my grasp of this fixed rule of meaning. Now if these rules of meaning exhaust what I know when I know the meaning then they must guarantee how my knowledge of the rule is reflected in my

Perception) of this paper is closely related to the position which I will adopt in the discussion below.

⁸² Both in S. Holtzman and C. Leich (eds.), *Wittgenstein: To Follow a Rule*. London: Routledge and Kegan Paul, 1981: 141-187.

⁸³ Wittgenstein, L. *Philosophical Investigations*. Oxford: Blackwell, 1958.

behaviour when I use words correctly. (The rules of meaning are intended to be exhaustive, they attempt to say what meaning *is*.) As McDowell puts it, this picture is captured by the thought that “the grasp of the rule is a matter of having one’s mental wheels engaged with an independently traceable rail”⁸⁴. This is well illustrated by the example at §1 of the five red apples. The word ‘five’, for the person who knows what it means, ‘tells’ them to perform a series of tasks — counting, etc. — in a certain way. If the meaning of a word is to be understood in this way then this is the picture of the mind for someone who understands what a word means — the rails are in place. The question then becomes ‘what kind of evidence could there be for this picture of the mind?’.

In this context the discussion of extending the series according to the rule ‘add 2’ at §185 is intended to show that the picture outruns the evidence. The pupil has learnt the natural numbers and been tested with various series up to 1000. When the pupil is asked to extend the series ‘+2’ beyond 1000 he does so by writing 1004, 1008, 1012. Importantly, the pupil thinks, when asked, that he was going on in the same way, as he was meant to. As McDowell points out, this imagined scenario highlights the difference between the limited amount of behaviour which we could have as evidence for the pupil’s ability to ‘go on in accord with the rule’ and the unlimited behaviour which the rule, as capturing the meaning of the instruction, is supposed to dictate. Thus, “the pictured state always transcends any grounds there may be for postulating it.” (p.147)

Seeing things in this light makes McDowell’s worry about taking the rule-following discussion as supporting a skeptical conclusion about our knowledge of other minds (or about the possibility of giving an account which relied upon rules)

⁸⁴ McDowell, *op.cit.*, p.146.

understandable. But as he points out “the argument aims, not to undermine [our confidence that another’s (or our own) behaviour will not come adrift], but to change our conception of its ground and nature.” (p.148) As I have set up Wittgenstein’s point, he is concerned to undermine a particular way of understanding what someone grasps when they know the meaning of a word. And this, in turn, is done in an effort to show that an account of meaning which is based on rules alone is inadequate because it is incomplete. Again, the inadequacy is illuminated by recognising that such an exclusive appeal leads to a picture of what is to understand a rule which ‘transcends any grounds there may be for postulating it’.⁸⁵

After responding to this possible objection McDowell then endeavours, with the help of a quotation from Stanley Cavell, to give an account of what underwrites our confidence that our own behaviour will not “come adrift”. This marks the move to McDowell’s positive suggestions and here is not the place to consider them.

⁸⁵ It is interesting to notice two parallels to this argument. First, in the Preface of the *Critique of Judgement* ((1790) trans. J. C. Meredith, Oxford: Clarendon Press, 1928.) Kant justifies the claim that the principle peculiar to the faculty of judgement cannot be “derived from *a priori* concepts, seeing that these are the property of the understanding, and judgement is only directed to their application. It has, therefore, itself to furnish a concept ... which it can itself employ as a rule only — but not as an objective rule to which it can adapt its judgement, because, for that, another faculty of judgement would again be required to enable us to decide whether the case was one for the application of the rule or not.” (p.5, Meredith edition)

The second parallel is with what is known as ‘the third man’ argument. This is particularly striking after considering McDowell’s discussion of a postulated mediating state.

We get a nice summary of the overall point of this discussion, as it applies to our problem, in the following dialogue with the interlocutor at §186:

“What you are saying, then, comes to this: a new insight — intuition — is needed at every step to carry out the order ‘+ n’ correctly.” — To carry it out correctly! How is it decided what is the right step to take at any particular stage? — “The right step is the one that accords with the order — as it was *meant*.” — So when you gave the order ‘+ 2’ you meant that he was to write 1002 after 1000 — and did you also mean that he should write 1868 after 1866, and 100036 after 100034, and so on — an infinite number of such propositions? — “No: what I meant was, that he should write the next but one number after *every* number that he wrote; and from this all those propositions follow in turn.” — But this is just what is in question: what, at any stage, does follow from that sentence. Or, again, what, at any stage we are to call “being in accord” with that sentence (and with the *mean*-ing you then put into the sentence — whatever that may have consisted in). It would almost be more correct to say, not that an intuition was needed at every stage, but that a new decision was needed at every stage.

In the paper mentioned above, Blackburn objects to McDowell’s use of Wittgenstein on two counts: he questions the appropriateness of this extension of the rule-following remarks and whether, even with these extensions, Wittgenstein can be taken as deciding in favour of the cognitivist. The latter is certainly the more significant of the two. Essentially, Blackburn is skeptical about McDowell’s positive proposals

concerning the ground of our confidence that our behaviour will not come adrift. It is this part of McDowell's paper where he attempts to show that, consonant with the Wittgensteinian project, the grounds of our confidence generate what is required for cognitivism. My claims for the rule-following remarks are not nearly so large, and certainly do not involve enlisting them to decide, either way, between cognitivism and non-cognitivism. In any case these criticisms involve McDowell's positive suggestions and, as was mentioned above, discussing them at this point would be an unnecessary deviation. I will, however, make a few remarks about McDowell's proposals at the end of this section.

The former of Blackburn's objections is of more consequence here and it serves well to illuminate my point. His complaint is that Wittgenstein is only concerned with rules which we follow as "a matter of course" and "where there is no element of inspiration or any 'hearkening' or special sensitivity".

The whole stress in Wittgenstein is on the automatic and compelling nature of rule-following. The mental life of one who refuses to compute as we do after the standard training in arithmetic may be indescribable, and it may be quite literally unimaginable.⁸⁶

This is opposed to the situation in ethics, where the mental life of someone who differs from us *is* not only imaginable but is perhaps commonplace.

It seems difficult to disagree with this last point, however I am not sure that Wittgenstein's remarks quite play the role that Blackburn emphasizes. I would suggest that Wittgenstein's concentration on those cases where 'everything is a matter of course'

⁸⁶ Blackburn, 'Reply: Rule-Following and Moral Realism.' *op.cit.*, p.170.

reflect, not a limit on the kinds of rules to which the considerations apply but a demonstration of the breadth of the concern. Although Wittgenstein's main interest lies with accounts of meaning, the focus on rules and the picture of their operation is general enough to be utilised with respect to any rules. If the 'standard' picture of the operation and role of rules in philosophical accounts can be undermined in the 'easiest' and least controversial cases then the power of the argument is strengthened.

Two kinds of cases are helpful in the ethical case. First, imagine a case where the precise set of moral rules *was* settled. This would not mean that their application was settled. There would still be a residual explanatory question. Alternatively, imagine that an individual was entirely convinced that they possessed a complete and worked out set of moral rules. Their application would still be an issue: decisions would still be required. These decisions could not be (by stipulation) about the rules themselves and so must be about how to apply a rule or which rule to apply. This is, I take it, Wittgenstein's point and this is part of why he only considers examples of settled and uncontroversial rules. In both of these imaginary cases we approximate the 'add 2' example.

Even granting all of this — that there is a gap between rule and application —, there is still the question of why we should think that the application is so important? This is important for our discussion because I want to show not only that how a rule is applied needs to be explained but that it is the most important feature. The general point which effects this stance is the non-normative requirement discussed in chapter one. An important, if not primary, concern of those interested in developing a fully specified set of rules has been to lay out what ought to be done or at least provide a schema according to which agents could determine what ought to be done. The rules have thus played the

role of giving content to morality. Now as we saw in chapter one, the interest of this project is sharply focussed on the structure of the judging rather than on the content of the judgement. The non-normative requirement is a consequence of this focus. If we are to get a satisfactory account of the act of judging, and specifically, the conditions of which make it possible, then we cannot privilege any particular account of the content of the judgements either at the beginning or over the course of the investigation. We cannot afford to risk a faulty account by privileging a particular conception of either the limits or the content of morality.⁸⁷ As a result, we are not concerned with which set of rules best captures what one ought to do or best allows one to determine one's obligations.

In addition to these more general points there are two sets of considerations that illustrate this extra claim about the importance of the application step. First, it is imaginable that someone could genuinely not possess (or be committed to) any rules and yet make moral judgements. [What enables us to notice this is the lack of normative content in what has been considered. Again, if we had stipulated what kinds of content were to count as moral then this observation may not have been possible.] That is, we can readily imagine someone who does not refer to any rules or see any rules as governing their moral behaviour or moral judgements — Jonathan Dancy and other particularists, maybe casuists and some projectivists would be prime examples. This means that it is possible to make moral judgements without thinking that rules play a role in those judgements. It is hard, then, to see how possessing rules is necessary for the making of moral judgements, without antecedently stipulating what counts as a moral matter. Notice that even if it turned out that all moral agents did in fact represent morality in terms of

⁸⁷ Recall the individual discussed in chapter one who takes matters of etiquette to be a moral issue.

rules this would not make the rules necessary for the possibility of moral judgement. One might think that anyone for whom rules do not play a role cannot make right judgements but this, of course, is a normative claim — the thinly veiled insistence that one ought to have rules.

A second kind of case which supports the importance of the application step in a similar way is the following. If we assume that two people possess the same rules then any differences in their judgements will only be explainable by referring to the application of those rules. This, it would appear, makes the application the key. Any doubts about the possibility of there being a case such as this should be dispelled by reconsidering the rule-following discussion. Notice here that again the importance of the application step derives from the combination of the possibility of divergent interpretations and our interest in the act of judging rather than the content.

In the face of these worries about the inability of rules to adequately explain our moral judging one might be tempted towards a thorough-going particularism⁸⁸. Thus, with Dancy (among others), we might suggest that there are no moral rules.⁸⁹ However there seems to be some ambiguity involved in this position: Is the claim that people do not actually use any rules or that rules will not feature in the true theory of morality?

⁸⁸ A temptation similar to this one is, it seems, at work in some interpretations of Wittgenstein's discussions of rules, particularly that of Saul Kripke in *Wittgenstein on Rules and Private Language* (Oxford: Blackwell, 1982.)

⁸⁹ Jonathan Dancy, 'Ethical Particularism and Morally Relevant Properties', *op.cit.*, p.530. Others of some note include E. F. Carritt and to a lesser extent, David McNaughton (*Moral Vision*. Oxford: Blackwell, 1988: chapter 13). Carritt is of particular interest here because he seems not to give *no* place to moral rules but consigns them to the role of 'ballast': "The usefulness suggested is, I suppose, that where the right course is obscure or the passions violent, rules may save the ship. So far as this is so I think that their function is that of ballast rather than compass." (*The Theory of Morals*. Oxford: Oxford University Press, 1928: 115.)

There are various ways of stating the latter alternative but all involve reference to the 'true' or 'right' content or limit the scope of what is to count as a moral matter. Clearly this puts the question too simply. The first of the disjuncts is an empirical claim and, moreover, a false one. There are moral rules and people do use them. The genuine claim is, I think, that people ought not to think of morality as rules-based or, better, that they are incorrect or making a mistake about morality in so doing. The more substantive point, then, of the question above is whether the mistake being made when people talk in terms of rules is one about the content of morality or about the structure of moral judgements. That is, when a particularist claims that there are no moral rules do they mean that: (i) when we really understand what morality is about, what things are right and wrong and what acts ought and ought not to be performed, then we will see that general rules are of no use in attempting to describe this content; or (ii), even though it appears that people use moral rules when they decide what they ought to do, they really are not or could not be doing so. Importantly, we would come to this latter conclusion by investigating what it is that people do when they make such judgements.

Noticeably, both the interests of my project and the discussion of rule-following above concern the latter of these alternatives since they too involve considering what it is that we are able to do when we make moral judgements and the role that rules can play in the account of this ability. However, the conclusion we drew from the rule-following discussion was *not* that people are not or could be really using rules even when they think they are. It was that even when the rule is clearly specified and the behaviour dictated by the rule seems blatantly clear, application still needs to be explained. What is involved in following the rule is not *only* explained by the possession of the rule. So, if we interpret

the particularist's claim to be one about the structure of moral judging, our arguments, at least, do not support the position.

In any case, Dancy's arguments seem almost exclusively concerned with the first interpretation of the particularist's thesis — the claim that the rules will not or could not figure in an adequate theory of the content of morality. Since we are not concerned with the shape of such a theory, his arguments do not immediately concern us. Having said this there is a further remark that is warranted regarding these arguments. A large proportion of Dancy's concerns about pluralism are concerned with how any set of rules is to be justified. That is, what kinds of reasons or evidence we could have for thinking a rule true. He discusses both 'generalist' and 'particularist' epistemologies as candidate ways of justifying a set of rules. Essentially, he is troubled by an instance of an application of a rule being also evidence of the rule's truth. Briefly, we might say that the justification of rules comes from our theoretical reflection or deliberation. These reflections and 'theory-building' influence but do not completely determine what we judge. We take a different stance when we reflect on our judgements as compared with our actual judging in the context. How we think about ourselves, our values, the kinds of judgements that we have made and those that we are likely to make could very well influence (rather than, determine) the kinds of general principles to which we are committed and, in turn, the judgements that we make⁹⁰.

In short, we should, in our case, resist the temptation to particularism. The rule-following discussion indicates that simply citing the rules of a normative moral theory

⁹⁰ This kind of suggestion seems similar in form, at least, to the conclusion reached by Blackburn on this issue: "In effect we want our sensibilities to be reliable and projectible (pulling in the direction of general rules) but also sensitive and flexible (pulling in the direction of emphasis on particular contexts)." (*Ruling Passions*, p.308.)

will not be enough to explain the judging. There will be the residual question of how to explain the application. Furthermore, even if it turns out that rules are needed in a normative moral theory, given the goals of this chapter, it is the application of rules in particular cases that matters most. Again, the clearest way to see this is by recognising the possibility that two people may differ about the shape and content of normative morality and yet both be making moral judgements. So, Dancy and, say, a pluralist may both make judgements of obligation whose content is dictated by the ‘shape’ that they give to the true or right moral theory and yet, in so judging, be, as it were, doing the same thing.

Before moving on it is useful to revisit McDowell’s suggestions. Not only is the discussion above relevant to his view but it will provide a nice link into the next chapter. When discussing McDowell’s account of rule-following we stopped short of going through his positive proposals. After setting up the rule-following issues and illustrating that we should not take the skeptical route, he gives an account of what grounds our confidence that our behaviour will not come adrift. He suggests that once we understand that the ground of our confidence is found by immersing ourselves within our own practices — in the way in which we do when we are not reflecting on the relationship between our judgements and the world as it is — we will no longer find the lack of ‘independently fixed rails’ disturbing. Thus, someone who has been taught and has become habituated in the practices of morality develops the ability to see what situations

require. McDowell (in papers other than 'Non-Cognitivism and Rule-Following') compares this ability to the ability to detect secondary properties.⁹¹

Now it should be pointed out that we are only inclined to travel the secondary properties route if we are interested in both what the agent has the ability to do and what is to count as the right version of that ability. The claim that being able to detect moral properties (in our terms, have the right 'take') is just like being able to detect colour properties is only tempting if we are immediately interested in whether some 'takes' are true (or *really* better) and others not and if we think that, actually, some 'takes' really represent the context as it is. That is, the analogy between 'The tomato is red' and 'X is wrong' or 'I ought to X' will be more impressive for those inclined to find that 'X is wrong' is true. If we have antecedently limited our concern to understanding the agent's ability and not whether the ability delivers the 'right' results the secondary properties analogy is not as captivating. There is a good deal more to be said about both McDowell's strategy and the 'secondary properties' route but, as I have already noticed, discussion of the issue of truth in moral judgements must come towards the end of the project.

C. Aristotle's Minor Premise.

It has been my contention that rules cannot provide us with an adequate explanation of what goes on when an agent makes a judgement of obligation. Indeed the explanation of how the agent comes to have the 'take' on the context that he has remains the most important task in the overall explanation of the making of these judgements.

⁹¹ Particularly, 'Values and Secondary Properties' in Honderich, T. (ed.). *Morality and Objectivity*. London: Routledge & Kegan Paul, 1985.

This is because the way that the agent 'attends to' the particularities of the context determines whether or not the pressure is experienced. Alternatively, the agent's judgement of obligation, of what ought to be done, depends on how the agent attends to the particularities of the context. Of course, the attention can still be cashed out as either a primarily cognitive ability or as an attitudinal one.

There is a good reason why this issue may not have been focussed upon in the past. Those concerned to develop a rule-based account have, it seems, been largely interested in getting the one true set of rules and so have looked for correlations between properties and judgement contents. Thus, the act of judging might not have been sufficiently attended to. If one were concerned to arrive at the right rule or set of rules, then how the agent applied a given rule would not necessarily be to the fore. Here we might recall Mandelbaum's discussion of Ross noted at the end of the previous chapter or the so-called problem of act-description which has plagued many deontologist theories, particularly Kant's. In the latter case, if the act is described in one way then the rule or a rule governs it but if it is described in another, perhaps equally plausible way, the rule does not govern it or another one does. Those concerned to provide an account based on general sentiment, not as interested in truth conditions for the content of moral judgements, have tended to concentrate on psychology — particularly, on causal explanations of why we react as we do. As a result, the quest is for those features of human psychology, facts about our psychological natures, which account for the reactions that we have to contexts which confront us. Thus the focus is on general features common to all rather than the seemingly smaller issue of how the agent relates, in a particular context, to that context in such a way that a moral judgement is made.

These remarks are obviously not intended to exhaust moral philosophy. As we shall see there are a number of significant exceptions to what are perhaps only general tendencies.

These claims concerning the importance of explaining of the agent's 'take' do not, in any way, preclude either rules (and their cognitive allies) or sentiments from playing an important role in the final story. The matter of an agent's 'take' as well as the background account of an agent's values can still be given in either of these sets of terms. We are given some inkling of how these details might go for the rule option in Martha Nussbaum's discussion of Aristotle in 'The Discernment of Perception'⁹². She writes:

We must notice first that rules could play an important role in practical reason without being prior to particular perceptions. For they might be used not as normative for perception, the ultimate authorities against which the correctness of particular choices is assessed, but more as summaries or rules of thumb, highly useful for a variety of purposes, but valid only to the extent to which they correctly describe good concrete judgements, and to be assessed, ultimately, against these.⁹³

The parallel kind of story for the role of sentiment, as we saw in the previous section, might allow that our reactions and responses become increasingly complex, manifesting a

⁹² Martha Nussbaum, 'The Discernment of Perception', *Love's Knowledge: Essays on Philosophy and Literature*. New York: Oxford University Press, 1990. Similar points are made in Chapter 10 ('Non-Scientific Deliberation') of *The Fragility of Goodness*. Cambridge: Cambridge University Press, 1986.

⁹³ Nussbaum, *op.cit.*, p.68. It is interesting to distinguish this view of the role of rules from E. F. Carritt's (mentioned in note #17). Carritt, with his ballast and compass analogy, suggests that rules steady the ship rather than guide it. On Nussbaum's view, rules can guide decisions but only when a further decision is made allowing them to do so. On my view both are possible and so neither is necessary for being able to make moral judgements.

versatility and subtlety that cannot be adequately captured by general sentiment types. The origin and basis of these more complex reactions, however, could still be given in terms of an over-arching feature of human nature but the importance of such a feature is diminished. Instead, the focus might be on those desires and attitudes which influence how the world appears to us in addition to those which determine the variety of our goals and aims.

In the paper mentioned above, Martha Nussbaum provides not only an account of the way in which rules might fit into this kind of picture but an account of Aristotle's views about the 'priority of the particular' which gels with many of the main points of this chapter. As such it provides the beginnings of a link between some of Aristotle's concerns in the *Nicomachean Ethics* and the discussions of this and the next few chapters. Although her focus lies more with the rule option than with the sentiment route, room, as we shall see, can be made for both.

Nussbaum's consideration of Aristotle's views on the 'particular' divides neatly into two parts. In the first part she shows that Aristotle does indeed suggest that the particular is prior to the general or universal. In the second, she answers the objection that Aristotle's remarks only apply to actual systems of rules and do not imply that a "science of ethics" would fail to arrive at a complete and correct system. In the first of these parts Nussbaum concentrates on three passages in particular:

- (1) But up to what point and to what extent a man must deviate before he becomes blameworthy it is not easy to determine by reasoning, any more

than anything else that is perceived by the senses; such things depend on particular facts, and the decision rests with perception. (1109b19-24)⁹⁴

- (2) When the Law speaks universally, then, and a case arises on it which is not covered by the universal statement, then it is right, where the legislator fails us and has erred by over-simplicity, to correct the omission — to say what the legislator himself would have said had he been present, and would have put into his law if he had known. (1137b19-24)
- (3) For when the thing is indefinite the rule also is indefinite, like the leaden rule used in making the Lesbian moulding; the rule adapts itself to the shape of the stone and is not rigid. (1137b28-31)

The significance of these three is quite clear. The first of these deals with the importance of deciding *in context*. Reasoning (Nussbaum translates this as ‘a principle’) is too abstract to be able to capture the concrete particularities involved in good practical decision-making. The second deals with the need for rules (in this case the Law) to be supplemented and corrected in cases where the rule is too simple. The third points to the importance of adapting to the context if we are determine what is right. Thus, “good deliberation, like the Lesbian Rule, accommodates itself to the shape that it finds, responsively and with respect for complexity.” (p.70)

It should be pointed out that, in the third passage, ‘rule’ is not to be understood in the way in which we have been using it. The actual rules that we have, like the Law, are

⁹⁴ For the sake of consistency I will stay with the Ross translation. The differences between this and Nussbaum’s should not effect the points that I wish to make. In *The Fragility of Goodness*, p.492, n29, Nussbaum gives the following list of additional places where Aristotle refers to practical wisdom as being concerned with ultimate particulars: 1110b6, 1126b4, 1142a25, 1143a29, 32, b4, 1147a3, 26, b5.

not being spoken of as indefinite. Instead, it is the rule which says what is right. Thus, the person of practical wisdom acts “in accordance with the right rule” (1138b24)⁹⁵. Just as the Lesbian ruler measures shapes which cannot be measured by a straight ruler, so the rule possessed by the practically wise adapts itself to what is right and this is indefinite.

Now the question arises as to whether these remarks imply only that our systems of rules, so far, fail to do justice to the complexity of moral decision making and not that the particular is, in principle, prior to the general. In response to this, Nussbaum cites Aristotle at 1137b17-19: “The error is not in the law or in the legislator, but in the nature of the thing, since the matter of practical affairs is of this kind from the start.” She then continues, in support of this by raising a further three sets of considerations concerning the “matter of the practical”: first, “practical matters are mutable, or lacking in fixity”; second, they have an “indeterminate or indefinable character”; and third, concrete cases may “contain some ultimately particular and non-repeatable elements” (pp.71-72). The first two of these ideas are fairly straightforward. Aristotle’s discussion of specific virtues reveals the variability which he associates not only with their application but with the business of doing the right thing. The indeterminate character of the practical matters in Aristotle is clearly illustrated in the example above of the Lesbian ruler.

Nussbaum expands the third of these ideas by referring to the ‘doctrine of the mean’. She alludes to the passage in book II (1106a35-b8) where Aristotle discusses the moderate diet for Milo the wrestler. The point here is that what counts as moderate for

⁹⁵ At this point in the text, Ross refers the reader to 1103b31, 1107a1 and 1114b29. The passage at 1107a1 is Aristotle’s definition of virtue. Alasdair MacIntyre, in *After Virtue* (2nd Ed., London: Duckworth, 1985: 152), translates the phrase ‘*kata ton orthon logon*’ as ‘according to right reason’ which I think does assist with the current point but does not make too much difference overall.

one individual will not be for another. we may add to this that the mean will also vary with contexts for the same individual, so that the moderate or 'right' amount of anger in a given situation, say, in the face of gratuitous bigotry is different to that appropriate when a child misbehaves. Additionally, our attention is drawn to Aristotle's remarks on friendship. Here, "features of shared history and of family relationship that are not even in principle repeatable are allowed to bear serious ethical weight" (p.72).

Little has been said in this discussion about the role of the sentiments. Nussbaum goes on in her paper to consider the role of the imagination and the emotions but her discussion is not quite so neatly tailored to explore the sentiment option⁹⁶. Aristotle, however, provides us with plenty of material to use in order to see the relationship between his claims about the 'priority of the particular' and the sentiments. This is most striking in his discussions of specific virtues in the latter half of book III and book IV. We might say that the point of this part of the text is to give an account of specific virtues in such a way as to illuminate, not only their objects but the accompanying kinds of attitudes, motives and emotions. The general form of Aristotle's account is reflected in the following remark with respect to courage:

The man, then, who faces and who fears the right things and from the right motive, in the right way and the right time and who feels confidence under the corresponding conditions, is brave. (1115b18)⁹⁷

⁹⁶ There is a debate about how best to understand the emotions which reflects, at least in part, the distinction between the cognitive and the conative. This debate should not effect our discussion largely because the split between these two is repeated in this debate. The argument is not about whether the emotions have a role to play but whether the role that they play adds credence to the cognitivist or the non-cognitivist story.

⁹⁷ It is interesting to notice that the passage from which this is taken is referred back to by Aristotle (and according to Ross) on four occasions over the course of the discussion of

This kind of remark can be found a total of seven times over the course of these books⁹⁸. The importance of this is to illustrate that the importance of the particular, with which Aristotle is concerned and which Nussbaum explores, applies as much to desires and attitudes. Thus, the man who judges rightly not only possesses the right view of the circumstances but desires to the right extent and in the right way “and so on”⁹⁹. The priority of the particular is reflected as much in desires and attitudes as it is in applying rules.

We concluded in the previous sections that how the agent ‘attends to’ a context, how he formulates, comes to develop, a ‘take’ is dire need of explanation, since without this we have not explained how the pressure is felt and hence, how the judgement comes to be made. The discussion of this section is intended to show the parallels between these claims and Aristotle’s view of the importance of the particular. In the practical syllogism the place of the particular is represented by the minor premise. The importance of the particular is thus translated as the importance, for practical wisdom, of being able to ‘get’ the minor premise. At 1141b14-23 Aristotle writes:

Nor is practical wisdom concerned with universals only — it must also recognize the particulars; for it is practical, and practice is concerned with

courage. Thus speaking of those who are sanguine, Aristotle says that “when their adventures do not succeed ... they run away; but it was the mark of the brave man to face things that are, and seem, terrible for a man, because it is noble to do so and disgraceful not to do so.” (1117a15-18) At the ‘mark of the brave man’ we are referred back to 1115b11-24. This adds to the importance of the initial remark.

⁹⁸ Specifically, at 1119a11-18, 1119b15-19 (both with respect to temperance), 1120a24-29, 1120b27-33 (both with respect to liberality), 1125b32-1126a1, 1126a4-8, 1126a32-1126b8 (all with respect to anger).

⁹⁹ The further we progress through books three and four the more inclined Aristotle is to abbreviate his list with something like ‘and so on’.

particulars. This is why some who do not know, and especially those who have experience, are more practical than others who know ... Now practical wisdom is concerned with action; therefore one should have both forms of it, or the latter in preference to the former.

Thus, the importance of explaining how the agent constructs a 'take' is equivalent to the importance of explaining how an agent gets the minor premise. On the face of it, we do not have as rich a notion of the *right* minor premise or of the *right* 'take', as Aristotle appears to have, however we do have the idea of a 'take' which leads to the experience of the pressure and one which does not.

It is, however, arguable that Aristotle does not have a rich notion of the *right* minor premise at all. It might be suggested that what is right, in the end, turns out, on Aristotle's view, to be whatever the man of practical wisdom would do. Such an interpretation, it seems, is consistent with that of both Nussbaum and McDowell¹⁰⁰. As we have mentioned, at a number of key points Aristotle places a significant amount of weight on the notion of the right rule. Early in book VI he acknowledges this significance and sets out to determine "what is the right rule and what is the standard that fixes it" (1138b34)¹⁰¹. The remainder of book VI is then spent distinguishing practical wisdom from other forms of knowledge and explaining how those who are practically wise both possess and set the standard. In effect, no definitive content is given to the notion of what is right (and this is not surprising given the discussion above of the priority of the

¹⁰⁰ In various papers including 'Virtue and Reason' and 'Are Moral Requirements Hypothetical Imperatives?'.

¹⁰¹ Notice that there is very little difference between this translation and MacIntyre's. On Ross' version however, the fact that Aristotle does not answer the first of these questions is more obvious.

particular). If we are inclined to accept this kind of interpretation of Aristotle's account then the main difference between Aristotle's way of talking about the right rule or the right action and ours lies in Aristotle's confidence that there is just one right rule or right action. Part of what we have set out to consider is whether and in what sense there is a right action — this was the issue of the truth of moral judgements. Thus we are entitled only to talk of right action in terms internal to an agent's view of morality.

CHAPTER 4: EXPLAINING THE EXPERIENCE (PART I).

“... the thing we should look for may be precisely what Aristotle provides — namely, a conceptual framework which we can apply to particular cases, which articulates the reciprocal relations of an agent’s concerns and his perception of how things objectively are in the world ...”

(David Wiggins, *Needs, Value, Truth*, p.237)

In this chapter we consider the first of two kinds of accounts (broadly, cognitivism and non-cognitivism) that can be used to explain the two features of the experienced pressure accompanying judgements of moral obligation: (i) that it is experienced as objective and (ii) that since it is a pressure *to act*, it is essentially related to motivation. Sections A and B prepare the ground for this discussion and that of the next chapter by providing, in the former case, a framework for the kind of explanation that we seek, whilst in the latter, a closer look at the kind of attention that we will attempt to explain. In section C I will present an account of the cognitivist position which seems to best satisfy the specifications of the project, followed, in section D, by a look at two major points of contention with this position.

A. Aristotle’s Practical Syllogism.

We begin with some considerations of Aristotle’s practical syllogism. The syllogism is, after all, an account of practical judgement of some sort (and the ‘components’ thereof) and these are our focus also. Moreover the syllogism looks to involve particularity and perception in a way that is consonant with what has been said to

this point. But besides continuity, there are two important reasons for discussing the practical syllogism here — both of these tie closely to the tasks of this chapter. First, and once we have overcome some interpretive issues, the practical syllogism will provide an heuristic with which we can draw out the kinds of ingredients that are available to explain the judgement and the experienced pressure. Here we encounter an important assumption that will inform the discussion of the following chapters: the tools that are available to explain both the judgement and the two characteristics of the pressure are states that are broadly characterized as belief-like and desire-like.

Second, in spite of the fact that there are other accounts of practical reasoning with which we could begin, the kind of ‘cognitivist’ accounts which are most relevant to the tasks before us, those of McDowell and Wiggins, are developed in close proximity to a reading of Aristotle’s practical syllogism. Much of the overtly cognitivist discussion takes place, for McDowell and Wiggins, in the context of judgements of value (and the analogy with secondary properties). It will be recalled from chapter one that Wiggins, in particular, makes a good deal of the distinction between judgements of value and deliberative or practical judgements. In their work on the issues most immediately related to their form of cognitivism both he and McDowell invariably concentrate on the former rather than the latter¹⁰². In order to understand the ways in which a cognitivist position can be developed along these lines we must utilise McDowell and Wiggins’ reflections on practical (deliberative) judgements. And in this context, the focus of discussion for both is Aristotle’s practical syllogism.

¹⁰² One exception to this is Wiggins’ paper, ‘Truth, and Truth as Predicated of Moral Judgements’ (*Needs, Value, Truth. op.cit.*, pp.139-84.). However even here Wiggins’ concern is with the ways in it might make sense to take practical moral judgements to be true rather than on the nature of the act of such ‘judging’ in our sense.

The discussion in the previous chapter almost exclusively involved the relationship between an agent's 'thoughts' about features particular to a context and those ('thoughts') that are general and are applied in contexts. We have seen the importance of the particularities of contexts and of explaining the agent's 'take' on particular contexts. In the context of Martha Nussbaum's discussion of 'the priority of the particular', I highlighted the close parallels between our conclusions about the importance of understanding the agent's 'take' and the importance, for Aristotle, of an agent's ability to 'get' the minor premise. Also, the concentration on judgements of obligation and the focus on the formation of a particular 'take' have played down the significance of the value 'component' of the agent's judgements. Some role must be given to these components if the account is to be satisfactory. The values that the agent has will give content to the judgements that the agent makes — that is, they will assist in determining what actually ought to be done. Having said this, we will not need to decide the issue of the nature of value or the nature of judgements of value. Instead, we are required to give such judgements, however they are to be accounted for, a role in the making of judgements of obligation.

A case can be made for thinking that one of Aristotle's major concerns in his discussion of the practical syllogism is to establish a parallel between theoretical and practical reasoning and so, between the theoretical and practical syllogism. Whether or not this is anything more than a parallel is the subject of dispute. In any case, this parallel or (perhaps better) ratio is given by Aristotle in the following passage from Book VI:

What affirmation and negation are in thinking, pursuit and avoidance are in desire ... of the intellect which is contemplative, not practical nor

productive, the good and the bad state are truth and falsity respectively...;
 while of the part which is practical and intellectual the good state is truth
 in agreement with right desire. (1139a21, 26-31)

Given this and the language that Aristotle uses in discussing the practical syllogism it is perfectly natural to take the practical syllogism as representing an argument form: major premise, minor premise and conclusion. Thus the practical syllogism might be intended somehow to ‘model’ our practical reasoning in much the same way as the theoretical syllogism is intended to model our theoretical reasoning. The model, we might think, gives us the structure of the process that is being modelled — what the parts involved in the process are and how they are related. The practical syllogism, then, since it is a syllogism, gives us some kind of conclusion and since it is practical some kind of conclusion about or related to action. For our purposes here we need not get into the specific kind of conclusion involved but *at the very least*, it is a practical judgement, that is a judgment about what ought to be done¹⁰³.

John M. Cooper, in *Reason and Human Good In Aristotle*¹⁰⁴ has noticed an ambiguity in how we might view the relation between practical reasoning and the practical syllogism. It is important to consider this ambiguity so as to clarify the role that the practical syllogism plays with respect to both practical reasoning and to deliberation. The practical syllogism could represent an account of how it is that practical judgements, practical conclusions or actions are actually arrived at or it might represent an account of

¹⁰³ The debate about whether or not the conclusion is an action or merely a judgement does not affect the discussion. The way in which I discussed ‘judgement’ in chapter one was intended to allow both sides of this debate.

¹⁰⁴ Cooper, John M. *Reason and Human Good in Aristotle*. Cambridge, MA: Harvard University Press, 1975.

how such “conclusions and actions (however arrived at) can be justified, or at any rate explained” (p.5). Interestingly Cooper also allows that, in some way, it might represent both. Let us put this alternative aside for the moment. Now if the practical syllogism were to be understood in either of the first two ways then we would be provided with two differing kinds of explanations of practical thought. In the first case, the explanation would be a causal/empirical account of our psychological processes. In the second, as Cooper notes, the explanation would be a justificatory or reconstructive account of how our practical thought might have gone. This would amount to an account of the structure of the kinds of reason-explanations that are available to the agent who seeks to understand or is called to justify his action. Thus Cooper writes

If a person does have reasons for acting, they will, when produced, constitute a deliberative argument in favour of the decision actually made; hence one can regard that argument as lying behind and supporting the decision even though it was not actually gone through in advance. (p.9)

Notably, it is only on the latter view that the practical syllogism represents anything akin to an argument. Just as the premises of a theoretical syllogism are taken to justify the conclusion, so, on the second, reconstructive reading the major and minor premises are taken to justify the practical ‘conclusion’ — by providing a reason-explanation. On the former construal, the major and minor ‘premises’ might be taken to be psychological states of the agent which combine to yield, as it were, causally, the practical ‘conclusion’.

It is the latter, reconstructive interpretation that is favoured by Cooper as well as, to a certain extent, both McDowell and Wiggins¹⁰⁵. Cooper essentially provides two reasons for favouring the reconstructive account. The first of these involves the relation of deliberation and the practical syllogism to practical reasoning whilst the second appeals to Aristotle's account of moral development and education. First, for Aristotle it looks as though deliberation and the practical syllogism play a central role in his account of practical reasoning. After all, choice is closely connected to what is or has been determined by deliberation (*NE* 1112a15). However as Cooper notes, deliberation is usually taken to be a 'figuring out' of what to do. But there are plenty of practical/moral decisions that we make without any such 'figuring out'. The decision, as it were, comes to us straight-away. Thus to take Aristotle's account of practical reasoning to be an account of how we actually arrive at practical decisions would be either to severely limit what is to count as a practical/moral choice or to suppose that all practical/moral choices are calculated or 'figured out'. The first of these disjuncts would make Aristotle's theory very weak. We are (presumably) interested in all of what we would ordinarily call

¹⁰⁵ Cf., McDowell, 'Virtue and Reason' (*op.cit.*) and 'Some Issues in Aristotle's Moral Psychology' (in *Mind, Value and Reality*. Cambridge, MA: Harvard University Press, 1998: 23-49) and Wiggins, 'Deliberation and Practical Reason' (in *Needs, Values, Truth*. 3rd ed. Oxford: Clarendon Press, 1998: 215-38.). McDowell in 'Some Issues in Aristotle's Moral Psychology' writes:

It seems best to give Aristotle a pinch of salt of this: to take it, for instance, that when he suggests choice (*prohairesis*) is the upshot of prior deliberation, his point is really that the conceptual structure that is characteristic of deliberation figures in the proper explanation of the relevant actions, whether or not prior deliberation takes place. (p.24)

This, as we shall see, is a slightly different story to Cooper's reconstructive alternative.

practical choices and not in some particularly narrow subset masquerading as practical choice. The second disjunct looks to be simply false.¹⁰⁶

Second, if Aristotle was suggesting that practical thinking/deliberation did *actually* take the form laid out in the practical syllogism it is unclear how this would sit with his account of moral development and education. As Aristotle presents it, this process is a gradual one where the student, through training, habituation and experience, learns not only what he is trained to do but to appreciate and understand why it is right. Thus, as Cooper puts it, “because this sort of understanding is arrived at gradually, it seems wrong to say that the virtuous man’s adoption of his virtuous policies is founded on an explicitly deliberated decision” (p.8).

However, as it stands such an interpretation of the practical syllogism and hence, such an account of practical reasoning is not of much help with our task. It will be recalled from chapter one that the task laid down was to explore the nature of the moral judging — where this was understood to be the necessary pre-conditions for the possibility of making judgements of obligation. Because this interpretation takes the syllogism to represent a reconstruction of the reasoning that *might* have taken place, there is no guarantee that it or anything like it actually did take place, let alone an indication that it provides the conditions which make such judgements possible.¹⁰⁷

More specifically, it is not at all clear that this interpretation can help with our current problem. We need to explain how (‘what goes on when’) an agent actually gets the so-called minor premise; that is, how the agent forms particular ‘takes’. One kind of

¹⁰⁶ Cooper writes that on this view “Aristotle would seem to hold that moral decisions are much more excogitated than in general they are.” (p.7)

¹⁰⁷ There is much more to be said on these issues. See below.

explanation that might be offered for this, consistent with the reconstructive interpretation, returns us to the considerations of moral development and education — upbringing. The key ingredients of ensuring that someone is brought up with this practical ability are habituation and experience. So, it is by being exposed to the right kinds of things, being taught to value them and, finally, by internalising them that one becomes able to ‘see’ situations in the right way. However, this will not do. Neither the appeal to habituation or to experience will satisfy our project since the key question remains, namely, ‘What is the ability that is thereby cultivated or developed?’ This ‘moral development’ response misses the target because it tells us what is involved in an agent’s coming to have the ability and not in what that ability consists. As might be inferred from parts of chapter three, I think that more can be said. I do not, however, intend to contravene Aristotle’s suggestions regarding upbringing, habituation and experience — indeed they will become very important in chapters six and seven. These matters do not constitute an answer to the immediate question.

The search for an adequate conception of the practical syllogism takes us back to the third of Cooper’s alternative interpretations that in some way, combines both the ‘causal/empirical’ and the ‘justificatory/reconstructive’ alternatives. This alternative, I will suggest, can manage to have a foot in both of these camps, as it were, by suggesting that the syllogism gives us the necessary structure of the kind of psychological explanation with which both of the other alternatives are concerned. What distinguishes this construal from both of the others is that it only attempts to give the structure of

psychological explanations rather than to be too specific about their content¹⁰⁸. What distinguishes it from the ‘causal/empirical’ model is that it provides the *necessary* structure of these explanations. What distinguishes it from the initial ‘justificatory/reconstructive’ model is that the structure it provides applies both to what was actually thought as well as to reconstructions and it avoids, to a certain extent, the question of justification. There is an argument implicit in what follows that this third alternative conception of the practical syllogism is the best interpretation of Aristotle. However, I will not push it here. I am primarily concerned to develop a schema in which we can tackle our question that takes its leave (as McDowell and Wiggins do) from these Aristotelian considerations.

An example will help prepare the ground for this third construal of the practical syllogism. Suppose that Barney suggests that he ate the last helping of ice-cream because he did not want it to melt and he could see that it was beginning to do so. If this is taken as (or meant as) an explanation of the first sort then it is a claim about the actual mental state that Barney was in just prior to his eating the ice-cream — let us say, the desire that the ice-cream not melt and the belief (resulting from the perception) that it was melting. Of course if Barney was not actually in this mental state then this explanation is no good. On the other hand, if this example is taken to be of the second sort, it does not require that Barney was actually in the mental state posited. Instead, the reasons are offered partly as what he might have thought and partly as a justification/defence of what he did. What is important to notice here is that whether or not this reconstructive/justificatory

¹⁰⁸ It is useful to observe here that Aristotle’s examples of the practical syllogism (see *NE*, books VI and VII and *De Motu*) are noticeably poor. One might suspect that if the practical syllogism were anything more than a structural account it would have been easier to give better and more substantial examples.

explanation is plausible as such depends upon how likely/plausible it is that this (the explanation) is what Barney actually thought. On the assumption that the kind of evidence that is available for a particular causal/empirical explanation is all and only Barney's behaviour, then this evidence is the same as that which will determine the adequacy of the reconstructive/justificatory explanation.¹⁰⁹

That there is a third alternative becomes clearer once we see this connection between the first two alternatives. If we think that the two construals are related in these ways then the third alternative — that which gives the necessary structure of psychological explanations — will be one which combines both by placing limits on how we explain what actually took place as well as how we might reconstruct particular instances of practical reason. What this alternative cannot do is enable us to fill in the content of the states mentioned in the explanation and because of this, it cannot provide a thorough-going defence or justification of the act it is intended to explain. Indeed because

¹⁰⁹ There is a good deal more that could be said here. These ideas tie quite closely to Hobbes' remark that "thought is quick" (*Leviathan*, III, paragraph 3 (Hackett, p.12)) which, I take it, can be read as blurring the distinction between the first two alternatives. If thought is quick then we presumably go through more of the reasoning that is supposed only to be reconstructed. It might also be noted that one may act deliberately without deliberating — this highlights the distinction between cases of wrestling, wondering and considering prior to acting (deliberating) and cases without these but where an agent acted intentionally, i.e. deliberately, although the same word is used for both. Recall, from the discussion in chapter two, that there can be a high degree of complexity in the 'considering' kind of case. One might make many judgements, both moral and otherwise, in the course of deliberating. Whilst the overall account must be able to deal with these sub-judgements, the central focus here is, let us say, the 'final' judgement of obligation. Thus although there may be some residual issues about the how much a proposed explanation reconstructs reasoning that might have taken place, the majority of the considering-type cases can be understood as a series of judgements (or, perhaps, a series of syllogisms). Importantly, these considerations essentially concern the nature of psychological explanations (cf. Blackburn, *op.cit.*, pp.51-9). Some of these issues are addressed below.

this alternative only provides the necessary structure of practical reasoning it cannot, by itself, explain any particular action whatsoever.

The interpretation of the practical syllogism that comes closest to this idea is that given by David Wiggins in 'Deliberation and Practical Reason'.¹¹⁰ This account is both widely known and respected. It is also the account that Wiggins and, to a large extent, John McDowell use as integral parts of their own positions. I will sketch the key features of the account that Wiggins offers. He writes that

The first or major premise mentions something that can be the subject of desire, *orexis*, transmissible to some practical conclusion (i.e. a desire convertible via some available minor premise into action). The second premise details a circumstance pertaining to the feasibility, in the particular situation to which the syllogism is applied, of what must be done if the claim of the major premise is to be heeded. (p.227)¹¹¹

This provides only the bare bones of the syllogism. The key question is how it works. Importantly, the problem of what to do is usually a problem that arises in response to a particular context rather than because there is something that one is antecedently required

¹¹⁰ Wiggins, David. 'Deliberation and Practical Reason' *Proceedings of the Aristotelian Society* vol.76 (1975-6). Reprinted, with revisions, in *Needs, Values, Truth* (1987) 3rd ed. Oxford: Clarendon Press, 1998: 215-37. References will be to this edition.

¹¹¹ He continues: "[It is natural] to describe the first premise of a practical syllogism as *pertaining to the good* (the fact that it pertains to some — either a general good or something which the agent has just resolved is good in this situation — is what beckons to desire) and to describe the second or minor premise as *pertaining to the possible* (where 'possible' connotes the feasibility, *given* the circumstances registered by the minor premise, of the object of concern mentioned in the first premise)." (p.227)

to maximise.¹¹² This requires that the agent be able to isolate the relevant features of the situation, not all of which will be obvious. “To see what [these features] are, to prompt the imagination to play upon the question and let it activate in reflection and thought experiment what ever concerns and passions it should activate, may require a high order of situational appreciation or, as Aristotle would say, perception (aisthesis).” (p.231) Of course in cases where there is no time to reflect the virtuous agent will still be good at isolating the relevant features and will be poised to notice those that are not obvious.¹¹³ Thus the ‘best’ practical syllogism has a minor premise that “records what strikes the person as the *in the situation most salient feature of the context in which he has to act*. This activates a corresponding major premise that spells out the general import of the concern that makes this feature the salient feature in the situation.” (p.233) The evaluation of a particular practical syllogism “is of its essence dialectical, and all of a piece with the perceptions and reasonings that give rise to the syllogism in the first place.” (p.234)

John McDowell’s version nicely complements this picture given by Wiggins. He suggests in ‘Some Issues in Aristotle’s Moral Psychology’ that

We can see the orectic state and the doxastic state as interlocking elements in a mechanism, like a ball and socket of a joint. But the capacity that determines which of the potentially motivating features of the situation is to serve as premise of the possible cannot be separated from the orectic

¹¹² This is in contrast with the situation in snooker (discussed by Wiggins (p.230-1)) where the end is to maximise points by choosing the best shot.

¹¹³ Though not mentioned by Wiggins, this is an important point for the interpretive issues discussed above. It tells, I take it, in favour of the third alternative.

state whose content constitutes the premise of the good. Having the right end is not a mere aggregate of concerns; it requires the capacity to know which should be acted on when. If that capacity cannot be identified with acceptance of a set of rules, there is nothing for it to be except the capacity to get things right occasion by occasion: that is, the perceptual capacity that determines which feature of the situation should engage a standing concern. So the premise of the good, and the selection of the right feature of the situation to serve as premise of the possible, correspond to a single fact about the agent, which we can view indifferently as an orectic state or as a cognitive capacity.¹¹⁴

There are three important aspects of this passage that warrant some attention. First, and as we shall see below, McDowell emphasizes the Aristotelian focus on the particular that was observed in the last chapter. The uncodifiability of what constitutes the right end¹¹⁵ means that the capacity possessed by the virtuous person is a case by case matter, keyed to the features of the contexts in which the agent finds himself. That the perceptual capacity possessed by the virtuous person is a perceptual capacity which interlocks with having the right concerns is, in part, derived from this uncodifiability.

Second and following from this, the importance of our being able to view 'the single fact about the agent' indifferently in the two suggested ways yields a kind of cognitivism with the potential to explain both the motivational character of moral

¹¹⁴ McDowell, John. 'Some Issues in Aristotle's Moral Psychology', *op.cit.*, p.30.

¹¹⁵ McDowell calls this the 'thesis of uncodifiability' in 'Virtue and Reason' (*op.cit.*). He goes into more detail there and also in 'Deliberation and Moral Development in Aristotle's Ethics' (*op.cit.*).

judgements and the objective character of the experienced pressure. The kind of cognitivism that we get turns on the perceptual capacity (possessed by the person of practical wisdom) in such a way as to explain Aristotle's suggestion that for that part of the soul that is practical and intellectual "the good state is truth in agreement with right desire" (1139a31). That we are well placed to explain the nature of moral judgements of obligation and the pressure on this account should be clear. On this view the motivation is provided by the 'major premise' component, having/engaging the right concerns, whereas 'viewing it the other way', namely in terms of the agent's perceptual capacity to get the minor premise, we have a cognitive capacity. The pressure is experienced as 'objective' because the situation is *seen* in a certain way (which is, of course, the result of this perceptual capacity).

The third point to note from the above passage is that the idea of the intertwined orectic and doxastic states, captured in the ball and socket metaphor, leaves room for an Aristotelian-style story to be told about moral education. If these two states function along the lines of the ball and socket metaphor, the process involved in a good moral education can be readily understood as one of coordinating and fostering the right concern relative to particular contexts with the ability to recognize the relevant 'engaging' features of those contexts. Thus, exemplification, practice and habituation, especially in the absence of a fixed, definitive set of rules, are of paramount importance.

The quotation at the beginning of this chapter is taken from a passage in which Wiggins is clear about the kind of account that he thinks can and should be given of practical reason and its connection with his account of the practical syllogism. Examining this passage will provide us with a more definitive sense of how the 'third alternative'

works both in the light of the Wiggins' reading of the syllogism and of how it ties to the concerns of this project. The passage I have in mind is:

Against this account, as I have explained it, it may be complained that in the end very little is said, because everything that is hard has been permitted to take refuge in the notion of aisthesis, or *situational appreciation* as I have paraphrased this. And in aisthesis, as Aristotle says, explanations give out. I reply that, if there is no real prospect of an ordinary scientific or simply empirical theory of all of action and deliberation as such, then the thing we should look for may be precisely what Aristotle provides — namely, a conceptual framework which we can apply to particular cases, which articulates the reciprocal relations of an agent's concerns and his perception of how things objectively are in the world; and a schema of description *which relates the complex ideal the agent tries in the process of living his life to make real to the form that the world impresses, both by way of opportunity and by way of limitation upon that ideal.* (p.237, original italics)

There are two key points to draw from this passage. The first involves the connection between the kind of account that can be given of practical reason and the possibility of giving an *exhaustive* specification of the ability possessed by those who are practically wise. I take it that for Wiggins the impossibility of such an exhaustive account is tied to there being no prospect of an empirical theory of all action and deliberation. This in turn is a consequence of the impossibility of giving a complete account of the agent's concerns — the thesis of uncodifiability mentioned above. This way in which our

explanations may give out is also related to the distinction between the content of the judgement and the act of judging. What is being limited by the thesis of uncodifiability is our ability to determine, at least prior to the agent's acting,¹¹⁶ the precise make-up of the agent's concerns — we cannot tell, once and for all, what those concerns exactly are and how they are ranked. This is connected to the idea that what we can say about an agent's concerns (and beliefs) is largely dependent on what they do (and then it is retroactive). The uncodifiability thesis therefore limits what can be said about the content of the agent's mental state in a way that restricts/prevents the ability to predict how the agent will act.

As should be clear, that explanations give out in these 'content'-related ways is no trouble for the project here. I have been concerned to avoid discussion of the content of moral judgement for the reasons outlined in chapter one. Also, the third alternative suggests that we understand the practical syllogism as giving the necessary structure of practical reasoning and thus, as purged of content. Finally the discussions concerning particularity in the last chapter are directly linked to the issue of uncodifiability and thus to 'content'-related explanations giving out. It is important to point out the parallels between the kind of conceptual framework that is suggested by Wiggins on Aristotle's behalf and the kind of quasi-transcendental argument outlined in chapter one. What I take to be the main, if not the only, difference between Wiggins' suggestion here and mine in chapter one is the involvement of the phenomenological features — the 'pressure'.

Second, it should be noticed that this first set of remarks (the third sentence of the passage) is separable from that to which it is a reply. The claims about the uncodifiability

¹¹⁶ There may also be some problems or limits to the precision/definitiveness of explanations of an agent's action after the fact.

of an agent's concerns rule out one kind of account of action and deliberation and (perhaps) make another feasible, namely the conceptual framework account. But, while this places some 'content'-related limits on what can be said about aisthesis, it does not mean that nothing can be said. Uncodifiability forces a certain kind of explanation of aisthesis to give out. This is illustrated by the fact that there remains, even within a conceptual framework, a distinction between aisthesis, understood as situational appreciation, and perception. In the picture that Wiggins gives of practical reasoning aisthesis is the key feature and, on the face of it, there is nothing yet that prohibits explaining aisthesis, by means of a conceptual framework, in terms of a perceptual capacity and the agent's concerns. Indeed Wiggins' further story (section C below) can help in saying more about the way in which a perceptual capacity and a set of concerns can combine to explain aisthesis in just this way. On my account here, the impetus to say more about aisthesis stems from the attempt to explain the experienced pressure. So, just as the uncodifiability thesis limits the kind of account that can be given of practical reasoning generally, the particular characteristics associated with the pressure limit and motivate the way in which we understand the "the reciprocal relations of an agent's concerns and his perception of how things objectively are in the world".

This discussion provides us with an entrance into the cognitivist positions of McDowell and Wiggins as well as a glimpse of how their Aristotelian approach gels with the current project. Aristotle's position is usually taken to be a cognitivist one even though one might doubt that the debate itself was one with which he was concerned. McDowell and Wiggins can be taken to start with this kind of Aristotelian picture and to proceed by developing the position more fully. We will examine their developments and

some further possibilities in chapter seven. Clearly though, this way of putting the task also highlights the key explanatory tools available for the current task: ‘the agent’s concerns’ and ‘the agent’s perception of how things objectively are in the world’. One of the novel aspects of the McDowell/Wiggins kind of account is the way in which these two ingredients are combined. The projectivist (or more generally the non-cognitivist) can be understood as organizing the same set of tools in a different way. We will turn, in chapter five, to the way in which Blackburn, in particular, arranges these explanatory tools.

As we have seen, the debate between cognitivists and non-cognitivists can be taken to be primarily one about whether or not the contents of moral judgements are capable of having truth-values.¹¹⁷ One way of attempting to decide this question is to consider the kind of mental state involved in or responsible for the moral judgement. The idea is that if the moral judgement is explained by reference primarily to cognitive or representational states then it is capable of having a truth-value whereas if the judgement

¹¹⁷ Some may object to this way of characterising the debate, perhaps justifiably. The debate about these issues in Britain in the seventeenth and eighteenth centuries, for instance, can reasonably be taken to concern the nature of our moral capacities rather than the possibility of there being no moral truth. This of course is a reaction to Hobbes’ perceived egoism. The ensuing debate was not about whether Hobbes was right [say, about: “Whatsoever is the object of any man’s appetite or desire, that is it which he for his part calleth *good* ... For these words of good, evil and contemptible, are ever used with relation to the person that useth them: there being nothing simply and absolutely so” (*Leviathan*, ch.13, cf. also ch.15)] but how best to show that he was wrong. In this debate, that moral judgements were either true or false was largely taken to be a given (Hume might be taken to be an exception to this). This part of history, at least, would be on the side of those who object to my characterisation. We might add that the question concerning the truth-aptness of moral judgements only became an issue with the advent of emotivism in the first part of the twentieth century (cf. A.N. Prior, *Logic and the Basis of Ethics*). In any case, whether or not we take the cognitivist/non-cognitivist dispute to be primarily centred around the possibility of moral judgements having truth-values does not matter here since it is only with the latter ‘mental state’ version of the issue that we are concerned.

is explained by reference primarily to non-cognitive states — that is, states understood as responses to the world — then it is incapable of having a truth value. These issues frame the debate (and have done so roughly since Hobbes) and give us the tools that we can use in our account. This is not to say that these distinctions are final. A failure to find an adequate account of the kind we require may well lead to a rejection of this way of thinking about the problem.¹¹⁸

Corresponding to this debate and the explanatory tools utilised in the practical syllogism, we might make the following distinction. A mental state that is belief-like is one whose role is to represent, or help in representing, the world: beliefs, knowledge and perception. A desire-like mental state is one whose role constitutes our responses to the world: desires, attitudes and sentiments. This distinction is similar to that suggested by Anscombe and discussed by Michael Smith and Mark Platts. Platts writes:

The distinction is in terms of the *direction of fit* of mental states to the world. Beliefs aim at being true, and their being true is their fitting the world; falsity is a decisive failing in a belief, and false beliefs should be discarded; beliefs should be changed to fit with the world, not vice versa. Desires aim at realization, and their realization is the world fitting with them; the fact that the indicative content of a desire is not realized in the world is not yet a failing *in the desire*; and not yet any reason to discard

¹¹⁸ It is here that we run into the issues raised by McDowell's attack on the 'eighteenth-century philosophy of mind'. It might be suggested that if we reject the idea that truth is the basis of the debate then the debate (if it is about anything) will be about the relative priority of belief-like states and desire-like states in moral judgement. But if this is the case, then it seems difficult to deny the distinction between belief-like states and desire-like states and still purport to be in the debate. See the final section of this chapter below.

the desire; the world, crudely, should be changed to fit with our desires,
not vice versa.¹¹⁹

The justification for assuming this distinction lies with the considerations in chapter one concerning the kind of psychological explanations with which we are concerned. Although there is a vast array of conceptual resources that can be employed in explaining behaviour, this broad distinction of roles does have *prima facie* plausibility.

B. Explaining Attention.

Wiggins' suggestions about the practical syllogism led, as he put it, to our explanations 'giving out' and particularly, with "everything that is hard taking refuge in the notion of situational appreciation". The view that I have been urging grants the reciprocal relation between an agent's concerns and his perception of how things are but seeks to explain those perceptions of how things are in terms of the agent's concerns *and* so that they issue in judgements of a particular kind, namely judgements of obligation. The grip that we can get on this kind of explanation, I have suggested, is given through the general feature of the experience which accompanies these judgements — the pressure. The idea is that this experience gives us a way of breaking into the circle that is formed by the agent's concerns and his view of how things are.

As I think both Wiggins and Aristotle would agree, the agent's 'take' and the agent's appreciation of the situation are effectively the same thing — the product of the agent's ability to attend to his context in a particular way. In our case this is a

¹¹⁹ Platts, Mark. *Ways of Meaning*. London: Routledge and Kegan Paul, 1979: 257. See also, Smith, Michael. *The Moral Problem*. Oxford: Blackwell, 1984: 111ff. and Humberstone, I. L. 'Directions of Fit' *Mind* (1992): 59-83.

consequence of the discussions of chapter three. We might still prefer to use 'take' in the place of 'situational appreciation' since the latter could be thought to bring with it some unwarranted (or, at least, unjustified) normativity. We might be tempted to think that someone could fail to appreciate a situation or, at least, fail to appreciate a situation fully. The search might then be for an account of the ability which issued in a full appreciation rather than for an account of the ability which issued in the experience of the pressure. It seems reasonably clear that Aristotle (as he is standardly understood) and perhaps Wiggins would not be concerned about this involvement of normativity but this distinguishes their projects from this one.

Again, what puts us in a position to give at least a partial explanation of this ability to attend is the feature which accompanies the product of this ability. Whatever this ability amounts to in the end, its product, the judging, is accompanied by the pressure. Thus, the investigation of the ability and the possible explanations that we can give of it must be such as to yield the relevant kind of accompanying experience. The explanation will only be a partial one because, again, we are only concerned with the 'structural' conditions which enable or make possible an ability which issues in a judgement accompanied by the pressure.

The problem which confronts us is not a problem about perception. Instead, it is, we might say, a problem about attention¹²⁰. It is illuminating to cast the matter in terms of examples like Wittgenstein's duck-rabbit¹²¹. Thus, the issue is not one of 'seeing' but of

¹²⁰ I do not mean 'attention' here in a strong sense akin to 'concentrating' or 'focussing'. Instead, it is more like 'noticing' or 'acknowledging'. I take it that this is implied in the notion of a 'take' and also in 'situational appreciation'.

¹²¹ Of note here and for what follows in chapters six and seven are the papers by Peggy DesAutels ('Gestalt Shifts in Moral Perception') and Ruth Garrett Millikan ('Pushmi-

‘seeing as’. The question is not about how one sees the “tangle of lines” as a rabbit or a duck *at all*, but how it is that one sees it as a duck, say, *rather* than a rabbit. Alternatively, we might ask ‘why one sees it as a duck *first*?’ or ‘why a duck in preference to a rabbit?’. Of course the idea is that seeing a context as being one way rather than another is analogous to initially seeing the diagram as a duck rather than as rabbit. In the moral judgement case, seeing a context one way leads to the pressure and a moral judgement whereas seeing it in another way might not lead to the pressure or it might lead to a different pressure, and so, to a different judgement. In the simple case, if I do not see the old lady as struggling with her groceries then I cannot feel any pressure to assist her in the struggle. There is a fairly clear parallel here with the issue of act descriptions. Without being aware of the old lady’s struggle, I could not think of my lending a hand as helping her with that struggle. Indeed we might go so far as to say that the act of my helping the old lady in her struggle could not be that act unless I was aware of her struggle. Having said this, what I do and how we describe what I do is not directly to the point¹²². We are concerned with the pressure which may or may not lead to action and a particular act description. Of course recognizing the old lady’s struggle does not

pullyu Representations’) in May, L., Friedman, M. and Clark, A. (eds.), *Mind and Morals: Essays on Ethics and Cognitive Science*. Cambridge, MA: MIT Press, 1996. There are clearly other examples of this kind besides the duck rabbit. See, for instance, Gregory, R. L. ‘Perceptions as Hypotheses’ in Brown, S. C. (ed.) *Philosophy of Psychology*. London: Macmillan, 1974: 195-210.

¹²² The issues of act-description may still enter. We might still say that there is a relation between particular instances of the pressure and what I *would* do were I to yield to the pressure. However this involves the characterisation of the pressure as being a pressure to perform a particular action. This is not a problem, it simply involves us in matters of content. The current point is only that the agent’s ‘take’ involves seeing the context as being a certain way and that there is a correlation between this way and the way that the agent describes either the actual or potential act.

guarantee that I feel pressured to help. I may go on to conclude that the struggle is deserved, that it serves her right for not getting the groceries delivered. That is, I may not think that the broader context is such that it warrants my assistance. This broader view is, in turn, dependent on what I see the broader context as being.¹²³

The analogy between the ‘duck-rabbit’ examples and an agent’s ‘take’ brings with it a number of useful consequences. In order to consider why a particular agent sees the tangle of lines as a duck rather than a rabbit we must assume that the agent is able to see it as both¹²⁴. It is important that, for the most part, people (with a usual amount of knowledge and adequately functioning sense organs) are able to see the diagram as a duck or a rabbit. Most people know what ducks and rabbits look like. Similarly in moral cases the vast majority of people are able to see contexts in various different ways (although it does not follow that all people will experience the pressure after altering how they attend to a context). Thus, it would be quite rare to find someone who was not able (in the sense that they did not have the requisite conceptual tools) to see that the old lady was struggling with her groceries.¹²⁵ In this simplified case, then, people generally are able to see the context as one which includes the old lady as struggling and as one which does not. An example of the latter sort would be failing to notice the old lady at all.

¹²³ This last remark is intended to register the idea that the broader context needs to be *seen as* being a certain way also.

¹²⁴ Of course if the agent is unable to see the diagram in one of the two ways then the question of why it is seen one way rather than another is moot. Notice that in such a case a problem about perception might still arise. Namely, how/why the agent sees the tangle of lines as anything at all?

¹²⁵ Notice that we would be begging the question if we assumed that people could see that help was required — the judgement that help is required is one that is accompanied by the pressure.

Further, there are two ways of going wrong on this score — that is, two ways of failing to judge that the one ought to help the old lady: (i) by failing to notice that the old lady was struggling (or by failing to notice the old lady) and (ii) by failing to be moved by the old lady's struggle (for example, by failing to experience the pressure in the face of the struggle). With the help of the analogy, the former of these conditions will not, for the most part, be the result of simply not being *able* to notice (it will be a failure to attend). We shall have cause, below, to return these issues. The question, again, is how to explain how it is that agents come to attend to contexts in such a way that they experience the pressure (rather than attending in another way, which does not yield the experience of the pressure).

Unfortunately, the analogy between the duck-rabbit case and the moral case breaks down. The two ways in which it fails, however, are illuminating. First, and quite clearly, we are more inclined to think that which of the alternative ways of seeing a context an agent adopts, *matters* in the moral case. For the most part, and certainly in the simplified case above, failing to notice that the old lady is struggling is perhaps a (moral) failing of some kind. If we are not inclined to call this failing a moral one then we might say that it is, for instance, preferable for an agent to notice the old lady's struggle. Either way, our calling it a failing at all presumes that we have already taken a normative stance. In the duck-rabbit case it does not matter which alternative an agent sees first. That is, we are less inclined (if at all) to adopt a normative stance with respect to the duck-rabbit case¹²⁶. Alternatively, in the 'moral' case, how the context is seen often determines what agents *do*. People often act on the basis of what they take their context to be. It would be

¹²⁶ Notice the connection with indirect judgements discussed in chapter one.

rare, however, for a ‘viewing’ of the duck-rabbit diagram to result in anything but mild intrigue.

The second way in which the analogy breaks down is somewhat parasitic on the first. In the duck-rabbit case we might not expect there to be an explanation (or perhaps just not an interesting explanation) of why someone tends to see it one way rather than another. It might seem accidental or random. In the moral case we are loathe to describe the agent’s attention as either accidental or random, largely because, as we have said, which way one sees a context *matters*. Since it matters in the moral case we are more inclined to seek or to be interested in an explanation. Additionally we might not expect an explanation in the duck-rabbit case precisely because it is random or accidental — we might be tempted to think that the outcome is not something for which the agent is responsible or over which the agent has control. Responsibility, being tied to blame and blameworthiness, is a key feature in the moral case but is absent in the duck-rabbit case. Ordinarily, someone can be blamed (held responsible) for not noticing that the old lady was struggling but not for noticing the duck before the rabbit.¹²⁷

We can however imagine a duck-rabbit case where we might expect there to be an explanation or at least think that there is a somewhat interesting explanation to be given. Suppose that someone consistently noticed rabbits, not just instead of ducks and not just in curious diagrams, but more generally. We might take the various manifestations of this ‘attentiveness’ to suggest some attraction to rabbits and then to use this attraction to explain why the diagram was seen as a rabbit rather than a duck. Alternatively, it might turn out that the person knows a good deal more about rabbits than about ducks —

¹²⁷ Think here of the retort: “How could you not have seen her struggling?”

through study or experience. This training, then, might be used to explain both the general tendency and this specific case. Both of these kinds of explanations are hauntingly familiar and look a good deal like the sort of situation and explanation we might expect in the moral cases. Perhaps then part of the dis-analogy between the two arises from our taking the 'attentiveness' in the case of the old lady to be part of a pattern, reflecting an ability which manifests itself across a large number and variety of cases and as such, demanding explanation. In comparison to this the duck-rabbit case is a 'one-off', isolated curiosity.

In spite of these various differences the point of the analogy stands. What we are concerned to explain in this moral case is attention rather than perception, where perception and attention are understood along the lines of the distinction between 'seeing' and 'seeing as' as it exemplified in the duck-rabbit example. Moreover, it looks as though the vast majority of those who are considered moral agents possess the requisite perceptual tools to be *able* to see in the required way. Although there are dis-analogies between the duck-rabbit and the moral case these look to be intimately connected with our preparedness to attribute some kind of responsibility (in virtue of what we take to matter) and hence our interest in finding an explanation. Another way of seeing the point of this section is as an initial investigation of the distinction between perception and aisthesis. Given the ideas involved in the agent's 'take', I have suggested that the 'seeing'/'seeing as' distinction would seem to apply here.¹²⁸

¹²⁸ There is a link here to the analogy with secondary properties. These considerations provide a way of by-passing the analogy without necessarily giving up on cognitivism (although this remains to be seen). The secondary properties analogy looks reasonable when we are considering the perception of features rather than the selecting or noticing of

C. Cognitivism and Sensibility.

Why think that the cognitivist can provide an adequate account of the kind of attention with which we are interested? (This is not to suggest that there is any reason for thinking that the cognitivist is in an antecedently worse position than the non-cognitivist.) Above, in making the duck-rabbit case one where we might seek an explanation, I mentioned two possible approaches. On the one hand, the contextual evidence might suggest that a person has a particular *attraction* to rabbits as opposed to ducks and this is why they see the diagram as a rabbit first. On the other, the person might have had more exposure to rabbits rather than ducks. The person's 'training', understood either in terms of knowledge or experience or both, might explain the tendency to see the diagram as a rabbit. If the cognitivist is to provide us with an adequate account in the moral case, we might suspect that it will bear some resemblance to this second way of making the duck-rabbit case interesting (perhaps to the exclusion of the first).

Let us consider an example where this kind of story might plausibly be told. A geologist looking at a cliff face might very well see the different periods of time that each of the layers of rock represents. Just from looking at these layers the geologist gets a picture of the history of that particular area. The lay-person sees layers of rock. (In my case, I see layers of rock and know that a geologist could tell me what each of them represents.) Of course the geologist and the lay-person are 'seeing' the same thing, the difference lies in what it is being seen as and this looks to be a result of what is known — the geologist knows about geology but the lay-person does not. Notice that what is

features. Of course the justification, now for a particular way of attending, will look very different.

known affects the attention (how the cliff face is seen) in two ways. That is, there are two ways in which the lay-person's lack of knowledge of geology affects the way in which the cliff face is seen. The geologist can see the cliff face as a picture of the history of that area but also, notices details which the lay-person does not. The lay-person, for instance, might not notice that one of the sandstone layers of the cliff contains small pieces of fossilized shells. Now of course the lay-person could see the shells but without direction might not notice them. Importantly, the lay-person can or could recognize the shell fragments but because of a lack of knowledge, does not.

We might tell the following story about this case. Our geologist went to graduate school where he learnt all about rocks and rock formations. In the course of learning and subsequently applying this knowledge he has come into contact with many variations of rock formations and countless rocks. As a result of his study and experience he is very good at noticing geological features (and he does so often). Indeed he rarely, if ever, fails to notice geological features. What is important here is that the geologist has a blend of knowledge of geology and experience — he has practiced applying his knowledge in cases and is in the habit of doing so regularly. One might suspect that the knowledge and the ability to apply it — and hence, the ability to notice certain features rather than others — go hand in hand. That is, one could have the knowledge without the ability — this would signal a lack of experience — but one could not have the ability without the knowledge.

There are then two differences between the geologist and the lay-person (corresponding to the two ways in which the lay-person's lack of knowledge effects the way the cliff face is seen). First, the lay-person is limited in the ways he can describe the

cliff face. He cannot read the cliff face as a history of the area, i.e., cannot describe the significance of each layer of rock. The lay-person is limited in what the cliff face can be 'seen as' because of limited knowledge. Second, the lay-person does not (as opposed to 'cannot') notice details that the geologist does, again, because of limited knowledge. In the latter case, the lack of knowledge on the part of the lay-person is the combination of a lack of knowledge of geology and the 'geological experience' which comes from practice and habituation in applying knowledge of geology.

There are two immediately apparent differences between this and the moral case. First, given what has been said in chapter three, in the moral case it seems unlikely that there is anything like the definite body of knowledge in geology that can be studied and applied. Second, it is not clear that there can be anything like the distinction between lay-person and expert that there is in the geology case. It is not as though the virtuous have at their disposal a whole set of descriptions that are simply not available to the non-virtuous. (This, of course can be seen to follow from there being no definitive body of moral knowledge.) The evaluative concepts that are used to describe a context (the features/properties selected in the formation of a 'take') are available to most (if not all) competent language users. This is the point that was made above about being able to see the tangle of lines as both a duck and a rabbit.

Neither of these differences are devastating but they do require adjustment if this story is to be applicable in the moral case. If the evaluative concepts, and hence the knowledge, are available to most then the 'studying' aspect of moral education must be downplayed. In its place, more emphasis is to be accorded to experience, practice and habituation — precisely as is done by Aristotle. What matters on this view is not that the

agent has poured over, as it were, the texts of normative moral theory but that the evaluative concepts and critical moral tools have been exemplified and applied over the course of the agent's upbringing. Thus, the agent comes to grasp the relevant concepts and tools by a piecemeal and gradual process of being exposed to their application accompanied by the agent's own attempts to do the same. What is thereby developed is a proper understanding of the concepts and a developed ability to recognise/notice/attend to contexts in a certain way. An agent who is not fortunate enough to have such an upbringing may well grasp the evaluative concepts and tools but, without the practical processes gained through exemplification, practice and experience, may fail to have developed the ability to attend in the requisite way. This is similar to the geology case where we saw that someone may 'know the geology' but, through lack of experience, practice and habituation, not have cultivated the ability to attend to geological features.

Having said all this, the non-cognitivist would be quick to point out (and be justified in doing so) that, we are still missing the most significant ingredient which separates the geology case from the moral one — that of motivation. In the moral case, judgements of obligation involve a call to action of some sort but the geologist is not in any way called to act or to respond to what he notices. In terms of the pressure, any judgement that the geologist makes on the basis of his knowledge and experience certainly seems to be a product of the way the world is and if we stretch the story a little, is experienced as 'objective' in just the way that the pressure is experienced. Nothing however, and this is the nature of the case, pulls or draws the geologist to act. The geologist's judgements about the cliff face are not practical in the way that judgements of obligation are. The second aspect of this account that the non-cognitivist would jump on

is the reference to the developed 'proper' understanding. In the geology case we presumably have a well developed and tested theory that explains, in the usual scientific way, the way the world is. There does not appear to be any obvious correlate in the moral case. What is required is some kind of theoretical base or core of conceptual and evaluative tools the possession of which constitute a 'proper' understanding.

What both of these objections require is an explanation, on the part of the cognitivist, of the nature of the evaluative concepts that constitute what is known or grasped by the 'virtuous' person. This explanation must be such as to include the requisite motivation and provide the tools for an account of the 'proper' understanding of these evaluative concepts. Just such an attempt is made by McDowell and Wiggins.¹²⁹

The central idea of the McDowell/Wiggins kind of story is that the kind of evaluative tools/concepts that are involved here are such that what one learns (and hence what one knows) — the perceptual capacity that is developed — cannot be separated, either in the learning or in the full possession of the capacity, from the kinds of responses or concerns that the properties involved elicit. That is, as Wiggins might put it, the properties and responses are inextricably paired together so that the full and proper understanding of the term used to designate the property brings with it a propensity to respond in a particular way. Thus, he writes that

If a property and an attitude are made for one another, it will be strange
for one to use the term for the property if he is in no way party to the

¹²⁹ We will concentrate on the story that Wiggins gives, particularly in 'A Sensible Subjectivism' (*op.cit.*, pp.185-211), largely because, as McDowell acknowledges, what he urges "is more negative than positive" (*Mind, Value and Reality.*, p.viii).

attitude and there is simply no chance of his finding that the item in question has the property. (p.199)

How does this inextricable pairing go? We begin with a whole slew of regular responses that we have to certain properties of those contexts in which we find ourselves. These properties get grouped together under one heading or name because they elicit a certain kind of regular response from us. Because the property-grouping depends on our reaction, the names for the groupings are “avowedly anthropocentric” (p.195). Also, because these features are gathered together under a name in virtue of our responses, there will not necessarily be any way to say what the names stand for independently of the particular reaction.

In disputes about whether something is funny (Wiggins’ example), there are many considerations that can be brought to bear on the question without these considerations adding up to or in some way amounting to a reduction of what it is for something to be funny. For instance, we might point out the unusual, curious, unexpected, ironic, etc., juxtaposition of features of a particular context. It is by means of these kinds of considerations that individuals can learn, develop and refine their sense of what is funny. They can become better at picking out funny things, making finer distinctions between funny things as well as finding different, perhaps more subtle things funny. In this way the considerations affect both the responses and the ability to discern the property.

What emerges from this process is a complex practice involving the use (the proper use) of a large and varied set of evaluative terms, the ability to isolate things or situations which contain the properties to which these terms refer, the responses to these properties and the patterns of justification and criticism by which both the presence or

identification of the properties and appropriateness of the corresponding responses are assessed.¹³⁰ The two most important features of this practice are the story that it provides about what someone grasps when they properly understand these evaluative terms and the way in which this ‘proper’ understanding is justified.

First, we can see from this how the use of an evaluative term is related to a property/response pair. The agent who is likely to use a term in a particular context will be one who is in some way party to the attitude and is able to find that “the item in question has the property”. This is part and parcel of what one learns when one learns how to use an evaluative term, in particular, and, more generally, how the network of evaluative terms fits together with the critical and explanatory tools that are available. Importantly, one is caught up in a social context, a “social morality”, in virtue of ‘where one is’ — that is, being involved in the social practices of relating to others. As such we are involved and participate in the “shared sentiments and responses that are presupposed to my having grasped the sense of the moral and political language that my membership in the social morality commits me to express myself in.”¹³¹ The key claim here is that the

¹³⁰ The example of ‘the funny’ does exemplify the process in question, however, we can point to an important difference between it and, what are usually considered to be, moral evaluative terms. In the case of a dispute about X’s being funny, being convinced that it is indeed funny will not normally result in the same kind of response as if one had found X funny in the first place. Think here of what tends to happen when a joke is explained. It does not have the same impact as it might have done. The distinction is between finding something funny and understanding why something is funny. In what we might take to be a paradigmatically moral case, say of someone’s needing help, being convinced that a third party needs or needed help is to come to find the situation to be one in which the person needs help and to understanding the situation as being that way. As Wiggins puts it “*finding* *x* to be Ψ is prior to (or at least coeval with) *thinking* *x* to be Ψ .” (p.202) My point is that for many of the evaluative terms that are often taken to be moral ones, the ‘finding’ and the ‘thinking’ are closer than they are in the case of the funny.

¹³¹ Wiggins, David. ‘Cognitivism, Naturalism, and Normativity: A Reply to Peter Railton.’ in Haldane, John. and Wright, Crispin. *Reality, Representation and Projection*.

story about the development of the concept, the property/response pairing, amounts to an explanation of what it is to come to understand the meaning of the evaluative term. In this way we can see the connection back to chapter three. The general problem raised there was how to explain the relationship between the possession of the rule and the pressure. By involving shared sentiments and responses along with the range of critical resources and the whole network of evaluative terms within the grasp of particular evaluative terms, the agent who possess an evaluative term and is likely to utilise it in particular contexts will also be “party to the [accompanying] attitude” and will find “that the item in question has the property”. Thus, instead of the rule being straightforwardly about one’s obligations, this approach develops a picture in which the rule that is grasped is at the level of concepts which influence the way in which an agent forms ‘takes’, i.e. at the level of the way in which an agent construes the situation that confronts him. The pressure on this view is a result of the sentiment/response that is embedded in the possession of the concept, caught up in the meaning of the term and the agent’s tendency to use it as applying to particular cases.

Second, in order to show how the appropriate understanding that one grasps is to be thought of as the ‘proper’ understanding Wiggins draws an analogy between evaluative terms and mathematics. The thought is that “*genuinely* funny things are things that not only amuse but have this effect precisely because they are funny” and that this is analogous to (“something like”) the way in which our tendency to think that $7 + 5 = 12$ is explained by there really being nothing else to think about what $7 + 5$ is.¹³² The parallel

New York: Oxford University Press, 1993 (301-13): 307. Hereafter referred to as ‘Naturalism’.

¹³² Wiggins, David. *Needs, Values, Truth*. (*op.cit*), pp.199-200.

between mathematics and the evaluative is an important one here for the cognitivist. We take arithmetical statements to admit of plain truth (as Wiggins would put it) and have various ways of demonstrating and explaining their truth all of which lie within the practice of mathematics. These explanations, importantly, play the double role of explaining and justifying — they explain the truth of $7 + 5 = 12$ and they justify our tendency to think that $7 + 5 = 12$. They are vindicatory explanations. So, the considerations (proofs, calculations, etc.) that can be adduced in favour of $7 + 5 = 12$ demonstrate that *there is nothing else to think* and this, in turn, justifies our thinking that $7 + 5 = 12$.

The case of evaluative terms is, for Wiggins, something like this. He uses the example of slavery. Here, there are a wealth of considerations that can be produced to show that slavery is unjust and insupportable and these considerations are such as to show that there is nothing else to think but that slavery is unjust and insupportable.¹³³ Raimond Gaita discusses slavery in a way that well illustrates the point here.¹³⁴ Speaking of a passage that he has just cited of Cavell's,¹³⁵ he writes

Cavell means, I think, that the slave owner denies that the slave has his
kind (the slave owner's kind) of individuality: the kind of individuality
that shows itself in our revulsion in being numbered rather than being

¹³³ Wiggins, David. 'Moral Cognitivism, Moral Relativism and Motivating Moral Beliefs'. *Proceedings of the Aristotelian Society* (vol. 91), 1990: 61-85. Hereafter referred to as 'Moral Cognitivism'. See particularly, p.71 and note 10 for Wiggins' suggestions about these kinds of considerations concerning slavery and the slave trade.

¹³⁴ Gaita, Raimond. *Good and Evil: An Absolute Conception*. London: Macmillan Press, 1991: 156-65.

¹³⁵ Cavell, Stanley. *The Claim of Reason*. Oxford: Clarendon Press, 1979: 376.

called by name; the kind of individuality that gives human beings the power to haunt those who have wronged them, in remorse. If the slave owner could be haunted by the slave girl he raped, then her days as a slave would be numbered. (p.159)

What this passage and the rest of Gaita's discussion brings out is how the *concepts* of 'individuality', 'human being' and 'person' themselves are invested with an emotional and responsive power that cannot be disentangled from them. These kinds of emotions and responses are caught up with the idea of taking someone seriously as an individual. They illustrate the full force, not of what it is to *treat* someone as a human being, but of conceiving of them as 'other to my one'.¹³⁶ This kind of investigation of, as we might say, the psychology of the slave owners shows how their understanding of those who were their slaves failed to approach their conception of those who were 'fully human'. It trades on an understanding of the gamut of our evaluative and emotional resources and seeks to show how these resources cohere. In the face of this understanding and in the face of this diagnosis of the slave owner we are to see that there is nothing else to think.

That his slave girl was a Christian and not merely a 'Christian', or that she was a proper respondent to love's claims, or that her body could invite, and worthily receive, a tender caress, ... If he came to see his slave girl like this then we would say that he sees them as human beings and we would be right to say that he had 'come to see things as they are', but the grammar of that should not be misunderstood. What he must learn has to do not with facts and the consistent application of principles, but with

¹³⁶ This is Cavell's phrase (*loc.cit.*), discussed by Gaita at (*op.cit.*) p.159.

meaning — with the meaning that the lives of his slaves can have, with what they can understand, feel and do, therefore, with what they can be.
(p.165)

Importantly, what is shown by these kinds of considerations is not that there is nothing else *for us* to think but that there is nothing else to think, simpliciter. This is a point about the way in which the considerations function. The thought is that the arguments, etc. are mounted from the level of a perspective that is not particular, that aims at (and, Wiggins thinks, sometimes achieves¹³⁷) an abstraction from the ‘for us’ to one involving simply one person to another. What is shown, then, by these arguments is that the price of not thinking that slavery is unjust and insupportable is that one opts out of the point of view that is common between people and which is aimed at by the considerations and arguments brought forward. This common point of view is not culturally or socially specific even though particular modes of expression and, as it were, conceptual organization might be.

The issues surrounding the kind of justification being offered by Wiggins will be discussed further below, however it is instructive, at this point, to say something regarding ‘opting out’ and relativism. Both the possibility of opting out of one’s own ‘moral system’ and worries about relativism stem from our recognition (either through reflection or through confrontation with alternative ‘systems’) of “the apparent contingency of the fact that it is these practices [my own] rather than some other practices I find myself in the midst of, and the consequential contingency of the fact that it is this

¹³⁷ This is one of Wiggins’ central tasks in ‘Moral Cognitivism’, pp.68ff. See below: section D.

rather than that I find that I am expected to do in circumstances *C*.”¹³⁸ The case of the felt contingency of an individual’s own moral system can be seen in some (if not all) cases of ‘teenage angst’. Here we might find the teenager questioning the set of values that he has been given and that are directly associated with parents and their generation.¹³⁹ Faced with this apparent contingency the individual (the teenager) has three options: “outright rejection”, “piecemeal criticism” and “simple acceptance”. Importantly, Wiggins only needs to demonstrate that the simple acceptance option is a reasonable one. He need not (and does not) deny that the first and second options are rational. This is a result of the task that he has set for the cognitivist: to show that moral judgements aspire to truth and in some cases achieve it. It need only be reasonable that an individual accept the set of practices he has been given. Thus Wiggins points out that

I accept these practices because I find that no other alternatives come to life for me ... [or] I may reflect that my own life is rooted in these practices and scarcely meaningful to me outside them.¹⁴⁰

The obvious source of the worries about relativism is our interaction with different cultures. To generate relativism this interaction can be at the level of a direct encounter or more commonly by simply imagining interaction. Again there are three

¹³⁸ Wiggins, ‘Naturalism’ (*op.cit.*), p.308.

¹³⁹ This might be a little generous for the run-of-the-mill angst-ridden teenager.

¹⁴⁰ *ibid.*, p.308. There is one reasonably important qualification on all of this. Simple acceptance is unreasonable if “the social order that sustains and is sustained by these practices is manifestly unjust”. This warrants some explanation since it seems to presume that there is a standard of justice that is independent of or accessible independently of any system of practices. Wiggins refers the reader to chapter one of *Needs, Value, Truth*.

options: incommensurability, perseverance and underdetermination.¹⁴¹ First it may be that the parties to the dispute come from such different ‘forms of life’ (civilizations or cultures) that “any semblance of agreement on the question what one ought to do or what is good is only a semblance.”¹⁴² In such a case,

the fact of them seeming to disagree cannot count against the cognitivist expectation that there will be convergence in belief among those who *do* understand the same thing. (‘Moral Cognitivism’, p.75)

Second, there may be a genuine question that both parties do share and where disagreement is non-trivial. Here Wiggins urges perseverance. The thought is that by uncovering, as it were, the ‘deep’ differences of perspective, both conceptual and responsive, one or both of the parties can come to see or understand the other’s standpoint. This seems a fairly common occurrence — one finds oneself (suddenly) seeing or feeling what it is like from the ‘other side’. The idea here is that perseverance

¹⁴¹ Wiggins, ‘Moral Cognitivism’, p.77. A good example here is the Azande and the discussions of them given by Evans-Pritchard and Winch. Winch takes Evans-Pritchard (*Witchcraft, Oracles and Magic among the Azande*. Oxford: Clarendon Press, 1937.) to task in his paper ‘Understanding a Primitive Society’ (*Ethics and Action*. London: Routledge and Kegan Paul, 1972: 8-49.) What is nice about this case is that we can imagine each of the three sorts of dispute arising between, say, Evans-Pritchard and one of the Azande tribesman. If Winch is correct in his criticisms of Evans-Pritchard (cf. the first part of Winch’s paper) then there is enough cultural divergence to provide cases of incommensurability. Similarly, there will be cases where perseverance would be productive and that are simply underdetermined — cf. the concluding section of Winch’s paper.

¹⁴² The example that Wiggins gestures to at this point is Peter Winch’s discussion of Billy Budd in ‘The Universalizability of Moral Judgements’ (*Ethics and Action*. London: Routledge and Kegan Paul, 1972: 151-70.) The relevant (though much too simplified) idea is that Winch could not do what Captain Vere did but that this does not mean that what Captain Vere did was wrong. It is a direct challenge to the claim that if I judge that *X* ought to be done then I think that anyone else in my position ought to do *X*. Winch particularly targets the question of when someone else is to be counted as being in my position.

would or could achieve agreement or a resolution of the dispute. Whether such a strategy works or not is more a matter of the individuals concerned and the circumstances of their involvement. In the enterprise of trying to understand one another egos, temperament, claims of authority and attitudes of superiority can get in the way. Finally, there may be situations where, although there is a real issue, a real question of what ought to be done, it is not clear that any of the parties is in a position to decide and to be justified — there is no manifest possibility of any winning set of considerations ever being mustered.”¹⁴³

Much more could, and indeed should, be said about Wiggins’ idea of there being nothing else to think and its connections with relativism and truth. To do so here would take us too far afield. However to assist immediately it is perhaps useful to contrast briefly the way ‘there is nothing else to think’ is intended to function (in ethics) with the way in which G. E. Moore uses ‘self-evident’ in *Principia Ethica*¹⁴⁴. There we are given the following account of self-evidence:

The expression ‘self-evident’ means properly that the proposition so called is evident or true, *by itself*, alone; that it is not an inference from some proposition other than *itself*. The expression does *not* mean that the proposition is true, because it is evident to you or me or all mankind,

¹⁴³ This final kind of solution can, I think, be illustrated in a non-dispute case where one simply cannot unravel the competing claims in an imagined case enough to be in a position to be able to say that this is what one ought to do or that such and such is right. This, I take it, is an example of a similar kind of indeterminacy or underdetermination to Wiggins’. Notice that this is not a case where two competing judgements are immeasurable but where we feel unable even to begin to decide what we would/ought to do. This has always been my reaction to cases like Williams’ Jim and the Indians (see Smart, J. J. C. and Williams, B. *Utilitarianism: For and Against*. Cambridge: Cambridge University Press, 1973: 98-9).

¹⁴⁴ G.E. Moore, *Principia Ethica* (1903). Cambridge: Cambridge University Press, 1962. (Hereafter, referred to as ‘PE’.)

because in other words it appears to us to be true. ... By saying that a proposition is self-evident, we mean emphatically that its appearing so to us, is *not* the reason why it is true: for we mean that it has absolutely no reason. (*PE*, p.143)

The most significant aspect of this account is that the claim of self-evidence *proves* nothing. An individual asserting that a particular proposition is self-evident, in effect, makes two claims: that the proposition is true and that no reason can be given that proves that it is true. As a result of the second claim, disagreements about what is self-evident and what is not (or what is self-evidently untrue) are simply conflicts of intuition¹⁴⁵.

One point of convergence between Moore and Wiggins here involves the relationship between their respective ideas and inconceivability. Neither Moore nor Wiggins intend their 'term' to be related to what is inconceivable. For Moore, something "would not be a self-evident proposition if we could say of it: I cannot think otherwise and therefore it is true." Similarly Wiggins is neither concerned with what one cannot think or with what is necessarily true.¹⁴⁶

On the face of it, however, the status of 'propositions'¹⁴⁷ for which 'there is nothing else to think' differ from those that are self-evident on the second claim — that no reason can be given which proves that the proposition in question is true. As we have

¹⁴⁵ Moore finds one such conflict with Sidgwick in chapter three of *Principia Ethica*.

¹⁴⁶ See Wiggins, 'Moral Cognitivism', p.70. This is also pointed out by A. W. Moore in his paper 'On There Being Nothing Else to Think, or Want or Do' (in Lovibond, S and Williams, S.G. (eds.) *Identity, Truth and Value: Essays for David Wiggins*. Oxford: Blackwell, 1996: 166-7) and agreed to by Wiggins in his reply (p.272). Moore also cites Wiggins in 'Moral Cognitivism', note 7, p.67.

¹⁴⁷ Wiggins, I'm sure, would not be particularly comfortable with talk of propositions — hence, the scare quotes.

seen, Wiggins is primarily concerned with vindicatory explanations. Since these appear to involve the giving of reasons (indeed the ‘wealth of considerations’ that can be brought to bear would appear to count as the giving of reasons), there being nothing else to think but that *p* seems clearly to differ from claiming that *p* is self-evident. Moore, however, adds an important qualification. Referring to the fundamental proposition of Ethical Hedonism he writes “I shall try to shew you why my intuition denies it, just as his intuition affirms it. ... I am bound to be satisfied, if I can [citing Mill] ‘present considerations capable of determining the intellect’ to reject it” (*PE*, p.75). Later, in chapter five, this point is reiterated, again with reference to Ethical Hedonism:

I could do nothing to *prove* that it was untrue; I could only point out as clearly as possible what it means, and how it contradicts other propositions which appear to be equally true. (*PE*, pp.144-5)

These two passages would appear to make Wiggins’ standard much closer to that of Moore’s. Moore thinks that there are considerations that can be produced, considerations involving meaning and (perhaps) holism, that can serve to convince, in this case, the Ethical Hedonist of their mistake. For Wiggins, as we have seen, “by drawing upon the full riches of our intersubjectivity and our shared understanding such a wealth of considerations can now be produced [that] it will become apparent that there is *nothing else to think ...*” (‘Moral Cognitivism’, p.70). Given the story about intersubjectivity and shared understanding — the story about evaluative concepts —, Wiggins’ standard seems also to rely on considerations involving meaning and a holism of conceptual resources.

What causes the initial confusion is Moore's restricted use of 'a reason' and 'prove'¹⁴⁸. He, in fact, thinks that there are three kinds of reasons. First, "a logical reason, or reason in the sense in which self-evident propositions have no reason, is a reason why *the proposition itself* must be true, not why we hold it so to be." (PE, p.143) The second kind of reason then is a reason why we hold a proposition to be true. These are reasons that are to be understood as "the evidence of a proposition" or "the *cause* of our asserting it". Third, "that a proposition is evident to us may not only be the reason why we do think or affirm it, it may even be a *reason* why we ought to think it or affirm it." (PE, p.143-4) Given these distinctions, there look to be plenty of reasons that are available for use in the defense of particular self-evident propositions. Interestingly, the two other kinds of reason that Moore's distinctions allow are (roughly) empirical and normative.

Whilst more work needs to be done to properly examine just how close the two ideas are, I do think that it is instructive to think of 'there being nothing else to think that *p*' along the lines of 'it is self-evident that *p*'. Having said this, we should be careful about how we think of self-evidence. It is tempting to think of claims of self-evidence as being, as it were, conversation stoppers — where the claimant, at least, regards the matter as closed and that nothing more can be added. But as Moore allows and as Wiggins spends some time suggesting, reasons can sometimes be produced that are persuasive (though they might not always persuade). Those cases where the proposed argument fails to persuade are then handled in the ways mentioned above. Wiggins insists that we may

¹⁴⁸ Moore's citation of Mill is intriguing. He does not say that he can produce reasons, rather, they are 'considerations capable of determining the intellect'. I think the best way to take this is as 'considerations capable of determining the intellect that are not logical reasons'. Thus, both 'non-reason' considerations and Moore's second and third kind of reason are available to demonstrate the truth of proposition that is claimed to be self-evident.

not be able to construct this kind of argument for many things of which we are convinced but, for the cognitivism that he is defending, what matters is only success in some of its aspirations to truth.¹⁴⁹ By the tone of the latter of the two quotations from Moore, he may well have agreed.

This comparison, I think, brings out an interesting point about the lineage of the Wiggins/McDowell story.¹⁵⁰ If the parallel between ‘there being nothing else to think’ and ‘self-evidence’ holds, they both sit within a tradition that includes, with Moore, Ross, Prichard and Carritt this century, as well as Sidgwick (perhaps), Reid, Price and Butler. And all of this fits well with those parts of Aristotle’s ethics that have a decidedly intuitionistic flavour. Moreover, Wiggins’ ‘sensible subjectivism’ can be understood as making the adjustments that were required after the initial attacks on intuitionism in the first half of last century — that is, adjustments prompted by the attack of emotivists and non-cognitivists like Ewing, Ayer, Stevenson and Mackie.

Having raised the historical connections, it is instructive to turn back, briefly, to the Aristotelian issues of moral education and upbringing. On the model of the geologist, we wanted an account of what the ‘virtuous’ person comes to know when they learn to be virtuous. This account could not involve a specific body of, as it were, propositional knowledge. This was a consequence of chapter three and represented an important difference between the geologist and the virtuous person. The natural move here was

¹⁴⁹ Again, this is not success in persuading but success in coming up with a vindictory explanation.

¹⁵⁰ Again I include McDowell here because I think that were he to focus more squarely on the issues that Wiggins does, he would tell (perhaps must tell) a similar story. Indeed McDowell does not have the connections to Hume that Wiggins does and so he might be a better fit with this lineage.

towards the kind of experiential knowledge that the geologist develops in the course of applying his theoretical knowledge and that is built up through time 'in the field', habituation and practice. We also required, after an initial objection by the non-cognitivist, that some room be made for the motivational aspect of moral judgements. This was the second difference between the geologist and the virtuous person. It was this difference that led us to consider Wiggins' suggestions regarding evaluative terms. Wiggins suggests that we understand evaluative terms as names for property/response pairs that develop within a complex set of critical and reflective practices. One who grasps the meaning of these terms both has an ability to isolate the relevant properties and will be party to the accompanying attitude. As Wiggins puts it, the resulting account is one

that arises from the way in which the whole range of judgements that a person makes are conditioned by a background of commitments and habits of concern. These are the commitments and habits that make it possible to find in the world what will give a meaning for his life and, in grasping the sense of moral language, to focus upon one set of saliences out of countless possible ones. (*Needs, Values, Truth*, p.331)

Perhaps in a more starkly Aristotelian tone, taking us back to section A of this chapter, McDowell writes that:

The result of habituation is a motivational tendency, but one with a conceptual and hence rational aspect. People with a properly formed character have learned to see certain actions as worth undertaking on the ground that they are noble; they have acquired that reason-giving concept,

in a way that is inextricably bound up with acquiring the propensity to be motivated by thoughts in which it is applied. The question of correctness is the question whether the actions they see as worth going in for in that specific way are really worth going in for in that way. We can resolve it into a series of piecemeal questions, whether this or that action is correctly seen as noble. And these piecemeal questions arise *within* the conceptual and motivational outlook that, according to Aristotle, ethical upbringing imparts.¹⁵¹

D. Two Points of Contention.

We can, I think, isolate a number of ways in which the non-cognitivist might attempt to find fault with a cognitivist position of the kind sketched above. However what should be remembered, and will become clearer, is that despite the differences in orientation and focus, the Wiggins/McDowell cognitivist position is not far at all from sophisticated forms of non-cognitivism, like the one we shall consider in the next chapter. As we shall see for instance, the account of justification, of what is to count as a 'proper' understanding of a concept, given above is particularly close to the kind of account of justification that is available to Blackburn's quasi-realist.¹⁵² Having said this, the key differences remain. The cognitivist position sketched above is firmly entrenched in the

¹⁵¹ McDowell, John. 'Deliberation and Moral Development in Aristotle's Ethics' (*op.cit.*), p.31.

¹⁵² Blackburn suspects that what is good in Wiggins story is a notational variant of his own position (*Ruling Passions*, p.99). It is easier to see the difference between McDowell and Blackburn because, as we have noticed, McDowell is primarily concerned to argue against non-cognitivism.

tradition of cognitivism and has its focus squarely on finding and accounting for the ways in which moral judgements admit of truth or, at least, have cognitive aspirations that are achieved in some cases.¹⁵³ The non-cognitivist thinks that moral judgements do not admit of truth, that they are, in complex ways, expressions of our attitudes and desires and must show how, being things of this kind, they can approximate and *appear* to have cognitive aspirations.

So on to the non-cognitivist's objections. There are three points on which the contentions will be focussed: (i) the issue of justification — how we come to have genuine standard of what is 'proper'; (ii) the kind of psychological account that is required — as it were, the 'rules' of psychological explanation; and (iii) the 'inextricability' of property and response that is claimed, by the cognitivist, to be built into evaluative concepts. I will address only the first two in this section, leaving the third for the next chapter. As will become apparent, the primary reason for this is that the first two objections end up being disputes that lie outside the scope of the current project whereas the third will lead us squarely into the non-cognitivist's position.

(i) Justification.

If there is to be any prospect for the kind of cognitivist account discussed in the previous section, there will need to be room for some story about what counts as a

¹⁵³ Wiggins, 'Moral Cognitivism', p.62ff. Of particular note also is Wiggins' paper 'Objective and Subjective in Ethics, with Two Postscripts about Truth.' (*Ratio* N.S.8. 1995: 243-58). There he distinguishes objective from non-objective and subjective from non-subjective. The point being that one can have an account that is both objective and subjective. Briefly, 'objective' is defined "in terms of there being nothing else to think" whilst 'subjective' is defined "in terms of answerability to a standard founded in sentiment" (p. 249).

‘proper’ understanding of an evaluative term. Wiggins puts this problem of justification in the following way. The non-cognitivist asks

how the cognitivist could show that just anyone who is willing to take up the point of view that shall be common between one person and another *must* subsume the institution of slavery under the concepts it is subsumed under in the argument for its injustice and insupportability. (‘Moral Cognitivism’, p.78)

[The italicised ‘must’ here is very important.] Wiggins’ response to this question is, in effect, a move to the whole system of evaluative terms. This makes the objection not one about slavery, in particular, but about the whole network of evaluative terms complete with the ‘wealth of considerations’ that can be called on and are utilised in the justification of particular judgements. The idea is this. The patterns of argumentation, evaluative terms, responses and properties all function together to justify, internal to these practices, the judgements and considerations that enable one to establish that there is nothing else to think. The challenge to show, over and above this network, that such and such a moral claim must be the case is a call to justify the whole practice in which the particular justification takes place. But now the objection can be moved up a level: if, in the end, at least some moral judgements will be true or false then the justification for them being so will need to transcend even the particularities of the network of evaluative terms and considerations that are used in particular cases. Otherwise the truth that any moral judgements attain will be relative to that particular network — i.e. not really truth but truth relative to the network. Whether someone takes a particular moral judgement to be true will depend on whether they happen to like/accept/adopt the network.

This is precisely the point at which, in terms of justification, this kind of cognitivist position diverges from the quasi-realist. The quasi-realist avails himself of all of the tools to which Wiggins appeals, it is just, in the end, a matter of attitude or stance which network of considerations we prefer or adopt. The very last point before the divergence is, I think, well captured by Wiggins:

Moral cognitivism cannot even pretend that it reposes upon an *a priori* foundation. The thesis is grounded, if in anything, then in the peculiar strength and dialectical and persuasive resources that can be mustered within the subject matter of morality itself.¹⁵⁴

The step that both Wiggins and McDowell take at this point is a familiar kind of move and one that has been made against the ‘epistemological’ skeptic. Recall that such a skeptic demands a justification or proof that we have knowledge of the ‘external’ world that is independent of our experience of it. The skeptic wants proof that the ‘external’ world really is as our experience tells us. The move which parallels that available to Wiggins in this case, is to deny the legitimacy of this demand: a proof of the kind that the skeptic demands cannot be given, not because there is no justification for our knowledge of the external world, but because the concepts of proof, justification or evidence arise within and are specific to the context of practices that are already up and running.¹⁵⁵ The ‘external world skeptic’ then, according to this response, is guilty of both seeking a

¹⁵⁴ Wiggins, ‘Moral Cognitivism’, p.79. Recall here the parallel with Moore’s notion of self-evidence in the previous section. If we understand Moore’s ‘logical reasons’ to include *a priori* reasons, then moral truths — being those for which there is nothing else to think or are self-evident and hence have no logical reasons — will not admit of *a priori* reasons.

¹⁵⁵ Notice here the parallels with Wittgenstein, Ludwig. *On Certainty* (1969). New York: Harper and Row, 1972.

justificatory standard that belongs to some other practice(s) (those in which *a priori* or deductive proofs are part of the practice — say mathematics or logic) and of seeking a standard that is independent of the justificatory standards internal to the practice.¹⁵⁶

Whatever we think about the prospects in the skeptic's case, Wiggins' account denies that such a justification can be given in ethics. Hence, the call for an independent story about why anyone *must* subsume the concept of slavery under the concepts suggested is an illegitimate one. Again, it is not that this is simply 'how we happen to do it' but it is part and parcel of the practice. Justification, and particularly, the kind of justificatory standard that is available, only makes sense internal to a practice. These practices, to be sure, arise because of our attitudes and responses (but not only because of

¹⁵⁶ The idea of a practice is doing a good deal of work in this discussion and requires comment (This was pointed out by Jonathan Adler). There are, I think, three components that are importantly involved within the idea of a 'practice' as it is being used here: (1) A practice will usually involve some form or part of discourse — this is certainly true of those under consideration here; (2) A practice will also have something like constitutive rules or conventions that will operate within it — rules about what counts as being a part of the practice or as being successful within it; and (3) The practice will involve some behavioural component — it will tend to be an activity.

In various ways McDowell, Wiggins and Blackburn exploit each of these in their respective accounts of the relevant terrain. In particular Wiggins' discussion of the development and enhancement of the basic property/ response pair in 'A Sensible Subjectivism?' (*op.cit.*, pp.195ff.) is a good example containing each of these features combined in a way consistent with this general idea of a practice.

There are two important 'traditions' behind this use of practice: MacIntyre and Wittgenstein. MacIntyre's definition in *After Virtue* ((2nd Ed.). London: Duckworth, 1985: 187.), though constructed with a different purpose in mind can certainly be taken to the use of the term here. I am reluctant without further discussion to co-opt the definition as it stands. Similarly Wittgenstein's idea of a language-game shares certain important features and resonates well with the concerns here.

Having said this it is important to be clear where the work is being done in these suggestions. The dispute between the cognitivist and the non-cognitivist here is not about the nature of a practice. The term is mostly being used to capture a key common thread on either side of the dispute. What is at issue is the status of justification and claims to truth that arise out of these practices.

them). But once these practices are established it is not our attitudes and responses that determine how we proceed or how it is proper to proceed within them.

You can say, if you like, that we create a form of life that invests certain features of people, acts and situations with the status of values. (Surely that is not false.) But this is not to say that *values* are created thereby. Rather values are discovered by those who live the form of life that is said to have been created. ('Moral Cognitivism', p.79)

McDowell, for his part, quotes Stanley Cavell:

That on the whole we do [go on in the same way] is a matter of our sharing routes of interest and feeling, modes of response, senses of humour and of significance and of fulfilment, of what is outrageous, of what is similar to what else, what a rebuke, what forgiveness, of when an utterance is an assertion, when an appeal, when an explanation — all the whirl of organism Wittgenstein calls "forms of life." Human speech and activity, sanity and community, rest upon nothing more, but nothing less than this.¹⁵⁷

There is a sense in which the disagreement here between the cognitivist and the non-cognitivist is not, as it were, a dispute about ethics, or, at least, one that is necessarily within or limited to ethics. Instead, it is a dispute about what is to count as 'the world'.

Darwall, Gibbard and Railton describe the cognitivist challenge in the following way:

At the same time, if non-cognitivism is to be distinct from cognitivism, it must insist that some judgements are cognitive in contrast with moral

¹⁵⁷ McDowell, 'Virtue and Reason', p.60. From Stanley Cavell, *Must We Mean What We Say?*. (New York: Charles Scribner's Sons, 1969) p.52.

judgements. Non-cognitivism is threatened not only by hardening values to make them like the hard facts of old; it is threatened by softening facts to make them more like emotive values.¹⁵⁸

What McDowell and Wiggins insist on (though McDowell, particularly, spends more time on this) is that what we take the world 'to be' is not sensibly thought of as being devoid of or completely independent of any kind of attitudinal investment on our part.

McDowell writes:

The idea of the world as motivationally inert is not an independent hard datum. ... If a conception of a set of circumstances can suffice on its own to explain an action, then the world view it exemplifies is certainly not the kind of thing that could be established by the natural sciences. But the notion of the world, or how things are, that is appropriate in this context is a metaphysical notion, not a scientific one: world views richer than that of science are not scientific, but not on that account unscientific.¹⁵⁹

The current point is exclusively concerned with the possibility of a "neutral conception of the facts"¹⁶⁰ or of the world as 'motivationally inert'. If there is no neutral conception of the facts and so, what we take to be a fact is invested with value in some way, then the cognitivist position we have been considering looks reasonable. That is, if cognitivists like McDowell and Wiggins can succeed in tying our concepts, stances, concerns, practices — the bundle of things that might be called 'forms of life' — to any attempt to

¹⁵⁸ Darwall, Gibbard and Railton, 'Toward *Fin de Siècle* Ethics', p.16.

¹⁵⁹ McDowell, 'Are Moral Requirements Hypothetical Imperatives?', *op.cit.*, p.83.

¹⁶⁰ *Ibid.*, p.81.

say ‘what the world is’, then the sharp distinction between facts and values might break down, as it were, from the fact side. It would be difficult, in such a situation, to resist Wiggins’ story about justification. However, if enough of the world is value neutral and we can give an account of it such that it is the intelligible object of our responses, then, pending such an account, the non-cognitivist would appear to have the upper hand.¹⁶¹ In this case, quasi-realism or something very much like it would be the order of the day.

In any case, we need not be too concerned about the issues surrounding this ‘over-and-above’ kind of justification since they lie outside the bounds of this project. They are *clearly* vital questions for settling the dispute between the cognitivist and the non-cognitivist but they are essentially questions about truth, or better, questions about the nature of the world ‘as it really is’. It will be recalled that the debate between cognitivist and non-cognitivist can be understood to take a detour through psychology. Rather than worrying, as it were, directly about the truth-aptness of moral judgements, we can investigate the kind of mental state that is ‘responsible’ for these judgements. The thought again is that belief-like or cognitive states, because they represent the world, will, at least in principle, lead us towards some notion of moral truth. Let us call the alignment of moral judgements with either belief-like or desire-like mental states stage one. Stage two will then attack the question of if and how this psychological account permits of moral truth in any way. These two stages very closely match those suggested by Wiggins:

¹⁶¹ Cf. Williams, Bernard. *Ethics and the Limits of Philosophy*. Cambridge, MA.: Harvard University Press, 1985: 129. “If human values are projected from human concerns, and are not a feature of ‘the world’, it does not follow that there is to hand a description of the world — or rather, of enough of the world — that is value free.” McDowell in ‘Aesthetic Value, Objectivity and the Fabric of the World’ challenges Williams on the point mentioned above: the distinction between the scientific and metaphysical conception of the world.

The first task of the cognitivist is simply to restore the cognitive aspiration of moral judgement. ... The moral cognitivist's second, more difficult task, ... is to consider how to make good on that aspiration.¹⁶²

Quite clearly, I think, the debate about justification, relativism and the proper understanding of 'the world' involves the second of these stages — the one that takes us from the suggested account of the psychology of moral judgement to the whether this psychology provides morality (and the content of moral judgements) with any kind of truth. As such it is beyond the scope of the current inquiry.

(ii) Psychology.

In addition to these large issues surrounding justification are contentious matters regarding the nature of psychological explanation or the proper account of the mind that is to be used in such explanations. McDowell spends some time attacking what he calls the eighteenth century philosophy of mind — that view of the mind as being comprised of two distinct faculties (reason and sentiment) and the corresponding ideas of belief and desire functioning with two distinct roles. It is not at all surprising that this is one of McDowell's targets. As we have seen, one of the key hurdles faced by any cognitivist account of moral judgement is motivation. That is, if moral judgements are taken to be the 'product' of cognitive or belief-like states, how can the agent come to have the motivation to act on that judgement? The cognitivist must either develop an account of

¹⁶² Wiggins, 'Moral Cognitivism', p.64. Notice however that there is a difference for the cognitivist and the non-cognitivist between the ways in which one might 'restore' the aspiration — the cognitivist seeks to restore the aspiration in such a way that the second task is an interesting one whereas the non-cognitivist needs explains the *appearance* of the aspiration.

moral judgement as belief-like and then, in some way, add the desire to do as one ought or develop an account in which the belief-like moral judgement brings with it enough motivation to act. It is the latter alternative that both McDowell and Wiggins take.¹⁶³ This course, however, seems to violate the typical understanding of belief-like and desire-like states (discussed at the end of section A of this chapter) which holds that belief-like states represent the world and desire-like states are responses to it. This 'typical understanding' is understood by McDowell to represent a prejudice encapsulated by the eighteenth century view of the mind.¹⁶⁴ Of course, the dispute about justification and the appropriate conception of 'the world' is very close at this point. If all (or enough) of our representations of 'the world' are infused with our values, stances and attitudes, then it may well be appropriate to understand some belief-like states as providing us with enough motivation to act.

In Wiggins the contrast between the two views of the mind is softened. Both Wiggins and McDowell, I think, have the same kind of picture in mind, as we have seen from their respective interpretations of Aristotle. However, McDowell is more concerned to play up his differences with the non-cognitivist and also, does not have the same links to Hume. Wiggins writes

¹⁶³ Given the project, it is also the route that we must take. Following W. D. Falk's original use of the terms, in this context the former is externalism and the latter, internalism (see "'Ought' and Motivation.' *Proceedings of the Aristotelian Society* 48 (1947-48): 492-510. Reprinted in Falk, W. D. *Ought, Reasons and Morality*. Ithaca: Cornell University Press, 1986: 21-41). This distinction goes to the heart of the issue about the 'pressure' in chapter one. In the end, the externalist must deny that the pressure, as characterized, accompanies all moral judgements of obligation.

¹⁶⁴ See McDowell, 'Non-Cognitivism and Rule-Following', pp.212-3.

What then are moral thoughts that motivate? They are states of finding that x deserves such and such a response. Or they are states of valuing x as having value v . It comes to the same thing. These states are comprised in larger motivational states. The full moral thought, if you like, is not just a belief. But it is not just a valuation either. The valuation itself can *be* the belief. ('Moral Cognitivism', p.83)

On this view, psychological explanations of moral judgements, in particular, do not resolve clearly into belief-like and desire-like states. This however is not necessarily to say that the distinction between the two kinds of states is useless, nor is it to say that the Wiggins/McDowell story involves the positing of a new kind of state, a "distinct existence". In this respect, the term 'besire' that has been used to characterize, particularly, McDowell's position is, at least, unflattering¹⁶⁵. Instead it represents a move to a holistic story. The claim is that the state of mind characteristic of the virtuous person is not reducible to, as it were, a relevant set of belief-like states and corresponding desire-like states. The criticism of eighteenth century philosophy of mind, then, is not that it limits states of mind to two 'functional directions' but that the division between the mental faculties (Reason and Sentiment) is too strict or discrete.¹⁶⁶

¹⁶⁵ Altham, J.E.J. 'The Legacy of Emotivism' in MacDonald, Graham and Wright, Crispin, eds. *Fact, Science and Morality: Essays on A. J. Ayer's Language Truth and Logic*. Oxford: Blackwell, 1986.

¹⁶⁶ Interestingly, this kind of position was proposed, in the eighteenth century, by Bishop Butler in his *Fifteen Sermons* ((1729) ed. Revd. D Matthews. London: Bell and Sons, 1953). Although Butler's aims were somewhat different, the problem with the account presented is effectively the same as the problem that the non-cognitivist sees with the Wiggins/McDowell account. Just as we might want to know more about the nature of conscience — how it works and what kinds of things affect it — so, the non-cognitivist demands an explanation, a further, perhaps reductive explanation, of the holistic state that relates 'the agent's conception of how things are in the world' with the agent's concerns.

The residual question, then, is about how any further resolution of the “full moral thought” would go and, indeed, if this kind of cognitivist is right to think that attempts to resolve the components of the full moral thought further are mistaken. Interestingly, this is closely connected to the claim, discussed earlier, that explanations give out. As we will see in the next chapter, the non-cognitivist thinks that more can be said and thus, that explanations do not give out at precisely the same point that the Wiggins/McDowell style cognitivist does. The non-cognitivist will (or should) agree that a complete account of the agent’s concerns cannot be given — i.e. that the thesis of uncodifiability holds. However the non-cognitivist cannot agree that we must leave the analysis of the relation between belief-like and desire-like states (or in Wiggins’ terms, from section A, a perceptual capacity and agent’s concerns) at the holistic level. This matter is importantly and tightly connected to the final point of contention between cognitivist and non-cognitivist that we will consider in the next chapter. If there are problems with the claimed inextricability of property and response involved in evaluative terms then we will have cause to suspect that psychological explanations do not quite give out where the cognitivist says they do. But of course, a failure to find a satisfactory way of giving further psychological explanation will suggest that after all these explanations do give out. We will be forced to return to the holistic position espoused by Wiggins and McDowell, at least on this score. For the purposes here and since we have not yet considered the non-cognitivist’s final objection and their version of the story, we cannot yet decide this matter.

CHAPTER 5: EXPLAINING THE EXPERIENCE (PART II)

Will you tell me, please, which way ought I to go from here? Alice asked.
 That depends a good deal on where you want to get to, said the Cat.
 I don't much care where —, said Alice.
 Then it doesn't matter which way you go, said the Cat.
 (Alice in Wonderland)

In this chapter we shift our focus from the cognitivist position to the non-cognitivist account. The shift will not be quite immediate because of the unresolved business from the last chapter. We will first tackle the third problem area that the non-cognitivist might find in the cognitivist account that was outlined in the last chapter. The investigation of this issue will lead us into the non-cognitivist account in section B. Section C of this chapter will turn back to the pressure to check how the non-cognitivist fares in explaining the phenomenon. We will find a way for the non-cognitivist but it will not necessarily prove to be a straightforward one. How this solution complicates the debate, its relation to cognitivism and the further development of the route itself will be considered in the remaining two chapters.

A. Shyness and Sensitivity.

One of the key moves in the Wiggins/McDowell cognitivist story is the claim that, as Wiggins puts it, the property and the response cannot be separated in an understanding/analysis of evaluative terms. This is the third point at which the non-cognitivist may attack. The non-cognitivist will claim here that not only does explanation *not* give out where the cognitivist claims it does but the explanation that can be given is a

non-cognitivist one — that is, including (the vital) desire-like state. Recall that part of the reason that the Wiggins/McDowell account could be said to a cognitivist one was related to explanations giving out at a certain point. It is partly because we cannot adequately resolve the virtuous person's state of mind into discrete belief-like and desire-like states that we end up saying that the virtuous person possesses a particular perceptual capacity. This is particularly evident in McDowell's paper 'Virtue and Reason':¹⁶⁷

The most natural way to press the [non-cognitivist's] objection is to insist on purifying the content of what is genuinely known down to something which is, in itself, motivationally inert (namely ... that one's friend is in trouble and open to being comforted); and then to represent the "perception" of a salience as an amalgam of purified awareness with an additional appetitive state. But what appetitive state? Concern for one's friends yields only the core explanation, not the explanation in which the "perception" of the salience was to figure. Perhaps the conception of how to live? That is certainly an orectic state. But given the thesis of uncodifiability, it is not intelligible independently of just such appreciation of particular situations as is involved in the present "perception" of a salience; so it is not suitable to serve as an element into which, together with some genuine awareness, the "perception" could be regarded as analysable. (p.70)

¹⁶⁷ John McDowell, 'Virtue and Reason' *The Monist* 62 (1979): 331-50. Reprinted in J. McDowell, *Mind, Value and Reality*. Cambridge, MA: Harvard University Press, 1998: 50-73. Page numbers are from the reprinted version.

This is an important passage and it is one to which we will return in the next chapter. Here however, the important point is that the unanalysable blend of belief and desire-like states is a vital step towards the final cognitivist account.

So why think that the cognitivist is mistaken about explanations giving out?¹⁶⁸ Simon Blackburn discusses a number of examples that look to be directed at precisely this point.¹⁶⁹ The first and, I think, most useful example is taken from McDowell's paper 'Are Moral Requirements Hypothetical Imperatives?'¹⁷⁰. McDowell writes:

In urging behaviour one takes to be morally required, one finds oneself saying things like this: "You don't know what it means that someone is shy and sensitive." Conveying what a circumstance means, in this loaded sense, is getting someone to see it in the special way in which a virtuous person would see it. In the attempt to do so, one exploits contrivances similar to those one exploits in other areas where the task is to back up the injunction "See it like this": helpful juxtapositions of cases, descriptions with carefully chosen terms and carefully placed emphasis, and the like. ...

Failure to see what a circumstance means, in the loaded sense, is of course

¹⁶⁸ Notice that if we can give further explanation it will most likely favour the non-cognitivist. The components of the cognitivist account are, as we have seen, precisely the components utilised in the non-cognitivist account (see section B of this chapter) — belief-like and desire-like states. If it turns out that the "perception" *is* analysable then it will be analysable into precisely these two components. And the resulting account will, at least, look decidedly non-cognitivist.

¹⁶⁹ Blackburn discusses a number of other examples that I will not explicitly address here. I will have something to say about them below.

¹⁷⁰ John McDowell, 'Are Moral Requirements Hypothetical Imperatives?' *Proceedings of the Aristotelian Society* Supp. vol.52 (1978): 13-29. Reprinted in J. McDowell, *Mind, Value and Reality*. Cambridge, MA: Harvard University Press, 1998: 77-94. Again, page numbers are from the reprinted version.

compatible with competence, by all ordinary tests, with the language used to describe the circumstance; that brings out how loaded the notion of meaning involved in the protest is. (pp.85-6)

[I will take it that the story behind McDowell's claims about this case is sufficiently clear from the previous chapter (particularly chapter four, section C).] Blackburn responds in the following way:

What this comes to, then, is that boys at a boarding school may know all the ordinary tests that a newcomer is shy and sensitive: they can sort out the shy and sensitive as well as anyone. But if they go on to tease him just because of that, this shows that in some special strong sense they do not know what it means: they lack the full amalgam which makes up the unitary psychological state that is only improperly split into perception and attitude. ... A natural way of saying it is that the boys see what the newcomer is like, and that is why they enjoy teasing him. This is what their cruelty consists in. They do not react to what they know as virtuous people should. (*RP*, p.100)

The most important point here is that this is a case in which the explanation, indeed the "natural" explanation, requires that perception and attitude are split. The feature of the example that makes it a strong one for our purposes here is that both the bully and the 'virtuous' person notice/see the same thing — shyness and sensitivity — and, by all accounts, they seem to understand shyness and sensitivity in the same way: (say) as involving some sort of weakness and potential for victimization. The difference, then, lies in the bully's response/reaction/attitude to the perceived shyness and sensitivity. It seems

hard to maintain that the difference between the two lies with their respective understanding of shyness and sensitivity.

This example functions by finding a case where the perceptual element is held constant while the attitudinal element varies. Thus the claim that the two cannot be separated or that our explanations give out looks to be undermined. Michael Smith expresses the form of the objection in the following way:

It is always at least possible for agents who are in some particular belief-like state not to be in some particular desire-like state; that the two can always be pulled apart, at least modally.¹⁷¹

Smith's worry seems to be that there might be cases where we cannot find a plausible, real-life example of someone being prepared to see a situation in certain terms but not having the required attitude. As a result, he insists that the separation is at least possible. This version of the argument is, I think, too weak to do the required work. The trouble is that if the "pulling apart" is only possible then the objection does not really get to the heart of the Wiggins/McDowell position. Wiggins would have no trouble granting that the attitude and the response *could* be pulled apart. To think otherwise would be to forget Wiggins' insistence that there are no *a priori* foundations for his view.¹⁷² Even if we supposed that there was nothing else to think but that a particular person was shy and sensitive and ought to be treated in a certain way, it would not follow that the perception of shyness and sensitivity *could not* be separated from the 'usual' accompanying attitude. To think that such a separation is not possible is to miss the contingency and

¹⁷¹ Michael Smith, *The Moral Problem*. Oxford: Blackwell, 1994: 119.

¹⁷² This is echoed by McDowell in, among other places, 'Virtue and Reason' (*op.cit.*, p.57-65).

anthropocentricity associated with Wiggins' story about the 'evolution' of our evaluative concepts.¹⁷³

The advantage of the example of the school bullies and their recognition of shyness and sensitivity is that it is more than a possibility, it is a plain and everyday example of how it can be that someone who has learnt and been inculcated into the language, can appropriately use a concept without the corresponding 'normal' attitude. The bullies are, we might imagine, normal in almost every respect. They have learnt the concepts of shyness and sensitivity in just the same way that everyone else has. They may very well, for instance, react differently to shyness and sensitivity recognized in their mothers or sisters. The bullies, indeed, look as though they satisfy Wiggins' requirements about the grasp of an evaluative term. Recall:

If a property and an attitude are made for one another, it will be strange for one to use the term for the property if he is in no way party to the attitude and there is simply no chance of his finding that the item in question has the property. (p.199)

The bully may well be 'party to' the attitude and certainly does find that his victim is shy and sensitive, but the attitude and the identification, the recognition of shyness and sensitivity, serves the desire to antagonize. What is the sense in which the bully is party to the attitude? He need not *have* the attitude in this case, but it may have come up in the course of his learning the concept. Indeed, having the ('normal') attitude may have been necessary for him to understand shyness and sensitivity in the first place. After that, things can go awry: the bully might not have the 'normal' association between the

¹⁷³ See chapter four, sections C and D. Also, Wiggins 'A Sensible Subjectivism' and 'Truth, Invention and the Meaning of Life' in *Needs, Value, Truth (op.cit.)*.

property and the attitude and he might be wrong in this respect. But this does not undermine the explanation of what is going on for the bully and equivalently, what is going on for the virtuous person. The difference looks to be something about their respective attitudes.

There are some very important things going on in all of this and most of them are related to the distinction between justification and explanation. In order to get at these it is useful to proceed through the question of meaning and understanding that is raised by McDowell's concession about the loaded notion of meaning involved in the shy and sensitive case. Recall:

Failure to see what a circumstance means, in the loaded sense, is of course compatible with competence, by all ordinary tests, with the language used to describe the circumstance. ('Are Moral Requirements Hypothetical Imperatives?', p.86)

The bully counterexample gives us reason to think that the meaning of shy and sensitive is not to be simply understood as an undifferentiable amalgam of the descriptive and the evaluative (belief and desire-like states). This is consistent with the concession. So how, then, are we to understand the sense in which the bully does not understand what it means to be shy and sensitive?

I think that for the moment we should separate the meaning involved (grasped) in ordinary linguistic competence from a broader (and more vague) sense of meaning that involves appeals to the justificatory practices and the interrelation between evaluative concepts as well as the 'whole wealth of considerations' accompanying the

Wiggins/McDowell approach.¹⁷⁴ The lack of understanding, given this separation, is caught up in the way the/our/one's normative and evaluative practices/concepts hang together rather than, as it might appear initially, in the linguistic meaning of shyness and sensitivity. Indeed if the Wiggins/McDowell story about justification ('there being nothing else to think', etc.) is right then there is a sense in which someone can fail to understand what it means to be shy and sensitive — of course with the proviso that 'means' here refers to the sense in which the/our/one's evaluative practices/concepts are related. But this does not jeopardize the point of the objection. The objection, as I have characterized it, primarily addresses the cognitivist's claim about where explanations give out and avoids issues of justification.¹⁷⁵ The point of this is simply to illustrate that

¹⁷⁴ I say 'for the moment' because the relation between linguistic meaning and this broader meaning needs to be explored. But even if it were right to challenge the adequacy of the 'narrower' conception of linguistic meaning and its connection to the broader idea, here it is not the place to do so.

¹⁷⁵ One of the reasons for not discussing some of Blackburn's other examples is related to this distinction between the explanatory and the justificatory projects. It is relatively easy to miss the distinction in some of the other examples and, it seems, their primary role is directed at the cognitivist's justificatory project. Consider the example of 'the cute' (roughly the same considerations apply to the 'fashion' example (*RP.*, p.96-7).

Cuteness, our man says, elicits and justifies various affective reactions. It is hard to specify them except as perceptions of cuteness, but perceiving cuteness in women, when all goes right elicits and justifies reactions along the lines of admiration and arousal in men, or envy or admiration, from women. ... If the last word is that these people perceive cuteness and react to it with the appropriate cuteness reaction, whereas others do not, we have lost the analytic tools with which to recognize what is wrong with them. (*RP.*, p.101)

If Wiggins' story about justification (there being nothing else to think, etc.) is correct then this story would appear to tell against those who apply 'cute' to women. That is, these judgements of cuteness (indeed the practice as a whole) fail in their aspirations to truth — i.e. fail in their cognitive aspirations — because the considerations that can be adduced in support of the judgements and of the practice fail to achieve the level of abstraction involving the common point of view of one person to another. That is, they fail to be something about which there is nothing else to think.

this example primarily engages the cognitivist in the justificatory issues rather than, what I have been calling, the explanatory ones. The shy and sensitive example strikes me as a particularly good one for bringing out the current objection but this is not to say that the same points could not be made using the ‘cute’ example.

The distinctions involved here are important for both sides of the debate. It is a mistake to think that the cognitivist claims that the attitude and the ‘perception’ are impossible to pull apart. Similarly, it is a mistake for the cognitivist to ignore the consequences of the distinction. At the explanatory level we do seem to need to posit an attitude to explain the difference between the bully and the virtuous person. Here, it looks as though explanations do not give out. Notice that if *these* arguments (namely, that there must be more that can be said) against the cognitivist do not go through for some reason, there is still some justification for pursuing further explanation, namely, reasons of symmetry. It seems odd to claim that when things are going well (i.e. an agent is virtuous) then our psychological explanations give out, but when things are amiss (as per the bully) then we can do a significantly better job.¹⁷⁶ It is often thought that one way of learning about a well-functioning organism/machine is by understanding what is awry when it is not functioning properly. This seems to be a common the approach in neurological science — a significant part of our more recently acquired knowledge concerning the operation of the brain comes from examining it when is malfunctioning. If this is a reasonable method to use, it would seem to be applicable here. Part of the problem for the cognitivist though is that the distinction between being virtuous and

¹⁷⁶ Notice that here I have been liberal with ‘justificatory’ terms like ‘proper’, etc. This has the advantage of granting the cognitivist their justificatory strategies whilst still getting the conclusion.

being vicious is not exclusive in the way that functioning and malfunctioning might be taken to be. As a result we cannot learn about the virtuous simply by looking at the vicious.

B. Non-Cognitivism and Projection.

At this point we turn to the non-cognitivist's position and given the discussion to this point it is natural to examine Simon Blackburn's version of it.¹⁷⁷ Blackburn's account may be described using three labels: projectivist, quasi-realist and expressivist. These labels, themselves, are unimportant but I will use them to signify three different aspects of the theory. First, let us treat projectivism (or Humeanism) as a view about the psychology of (or the psychological story behind) moral judgements: namely that the evaluative properties that we seem to 'see' in things are the result of a projection of our attitudes onto those things. Second, let us take quasi-realism to be a thesis about the metaphysics of morality — the nature of 'moral reality' or the 'truth status' of moral judgements: namely that although moral judgements are not capable of being true or false, the practices and modes of discourse that have developed allow us to talk as though they are. Finally, let us understand expressivism to be a claim about the way in which moral judgements get their content: when we make a moral judgement, say, that *p*, we are 'expressing' or 'giving voice' to our attitudes/responses that *p*. These three kinds of

¹⁷⁷ Alan Gibbard (*Wise Choices, Apt Feelings: A Theory of Normative Judgement*. Cambridge, MA.: Harvard University Press, 1990.) is an exponent of a very similar view. Blackburn's version of non-cognitivism is appropriate for a number of reasons. First and most obviously, Blackburn engages directly with the positions developed by McDowell and Wiggins. Second, Gibbard's approach (in particular) tends to have a more linguistic orientation. As such he has more to say about the expressivist side of the position than the projectivist one.

positions are clearly interrelated in various ways. It is certainly possible to be a projectivist without being a quasi-realist, for instance, but it might not be possible to be an expressivist without being a projectivist.¹⁷⁸ Again, what I take to be important about each of them is their respective starting points or concentrations. Thus, the expressivist is primarily concerned with moral language, the projectivist, with the psychological underpinning of morality and the quasi-realist, with the truth and objectivity of moral judgements.

In chapter one, I distinguished various ways of approaching the nature of judgements of obligation. I suggested that we approach moral judgements of obligation through psychological explanation appropriately construed to fit the nature of the enterprise and the starting point. Thus, our primary interest in Blackburn's account lies with his projectivism since it is here that we see the distinctive, psychological account given of our moral judging. This is not quite all that is of interest however. The last two sections of chapter four discussed, to a limited extent, the potential of the Wiggins/McDowell position to achieve the kind of truth and objectivity for moral judgements that is characteristic of cognitivist positions. For the sake of completeness, it is useful to examine, briefly, the quasi-realist aspect of Blackburn's position. We will not, in what follows, address Blackburn's 'expressivism'.

1. The Picture of the Mind.

The first order of business in examining Blackburn's account is to be clear about the picture of the mind into which his projectivism fits. There are three components that

¹⁷⁸ Indeed these two terms occasionally seem to be used interchangeably. My main point here is to signify different ways of approaching moral judgements, here, through moral language or through psychology. Cf. Blackburn, *The Oxford Dictionary of Philosophy*. Oxford: Oxford University Press, 1996: 132, 305, 315.

we will briefly consider here: holism, interpretation and the distinctness of belief and desire.

First, holism is the view that an individual's 'mental life' operates as a field or network in which particular mental states are to be identified by their relations with other mental states. Thus,

different beliefs and desires (and perhaps other states, such as emotions, attitudes, wishes, fantasies, fears, and of course values) come together to issue in action. But the contribution of any one of them will vary according to what else is in the mix, and therefore resists definition in terms of behaviour.¹⁷⁹

To this is added a (connected) view about the way in which we attribute mental states to others and hence, how, psychologically, we understand or interpret the behaviour of others. The thought here, taken from Davidson and Dennett¹⁸⁰, is that in order to interpret another's action we combine what we know about their situation (their context and, perhaps, what they have said or done in the past) with what we take them to have done. We then attribute to them beliefs and desires according to what it *makes sense* for them to have thought and wanted given what they did. Of course we can use the same set of considerations to 'solve' for a different unknown. Thus, "we can say what a person *ought* to do or what it *makes sense* for them to do if they have such-and-such a belief, coupled with such-and-such other beliefs and desires." (p.53) The combining principle,

¹⁷⁹ Blackburn, Simon. *Ruling Passions*. Oxford: Clarendon Press, 1998: 52. Hereafter, unless otherwise noted, page references will be to this work.

¹⁸⁰ Among other places, see Davidson, Donald. *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press, 1984 and Dennett, Daniel. *The Intentional Stance*. Cambridge, Mass.: MIT Press, 1987.

the principle of charity or humanity, facilitates the interpretation by providing a normative standard — what it makes sense to do, believe, etc. or what ought to be done, believed, etc. This principle is captured by Blackburn in the *a priori* principle of interpretation:

It is analytic that creatures with beliefs, desires, and other states of mind, behave in ways that (best) make sense (and not in ways that make no sense), given those states of mind. (p.55)

It “tells us what it is to have beliefs and other states of mind” and is a constitutive rule of the practice of interpreting others. (p.55) It means, again, that when we interpret people and give psychological explanations of their behaviour we are involved in a normative practice.¹⁸¹

The third component of the story is the distinction between, as Blackburn puts it, states that ‘belong’ to Apollo and those of Dionysus (*RP*, p.88). Thus, representations, beliefs and knowledge (perception), fall into the Apollonian camp whilst, desires, attitudes and responses (impulses, stances, dispositions, sentiments) belong in the Dionysian camp.¹⁸² Primarily, this division parallels that suggested in the previous chapter between belief-like and desire-like states. There I referred to the distinction between the directions of fit of each of the two kinds of states that was originally derived

¹⁸¹ How does this relate to (non-normative) functionalism, that is, the idea that mental states are identified by the role that they play in a casual network (p.56)? The idea is that there is an isomorphism between the normative structure and the causal structure: “Fred’s desires and beliefs form a causal structure that mirrors the structures of what it makes sense for him to do. If they don’t we have to reinterpret him.” (p.57)

¹⁸² Again, I have left out emotions because there is some disagreement as to where they belong. I would claim that however we understand emotions, their analysis will be given in terms of belief and desire-like states.

from Anscombe's example of the two ways of understanding a shopping list.¹⁸³ The shopping list may represent what the shopper has already purchased — the contents of the shopping bag — or it may represent what the shopper is to purchase — what is to become the contents of the shopping bag. Thus, the distinction is between “how we understand the world and what we are motivated to do about it.” (p.82)¹⁸⁴

2. The Account.

Having briefly sketched this background material, we can turn to the account itself. Blackburn's essential claim is that because ethics is practical, we are to understand moral judgements as essentially involving those states on the side of Dionysus: moral judgements are not to be taken as representing the world but as the products of our responses and reactions to it. It is desire-like states, on this view, that will play the central role in our explanations of moral judgements of obligation. Thus, Blackburn suggests that

¹⁸³ G.E.M. Anscombe, *Intention*. Oxford: Blackwell, 1957. As before, see also Platts, Mark. *Ways of Meaning*. London: Routledge and Kegan Paul, 1979: 257, Smith, Michael. *The Moral Problem*. Oxford: Blackwell, 1984: 111ff. and Humberstone, I. L. 'Directions of Fit' *Mind* (1992): 59-83.

¹⁸⁴ There is a question that might be raised about the combination of Blackburn's picture of the mind and his position in ethics. The position requires that input and output are always separable. Normativity, value and evaluation are each located on the output side, they each involve 'what ought to be done about the world'. But the picture of the mind described above finds normativity everywhere in the domain of psychology. Given this, it looks as though how we represent the world and what it is to represent the world involves normativity too. Blackburn's response is that the presence of normativity in the ways in which we represent the world does not jeopardise the distinction between

how we understand the world and what we are prepared to do about it. The understanding is practical within the sphere of understanding; our evaluative and ethical life is practical *tout court*. ... In all spheres of thought we make judgements, and judging is subject to standards of correctness. To represent the world as one way or another is to stand ready to be corrected. But in so far as this breaks down a distinction between representation and evaluation, it does so on behalf of evaluation. (p.82)

we can usefully compare the ethical agent to a device whose function it is to take certain inputs and deliver certain outputs. The *input* to the system is a representation, for instance of an action, or a situation, or a character, as being of a certain type, as having certain properties. The *output*, we are saying, is a certain attitude, or a pressure on attitudes, or a favouring of policies, choices and actions. (p.5)

We are not however to understand this too simply. Blackburn issues two warnings. First, this does not mean that what goes on between input and output is simple or “automatic” though it may be in some cases. We may dwell on our situation, reflect on our options or reconsider how we originally understood our position. Of particular relevance to us:

Selecting certain features of a situation as the ethically salient ones is a process that we practice, and that will change with education and experience. ... All that the input/output terminology insists upon is that we recognize the distinctness of the starting-point in the features of a situation that we believe to be present, and the upshot of taking them as salient, which is the output of practical policy, attitude, or emotion. (p.5)

Second (and not unrelated), the input/output idea is not meant to imply, as it were, a “one-way street”. “Attitudes and emotions determine the features of things and people that we notice. They organize our experience, determining how we construe situations.” (p.6) Again, importantly, our experience, properly understood as a creature of Apollo, is inert and is, at least in part, organized and shaped by our attitudes.

So how are we to organize our responses? That is, some of our responses, by our own lights, are moral and some are not. Since we are interested in moral judgements, we

need to say something about how they might be distinguished. Blackburn suggests that we think of a staircase of practical and emotional ascent. At the lowest levels we find 'weaker' attitudes like whims and impulses. As we climb, our attitudes become stronger, say, likes, dislikes and preferences. Still further up we find prohibitions, requirements and injunctions.

At the top end of the emotional scale are cases of harm and evil where dissent is not tolerated. I think it is wrong to hurt children, and here there is nothing left of *de gustibus non disputandum*. (p.12)

Typically, the attitudes that are associated with moral judgements are those at the top end of the staircase. However,

practical life comes in many flavours and there is no one place on the staircase that identifies a precise point, before which we are not in the sphere of the ethical, and after which we are.¹⁸⁵

It is perhaps useful to mark the gradations of attitudes on the staircase using a single example. Thus, if I do not mind raspberries or strawberries but have to choose, I might opt for raspberries today. This is the lowest level. Alternatively, I might prefer raspberries to strawberries and so (regularly, always) choose raspberries over strawberries. Continuing to climb we might get to the following attitudes: I might like or approve of raspberries to such an extent that I encourage others to have raspberries over strawberries; I might like or approve of raspberries to such an extent that I think that others should have raspberries; It might be unfathomable to me that someone could not

¹⁸⁵ *RP*, p.14. Cf. Chapter one.

choose/like raspberries; I might not associate with people who do not choose/like raspberries; I might despise people who do not choose/like raspberries.

Of course it is a little odd to consider attitudes to raspberries at the higher end of the staircase, obviously because it would be a little odd for someone to care that much. However, were we to change the example to one that sits more naturally at the higher end — say, our attitudes to knowingly causing the death of an innocent bystander —, we are likely to find it (at least) odd (more likely objectionable) that someone could think it only a matter of whim or fickle preference. This, I take it, bears out the idea of the staircase — our attitudes to both the individual who is indifferent about causing the death of another and raspberry lover are at different levels on the staircase.

It is crucial to this staircase idea that it includes both the practical and emotional components. Climbing the staircase is not only measured in terms of how strongly we feel. Sometimes our preparedness to act is the appropriate measure. We can be quite ‘cool’ about censuring the racist or about taking steps to avoid his company.

An ethic may be shown in perfectly calm. I may not be angry at someone who steals my goods, nor even hold him in contempt, but nevertheless think he ought not to do it, and here the output may be expressed in terms of my preparedness to encourage restraints and boundaries within which people should be forced to act. (p.13)

We might think of our ‘strength of feeling’ and our preparedness to act as two ‘variables’ that are involved when we talk of our attitudes. It is the combination and adjustment of these two variables that enables us to ‘locate’ an attitude somewhere on the staircase. There is one final variable that is available to the non-cognitivist: the objects of

our attitudes. Just as we find variation of the strength of feeling of an attitude and our preparedness to act on it, we also find variation in the kinds of things that are the object of our attitudes. Here we can speak of both higher order attitudes and more or less general attitudes. This variable registers the fact that I may have attitudes to attitudes, attitudes to, as it were, sets of beliefs (as opposed to particular beliefs or representations) and also, attitudes to sets of attitudes. So, in the example above, I have an attitude to raspberries for me today, to raspberries for me generally, to raspberries for others today, and so on up to raspberries for anyone ever. I might also, of course, have attitudes to all summer fruits or, more generally still, to all summer foods. Similarly, I may have attitudes to attitudes. I may wish that I did not like raspberries or I might take my attitude towards raspberries to be one of my most cherished — my attitude to raspberries is tied to who I am.¹⁸⁶ It is important to note that neither an increase in generality nor in order of attitude necessarily implies a move up the staircase. These three variables, in virtue of being variables associated with attitudes, constitute the tools available to the projectivist to explain moral judgements of obligation.

We can see each of these variables in operation in Blackburn's account of values:

To hold a value is to have a relatively fixed attitude to some aspect of things, an attitude with which one identifies in the sense of being set to resist change, or set to feel pain when concerns are not met. That fixed attitude typically issues in many dispositions, at various places on the staircase ... When things are not out of joint, values align with motivations. (p.68)

¹⁸⁶ Cf. Blackburn, *RP.*, p.66 and H. Frankfurt, *The Importance of What we Care About*. Cambridge: Cambridge University Press, 1988.

Values, then, are typically higher up on the staircase but, again, they need not be associated with strength of feeling. Hence they are “relatively stable dispositions to conduct practical life and practical discussion in a particular way” (p.67). We might also suspect that values tend to be more general or involve higher order attitudes. In this way they are well placed to issue in dispositions at various places on the staircase.

This is not to suggest that values and motivations are always (perfectly) aligned nor that we are always aware of what we value and the extent to which we value it.

On occasion, clearly, a person can act, knowingly and intentionally, against her values: she may have desires that overcome her scruples, or just knowingly succumb to temptation. ... But if the values are really there, we will expect them to manifest themselves somewhere else: in regret or remorse, or in guilt or shame. (p.61)

Sometimes we are surprised by how much we care about something (or that we do). We might discover this particularly by finding ourselves more reluctant to allow things to go wrong or more disappointed (remorseful, ashamed, guilty) when they do.

Perhaps the most distinctive feature of Blackburn’s account of values is their relative stability. Since values are to be thought of as relatively stable dispositions there are limits on how fast and (perhaps) how often we can change what we value.¹⁸⁷ These limits are related to the kinds of things that we take to be evidence for the attribution of

¹⁸⁷ Some values may of course change without our doing anything.

value in the first instance, namely settled patterns of behaviour that reflect the settled attitudes in question.¹⁸⁸ Thus:

Someone who one day professes one set of concerns and priorities, but another day quite a different set, is not someone who has many values, albeit for a short time each, but someone who has no real values at all.

(p.67)

3. Two Concerns.

As we have seen, we can understand there to be two significant intuitions that incline people away from non-cognitivism: the phenomenology of moral judgements and their apparent truth. Any non-cognitivist position must have something to say about both.

(i) Phenomenology.

There has been a good deal of discussion about the phenomenology of judgements of value, perhaps most explicitly in connection with the first chapter of J. L. Mackie's *Ethics: Inventing Right and Wrong*.¹⁸⁹ However, since our concern is strictly with the phenomenology associated with judgements of obligation we will not consider these issues here. The concern for us is how the non-cognitivist can provide an account consistent with the experienced pressure accompanying judgements of obligation. The grounds for the distinction here between these two kinds of judgements goes back to chapter one. How the non-cognitivist can account for 'the pressure' will be the subject of the next section.

¹⁸⁸ Of course for the most part we would believe someone's professing to hold a value. Evidence of such settled behaviour might only be required in cases where we, as removed observers, are doing the attributing or in cases that involve change or doubt.

¹⁸⁹ J. L. Mackie, *Ethics: Inventing Right and Wrong*. London: Penguin Books, 1977.

(ii) Disagreement/Deliberation.

There have been various ways of spelling out the second objection to non-cognitivism, concerning the truth of moral judgements, but most, I think, trade on the nature of moral discourse in some way. When we disagree about moral matters, for instance, the nature of our disagreement and the way these disagreements go, is not the way one would expect them to go if we were disagreeing about (mere) attitudes.¹⁹⁰ A major challenge to non-cognitivist positions, along these lines has been to provide an account of the practice of moral argument, disagreement and moral discourse more generally.

Blackburn's quasi-realism (combined with his expressivism) is designed precisely to meet this challenge. He points out that what we need is "an instrument of serious, reflective evaluative practice, able to express concern for improvements, clashes, implications, and coherence of attitudes."¹⁹¹ The central idea of quasi-realism is that particularly for those attitudes higher up the staircase, those that we count as moral attitudes, we could "invent a predicate answering to the attitude, and treat commitments as if they were judgements, and then use all the natural devices for debating truth."¹⁹² I have already mentioned the parallels here between the quasi-realist strategies concerning justification and the practices appealed to in the Wiggins/McDowell position. It is here that the 'wealth of considerations' that Wiggins makes use of are similar to those available to the non-cognitivist. Of course the difference is the status that these

¹⁹⁰ Cf. Ross, W.D. *Foundations of Ethics*. Oxford: Clarendon Press, 1939: 38-41.

¹⁹¹ Simon Blackburn, *Spreading The Word*. Oxford: Clarendon Press, 1984: 195.

¹⁹² *loc.cit.*

argumentative/justificatory practices can achieve. Again, the cognitivist thinks that they give us right to speak of truth, the non-cognitivist does not. This matter, as I have already pointed out, is beyond the scope of the project here.

C. A Return to ‘The Pressure’.

We can understand what has occupied us in chapters two, three and four as establishing that attention, the agent’s ability to form certain ‘takes’, is of singular importance to the project. The strategy in what follows then, is to ‘return to the pressure’, to bring this phenomenon to bear on the problem of understanding the ability to attend. What is yet to be addressed in the projectivist story is the fact that the pressure is experienced as objective or as originating independently of the agent — the concern about the phenomenology associated with judgements of obligation. Can someone attracted to the projectivist story provide an adequate explanation of this vital, indeed, defining feature of the pressure accompanying moral judgements? One might be tempted to suggest that the claim that such a explanation cannot be given is and has been the, at least implicit, rallying cry for those of a cognitivist or deontologist ilk throughout the history of ethics (of course, along with claims regarding the truth of moral judgements).

The preceding discussions, particularly those of chapter four, section B — those surrounding the explanation of the agent’s ‘take’ — have complicated matters to a significant extent. In order to conceptualise the competing issues it is instructive to have a working, simple model — one that glosses the complexities with which we are ultimately concerned. Of course once we have made some progress at this simpler level, it will be

necessary to check that the same considerations apply at the more involved level: how the agent attends to the context confronting him.

The complex problem concerns how the agent comes to have the 'take' that he has. That is, should we understand the 'take-formation' to be a matter of what the agent knows, and his attention affected because of it, or should we understand this process to be affected by the attitudes that the agent has? In short: what explains why it is that some people attend to contexts in such a way that they experience the pressure whilst others do not (or do not consistently)?

The simplified version of the problem drops the specific reference to attention or to 'take-formation'. Thus the question is: what explains why some judge that they ought, say, to help the old lady and others do not?

Here then is the central problem in its simplified form: My moral judgements are to be understood as the result of a projection of my attitudes onto a representation of the context in which I find myself. Alternatively, I represent the world and (then) react to it. The product of my reaction is the judgement. But as we have observed, the experience accompanying the judgement, the pressure, is one in which it feels as though it is the situation which is drawing, compelling or pushing me to so judge. This kind of experience, that accompanying normative judgements and, in particular judgements of obligation, is contrasted with that accompanying non-normative judgements where the apparent source of the pressure is something, as it were, internal to me — *my* attitudes, concerns and desires. The difference lies in where it feels as though the pressure comes from. The key idea here is that it is my desires, attitudes and concerns, and only these, that are responsible for my acting or my being willing to act. However, on the projectivist

story the only source of any kind of pressure is something that is in fact internal to me. The trick is then to show how only 'pressures whose source is internal to me' can generate a pressure that is experienced as 'external' to me. This is not simply a problem for the projectivist but is a version of the problem we were left with after chapter two. (Recall how McDowell's story would solve this: the idea is that by fusing cognitive and conative elements — or allowing that some cognitive states can 'generate' pressure, the felt 'externality' of the pressure can be accounted for.)

Since we have already discussed Blackburn's staircase of practical and emotional ascent, it seems fitting to begin understanding the projectivist's options with this idea. Once again, the suggestion is that as we climb the staircase we find attitudes and desires that are stronger and to which we are more committed. Thus, by and large, we do not particularly mind whether we get the second bowl of ice-cream but we do mind when someone around us is overtly racist or, when confronted by a manifest injustice. This, however, only captures the 'emotional' component of the staircase model. The second component is the practical element — the preparedness to act on our attitudes. What we are willing to do and what we actually do can also reflect a move up the staircase. Again, we may not be prepared to intercept the last of the ice-cream as it goes to someone else but we may be prepared to censure the racist or to ban him from our company. The need for the two components stems from the fickle nature of our experience of our own attitudes and desires. We do not always *feel* most strongly or most passionately about that to which we are most committed. (This is particularly important, as we shall see, for Michael Smith's account.) On this model, we are in the realm of moral judgements when the judgements we make are based on (or, in some way, the product of) attitudes/desires

that are high enough up the staircase to warrant our calling them such. Of course there is no set level above which attitudes are moral attitudes. There will be much variation across individuals, cultures and the like. Our moral judgements are based on/products of our deeper concerns and attitudes (where ‘deeper’ is intended to incorporate both the practical and emotional components in the appropriate way).

A first (and perhaps most natural) attempt to explain the experience of the pressure in the simple case might be to suggest that there are two competing desires: one which is, effectively, not to help the old lady and the other, to help (using, of course, the now familiar example). Thus, the pressure is explained by my desire to help the old lady. But as we have already suggested, the pressure is not captured by any occurrent desire that I may have and of which I am aware. When I judge that I ought to help the old lady across the road, the accompanying pressure is one that is felt to be placed on me or demanded of me *by the old lady's predicament*. Nowhere in the pressure is there any ‘reference’ to my attitudes or desires. The characterisation of the experience of the pressure is such that in many cases it opposes what I desire — or better, it is felt as being distinct from what I desire. As we discussed in chapter one, there may be cases where one has a desire to help and experiences the pressure. In such cases the two are distinguishable: “Even if I didn’t want to help her I still should”. Thus, in the case before us we might imagine that not only do I have a strong desire to continue on my way, but I am actually repulsed by the lady. She is smelly, unkempt and looks to have quite an unpleasant demeanour. Yet I may still feel the pressure to help¹⁹³. No conscious first

¹⁹³ It is interesting to notice that doctors might face this kind of situation reasonably often: the patient is in many ways unpleasant and may also be non-compliant and yet the obligation to treat remains (and is felt).

order attitudes can explain the pressure because I have no such first order attitudes to help.¹⁹⁴

Perhaps then we might say that it is on some other level that I want to help her and it is by reference to this 'other level' that we can explain the pressure — although it does not get us far, this certainly seems to be another way of capturing what is meant by 'deeper', above. There look to be two sorts of 'other level' available: higher order desires and unconscious desires. It is to these that we now turn.

Will higher order desires (more specifically, second order desires¹⁹⁵) assist in explaining the experience of the pressure? That is, can my desiring to desire to *X* explain how I come to experience the particular pressure to *X* (and sometimes to judge that I ought to *X*)? It seems not. It is natural to think that if I had both the desire to help the old lady and the desire to continue on my way then the second order desire (my wanting to want to help) would kick in and put whatever weight it had behind the desire to help. The desire to help might triumph. But again, I have no conscious first order desire to help. Alternatively, the second order desire might be satisfied and so generate the first order desire which was its object. In this event I would then have the conscious first order desire to help. Thus, second order desires could only be of use if they generated or affected first order desires but since I have no conscious first order desire to help in this case and since no such desire is generated, there can be no appeal to a second order

¹⁹⁴ Notice that even though there are no obvious first order desires that will account for the pressure, the strength of the pressure and strength more generally might still be an issue. I will help only if the pressure felt on me to help is stronger than any other pressures (no matter what their apparent source).

¹⁹⁵ It is, I think, easiest to imagine second order desires playing this role. In any case, the considerations which follow apply equally well to desires of higher order.

desire. These remarks, of course, only apply to cases where the generated or affected first order desires are desires of which I am aware. We will return to consider the version of this case with unconscious first order desires below.

There is a more subtle use of second order desires that is available here. It may be suggested that the pressure originates in some way from the combination of an awareness of a second order desire and the absence of the relevant first order desire. In such a case, the second order desire, it might be claimed, has an 'indirect' object and would be sufficiently disassociated to feel as though it comes from without. The second order desire is indirect because its object is not helping the old lady but the desire to help. It is disassociated because it is not satisfied¹⁹⁶. These two factors might be taken to combine to yield the seemingly mysterious pressure accompanying judgements of obligation. However we should be careful here not to confuse the second order desire with a judgement of obligation which has the first order desire as its object. That is, my wanting to want to help the old lady should be kept distinct from my judging that I ought to want to help. This aside, it remains difficult to see how the experience of the combination of a simultaneously unfulfilled and alternatively directed desire could produce the experience of a positive pressure to act which feels as though it originates from the context itself.

The objection which is, I think, most telling against this suggestion is that if any of the second order desires in question were ever satisfied then the pressure would not be experienced. Since we do experience the pressure (just as we do make judgements of obligation) and continue to do so, there must be something special about those second order desires that are responsible, in the end, for judgements of obligation such that they

¹⁹⁶ Thus, unconscious first order desires are not available for this option.

are incapable of being satisfied. Let me explain. If the account is that the reason I feel the pressure is due to my having a second order desire to desire to help the old lady but no first order desire to do so, then it would seem that whenever I judge that I ought (morally) to *X*, it is always because I desire to desire to *X* but do not desire to *X*. But occasionally, one would think, desiring to desire to *X* would result in my feeling the desire to *X*. Should this ever occur I would then end up wanting to help the old lady and not experience the pressure to do so. On this account it could never happen that both my second order desire be satisfied and the resulting judgement remain a judgement of obligation complete with the accompanying experience of the pressure. Thus, in order to preserve the experience of the pressure, there must be something special about *these* second order desires such that they cannot be consciously satisfied. It is hard to see what the principled and non-circular difference between these second order desires and all others could be such that these are incapable of satisfaction. On the face of it, there is nothing about the desire to desire to help the old lady which makes it unable to be satisfied.

The basic thrust of these suggestions can be summed up in the following way. Any conscious desire that I have will feel like a desire *that I have* and so will feel as though it comes from me. Thus, it does not look as though any conscious desires will be able to explain the experience of the pressure. Recall (from chapter one) that the degree to which I can recognise that it is a desire of mine that is the source grounding the 'ought' is the degree to which the experience of the pressure and the stringency of the 'ought' can be dissolved. Thus, in examples of, say, cleanliness, if I realise that the only reason I judge that I ought to clean my room is, in the end, because I have a general desire to keep

things in order, then the pressure which accompanies the initial judgement of obligation dissipates and is replaced by whatever impetus accompanies the desire.

All of the above, however, presumes that the first order desires that I have are conscious ones and hence, desires that I 'feel'. However, it clearly may be the case that I have first order desires of which I am unaware. On this view, I do in fact have a first order desire to help the old lady but I do not believe that I do because the desire is not felt. Thus, we come to the second way of explaining what is meant by wanting to help the old lady at some 'other level' — unconsciously.

The most significant problem with this suggestion is that unconscious desires will not do the explanatory work that is required of them. Our task was to explain the origin of the experienced pressure to help the old lady. However, if my desire to help the old lady is an unconscious one then this will only explain why I do not feel as though I want to help. It will not help to explain why I experience the pressure *to help*. Alternatively, we may think of the experienced pressure as a 'positive' phenomenon and the lack of a desire to help as a 'negative' phenomenon. Positing an unconscious desire will only explain the negative phenomenon — why I do not feel as though I want to help. It does nothing to explain the positive phenomenon.

A further word on second order desires is required here. One alternative which we passed over in the above discussion of the second order desire option was the possibility that I have a second order desire — the desire to desire to help — which is unconsciously satisfied. This falls prey to the same worry as the positing of unconscious desires generally. Not being aware that my second order desire is satisfied explains why I am not aware of my desire to help but does not explain why I feel the 'positive' pressure to help.

The desire of which I am aware, the second order desire, though now satisfied will still feel as though it is one of my desires and so, will not have the requisite 'external feel'.

To this point we have not been specific about accounts of 'desires' that might be offered. It should be pointed out that, insofar as the current project is concerned, the particular account of desiring that is to be preferred does not matter. Instead, we are considering whether whatever account is suggested can explain the phenomenon of the experienced pressure. There seem to be two possibilities. Either we have an account which in some way associates the phenomenology of desiring with what it is to have a desire ('the phenomenological conception of desire') or we have an account which makes the having of a desire not necessarily related to the phenomenology of desiring. The discussion above has largely been concerned with the former of these alternatives. The latter alternative takes us in the direction of dispositional accounts of desire and, in particular, the account offered by Michael Smith. It also leads us to focus more exclusively on the practical component of Blackburn's staircase.

The dispositional account of desire, as Smith presents it in *The Moral Problem*¹⁹⁷ trades largely on the significant problems faced by a phenomenological conception of desire and the idea of directions of fit, mentioned earlier.

(1) Problems with the phenomenological account.

(Again, I present these considerations, not for the purpose of advocating the dispositional account but simply to illustrate the motivations and capabilities of the dispositional account.) Smith suggests that there are two kinds of problems that face a

¹⁹⁷ Smith, Michael. *The Moral Problem*. Oxford: Blackwell, 1994.

phenomenological conception of desire. The first is related to the epistemology of desire whilst the second concerns the propositional content of desires.

The problem in the former case turns on a lack of correlation between the phenomenology of desire and when it is sensible to attribute desires to both ourselves and to others.

Consider, for instance, what we should ordinarily think of as a long-term desire: say, a father's desire that his children do well. A father may actually feel the prick of this desire from time to time, in moments of reflection on their vulnerability, say. But such occasions are not the norm. Yet we certainly wouldn't ordinarily think that he loses this desire during those periods when he lacks such feelings. (*The Moral Problem*, p.109)

In addition to this kind of case there are also more mundane examples such as walking or writing where we act on the basis of a desire but without awareness of any special accompanying 'feel'. Moreover, sometimes we simply make mistakes about the desires that we have. Smith gives examples which illustrate both that an agent may believe that she has a desire even though all of the behavioural evidence would suggest the opposite and where an agent believes that she does not have a desire but again, where the behavioural evidence clearly points in the contrary direction.

The second problem that confronts the phenomenological conception of desire involves the propositional content that desires have. Smith suggests (p.107) that an account of the nature of desire needs to capture or allow explanation of both 'the phenomenological content' and the propositional content. The comparison here is between desires and sensations, like pain. Both have phenomenological content, on the

current suggestion, but only desires can be ascribed to agents in the form 'A desires that p'. Hence, the problem with accounts which identify the nature of desires too closely with their phenomenological content is that they are in no position to explain how it is that desires have the propositional contents that they do.

(2) Directions of Fit.

Capitalising on Anscombe's idea (via Mark Platts¹⁹⁸), Smith uses the metaphor of a mental state's direction of fit with the world to introduce his account of desires. Again, the basic idea is that beliefs are supposed to fit the way the world is whereas the world is supposed to be fitted to our desires. In order to give substance to this metaphor Smith suggests a conception of desires which is based on their functional role — the metaphor for the differing roles given by the differing directions of fit. Thus,

we should think of desiring to ϕ as having a certain set of dispositions, the disposition to ψ in conditions C, the disposition to χ in conditions C', and so on, where, in order for conditions C and C' to obtain the subject must have, *inter alia*, certain other desires and also certain means-ends beliefs, beliefs concerning ϕ -ing by ψ -ing, ϕ -ing by χ -ing, and so on. (p.113)

This gels with the direction of fit metaphor by explaining the idea of a desire as a state with which the world must fit in terms of the dispositions of the agent to act on the world.

Roughly, ... a belief that p tends to go out of existence in the presence of a perception with the content that not p, whereas a desire that p tends to endure, disposing the subject in that state to bring it about that p. (p.115)

¹⁹⁸ Platts, Mark. *Ways of Meaning*. London: Routledge and Kegan Paul, 1979.

The difference in the direction of fit of beliefs and desires that are attributed to an agent, is thus given in terms of the different kinds of counterfactual that are true of that agent.

This account of desires can, according to Smith, solve both of the problems that faced the phenomenological conception. First, the epistemology of desire required that agents were fallible about their desires — thinking that they had desires which they did not and thinking that they did not have desires that they in fact had. On this view, it can clearly be the case that an agent can be mistaken about which counterfactuals are true of her and which are not. Second, the propositional content of a desire on this view would be determined by its functional role.

However, of greatest significance here is the position that this account allows with respect to the phenomenology of desire. It does not prejudice the phenomenological issue at all. It allows that some desires have phenomenological content essentially whilst others do not. As Smith puts it,

desires have phenomenological content just to the extent that the having of certain feelings is one of the things that they are dispositions to produce. Some desires may be dispositions to have certain feelings under all conditions: these have phenomenological content essentially. Other desires, though they are dispositions to behave in certain ways, may not be dispositions to have certain feelings at all: these lack phenomenological content altogether. (p.114)

What is important about this is that Smith's account does not call into question the phenomenological observations which we made in the first chapter. Indeed, it allows room for them. My desires, when I feel them, feel like they are mine — they, as it were,

come from me. This is what was meant by 'internal'. We can represent how this all plays out by combining what was suggested in chapter one, regarding the general experience of desiring and Smith's dispositional account. We can thus isolate three broad kinds of desires:

- (i) Desires that are always felt are given as dispositions to act and to produce feelings of the 'internal' sort.
- (ii) Desires of which we are sometimes aware are given as dispositions to act and occasionally to produce feelings of the internal sort.
- (iii) Desires of which we are never aware are given only as dispositions to act.

Additionally, given what we have said regarding the phenomenology of judgements of obligation, together with the assumption that judgements of this kind are the result of desires, characterised as the dispositional account would have us do, we can add a fourth kind:

- (iv) Desires with the associated 'external' pressure are given as dispositions to act and to produce feelings of the 'external' sort.

That this is an unsatisfactory result for our project should be clear. Consider our answers to the following questions. Why do we help the old lady across the road? Because we desire to. That is, because we are or were disposed to do so. Why do we feel the accompanying pressure to help the old lady across the road? Because part of the disposition to help the old lady includes the disposition to produce the feeling of pressure. The task at hand is to explain what 'goes into' the judgement in such a way that the *accompanying* experience (and by assumption, the necessarily accompanying experience) can also be accounted for. Here, however, we take our desires to explain the judging but

at the cost of explaining the accompanying phenomena. We might go so far as to call the fourth kind of desire an ‘obligation desire’ — the kind of desire that I have when I judge that I ought to perform some act.

The form that this problem takes is very similar to the objection that was made to the ‘unconscious desire’ suggestion above. There, the unconscious desire explained why we often do what we take to be our obligations but could do nothing to explain the accompanying phenomena. Here again, the phenomena are not explained. They are simply built into the dispositions.

The consequence of all of this is that the explanatory tools of the projectivist cannot be used in what we might take to be the usual way. Our initial attempt, on behalf of the projectivist was to endeavour to explain the experienced pressure on a simplified model. The problem had been complicated to a significant extent and it seemed reasonable to develop a model in which the subtleties involving the agent’s capacity to attend were glossed. We have now seen that the projectivist cannot provide the required explanation at this simpler level and so we must return to the complex version of the problem to see how the projectivist fares. Thus, in order to explain the experience of the pressure in projectivist terms, we must look to an account of how it is that I come to recognise what I do about the old lady’s predicament. We are back to attention and to the way in which “attitudes and emotions determine the features of things and people that we notice”.

In order to signal this return to attention I suggest that we introduce a metaphor that will capture the current issue. The metaphor which is, I think, useful at this point is

the idea of a 'filter'. Now, nothing of course, hangs on the metaphor that we use since, as we have mentioned, what matters is how such metaphors are cashed out. Thus, according to this picture and the projectivist model, our attitudes, as it were, sift through what is given to us, playing down some features and emphasising others. Blackburn suggests this very kind of thing when he discusses the importance of splitting 'input' and 'output': "Refusing to 'split' begins to sound like a refusal to think, perhaps symptomatic of a complacent belief that the emotional and moral lenses through which we see deserve no critical attention themselves." (*RP*, p.7) It is the emotional and moral lenses that are of interest. Clearly, this does not jeopardize the input/output model but it does signal a variation on the projection idea.

Thus, on this model the input is given by the context in which one finds oneself and the filter affects our attention to this context. The product of the filter and the input, the attention and the context, will be the agent's 'take' on the context — that is, how the agent 'sees' things, how the agent represents the world which confronts him. The output will remain, on the projectivist model, the judgement or the action characterised as "a certain attitude, or a pressure on attitudes, or a favouring of policies, choices and actions" (*RP*, p.5) Of course, other attitudes and beliefs can impinge on the agent's 'take' between the filter and the output. One way of construing the failure of the projectivist's story to explain the pressure on the simple model is in terms of a failure of this kind of 'secondary' attitude to explain the pressure. This new filter model is highly suggestive and, I think, contains the beginnings of a solution to our current problem. Blackburn, instantiating this picture, points out that

Loving or hating someone we highlight, perhaps unconsciously, features that make them lovable or hateful, sometimes even inventing ones for the purpose and suppressing what does not fit. In the light of emotion things which we would otherwise see become invisible, while others thrust themselves onto our attention. (*RP*, p.6)

The question that remains here is whether this kind of suggestion will carry over to the judgement of obligation case and assist in the explanation of the accompanying pressure.

It certainly seems to work for the bully in the 'shy and sensitive' example. However, in order for this example to be of most use here and in what follows, we need to fill in the story somewhat. Since we are concerned with judgements of obligation, we need, on the side of the virtuous person, to include such a judgement. Thus, let us imagine that the bully's intended victim has been placed in a potentially compromising situation. Staying with the playground example, suppose that the 'victim' has been instructed to perform some potentially demeaning, public task, say to pick up litter in the playground during the lunch recess. For the bully, this is a chance not only to put the victim further on display, but perhaps to make the task more difficult than it would otherwise be. Alternatively, for the virtuous person, the obligation, given the shyness and sensitivity of the 'victim', is to assist in the performance of the task, thereby (hopefully) reducing its unpleasantness and duration. This filled out version of the example does, I believe, preserve the relevance of both McDowell's and Blackburn's remarks whilst including, on one side, the judgement of obligation which is our target.

Returning to the point at hand, the bully likes to victimize those who are shy and sensitive and, on the 'filter' model, this explains why he is particularly good at noticing

shyness and sensitivity. Features that might suggest that someone is shy and sensitive are magnified by the bully's filter. The bully wants to antagonize (he has the desire) so, we can say, he notices the shyness and sensitivity because of his desire. That is, what explains the bully's ability to notice characteristics indicative of shyness and sensitivity is his desire to antagonize, the possession of which is part of what makes him a bully. So, extending this a little and using the language of the dispositional account, we might say that the bully's disposition to antagonize leads him (in some way) to be disposed to notice/recognize those features of contexts which will enable him to antagonize. It is because he always wants to antagonize that he is always alert, he is always looking or he always wants to notice certain kinds of features.¹⁹⁹ At this point in the proceedings I would like to leave this kind of desire — the desire to notice, etc. — unanalysed. Much more will be said on the matter both below and in the next chapter.

When we turn to the virtuous person making the moral judgement things become a little more tricky. The virtuous person, unlike the bully, does not have any conscious desire to help. This is a consequence of what we have observed about the experienced pressure and the judgement of obligation that the virtuous person makes.²⁰⁰ Thus, it is not the virtuous person's conscious desire that enables the attention as it did with the bully.

The natural move here is, again, to suggest that the virtuous person has an unconscious desire to help. On the face of it the move to unconscious desires here does

¹⁹⁹ We might take this to refer only to the 'ideal bully'. Real bullies get tired and distracted.

²⁰⁰ It is, as we have seen, possible that the virtuous person does have a conscious desire to lend assistance but this would be a case either where interest and obligation coincide or where the degree to which the virtuous person has the desire is the degree to which the obligation is not felt (i.e. the pressure is not experienced). Recall from chapter one: "Even if I did not want to lend assistance, I would still be obliged to do so."

not raise the same issues that it did in the previous, simpler model. The key difference is the role that these desires are supposed to play. In the simpler model, the unconscious desire was intended to explain the (potential) motivation to act and, along the way, the judgement. The problem was that the unconscious desire explained why there was no *felt* desire to help but did not assist in explaining why the pressure was experienced. On the 'filter' analogy the desire in question (the 'general desire to help') is not required to play this role. The filter is intended to sift the elements of the context, presenting the agent with certain features rather than others as salient. The unconscious desire here, would enable the 'colouring', the particular attentiveness, which determines the agent's 'take'. Thus the posited unconscious desire enables a similar kind of 'desire to notice' as did the bully's conscious desire to antagonize. The virtuous person, like the bully, notices those who are shy and sensitive. The burden to explain the attentiveness and, specifically the pressure does not lie with the unconscious desire. Instead it rests with the desire to notice that is enabled by that unconscious desire.

As a result of all of this we can develop the following picture. The bully, as we have seen, has an all too conscious desire to antagonize at every opportunity. Because of this and on the filter analogy, the bully is particularly alert to contexts which provide opportunities to bully. Cashing out the metaphor, the general desire to antagonize enables a desire to notice those who are potential victims. Thus, as in the case of loving and hating²⁰¹, the operation and satisfaction of the desire to notice enables some features of contexts to be 'played up' or enhanced, bringing them to prominence for the bully. The

²⁰¹ For a lover, the desires and concerns for the beloved lead to a heightened awareness — an alertness — of features associated with the beloved. Notice that this is not just at the descriptive level, i.e. not just at the level of highlighting the attractive features of the beloved.

bully's take on a context thus includes features indicative of shyness and sensitivity. The virtuous person, on the other hand, does not have the conscious desire to help but still notices those who are shy and sensitive. Thus, we can suggest that the virtuous person has an unconscious desire to help which plays the role of enabling the desire to notice that issues, for the case before us, in the recognition of shyness and sensitivity. Once again, the desire to notice plays the role of the filter.²⁰²

It is important to notice how these suggestions have the potential to give us the kind of account of the pressure that we sought. Crucially, the lens through which we see is not part of what we see. The colourings, weightings and saliences that make up that to which we attend, our 'take', are not registered as peculiarly ours — they feel as though they are a part of the 'world'. Thus, some features of the context are found to be particularly compelling. They are emphasized and are 'thrust onto our attention'. At this preliminary stage, it looks as though these features might have what it takes to explain the pressure.

Some initial remarks are warranted regarding, what I have called, 'the desire to notice', especially given the role that it plays in all this. First, there is nothing preventing this desire from being an unconscious one — that is, from being one which is not felt or at least not felt in a particular case. Indeed, if an agent were aware of such a desire not only would the pressure not be experienced, but the 'take' that the agent forms would have a distinctively 'subjective' feel.²⁰³ The agent might then think that he was seeing the

²⁰² If the virtuous person did have the conscious desire to help then the ensuing judgment would not be one of obligation — there would be a felt desire to help. This would then be a case just like that of the bully.

²⁰³ We might also add that were the desire to notice conscious, it would no longer be playing the role of the filter.

situation as he wanted to rather than ‘as it is’. The idea is that, by playing the role of the filter, the ‘desire to notice’ is already in place and operates, as it were, to select from the range of features that make up the context. Thus, we might be inclined to label it a *standing* desire to notice or to see the world in a particular way.

The second issue to address here is the relation between this ‘desire to notice’ and the agent’s ability to notice. As we saw in the last chapter, the cognitivist, following McDowell’s lead, suggested that it was the virtuous person’s ability that separated him from the non-virtuous. However, it seemed as though this ability could be characterized in terms of both a cognitive capacity and an extra non-cognitive state. This, with the appropriate qualifications, was the lesson of the ‘shy and sensitive’ case. The bully was able to recognise shyness and sensitivity just as well as the virtuous person but took such opportunities to antagonize rather than to help. This is why we looked to attitudes, desires and dispositions in the first place — to explain this difference. The problem is that being able to notice, say, shyness and sensitivity is clearly not the same as having the desire to notice shyness and sensitivity. Alternatively, it just does not follow that if I have the desire to notice then I will notice. But the virtuous person *does* notice and this is the special characteristic that the virtuous person possesses. It looks as though the desire to notice is not enough.

Returning to the analogy with the duck-rabbit will help. One of the favourable aspects of the analogy was the parallel between the amount of knowledge that people, in general, could be assumed to have. In the duck-rabbit case, the sense of the question ‘why is the tangle of lines seen as a duck first or in preference to a rabbit?’ depends on the assumption that the individual concerned is able to see it as both. In the ‘moral’ case

there was a similar kind of background assumption. We noticed that it would be quite rare to find someone who did not possess the requisite conceptual and perceptual apparatus that would allow them to see that the old lady was struggling with her groceries²⁰⁴. Being able to recognise that someone is struggling or that someone is shy and sensitive seems to be the same as a “competence by all ordinary tests, with the language used to describe the circumstances”²⁰⁵. Thus, it is reasonable to assume that people generally²⁰⁶ are able to perform these cognitive tasks. The explanation for why someone does not utilise these abilities or utilises them in differing ways, it seems, lies in some way, with their desires (or with desire-like states).

These suggestions, I think, provide us with the beginnings of a solution to our problem. They are, or at least seem to be, in the spirit of the projectivist account. How far and to what extent they remain consistent with this ‘spirit’ remains to be seen. It will turn on how we answer the two major questions with which we are left: (1) What is involved in ‘the desire to notice’?; and (2), What is the relationship between the more general desire (eg. the desire to help) and the desire which it enables — the desire to notice? It is to these questions that we turn in the next chapter.

²⁰⁴ What was not pointed out in last section is the connection between these observations and blame and responsibility. If we did not think that an individual was able to notice that the old lady needed help then would not blame them for, at least, failing to notice but perhaps also, for failing to help.

²⁰⁵ McDowell, John. ‘Are Moral Requirements Hypothetical Imperatives?’, *op.cit.*, p.22. My use of this remark in this context points to the criticisms made by Blackburn of McDowell’s position (discussed earlier). Thus, the more pertinent reference might be to Blackburn, *RP*, p.100.

²⁰⁶ Perhaps, we might say those adults who are not relevantly physically or mentally handicapped or at least, those with some minimum level of linguistic competence.

CHAPTER 6: THE DESIRE TO NOTICE.

“Though a man hath the best eyes in the world, he
cannot see any way but that which he turns them.”
(Bishop Butler, Sermon X, §4.)

There are three tasks to be completed in these final three chapters. First and most immediately, we must address the two questions with which we ended chapter five. This will occupy the entirety of this chapter. In the next chapter we will return to the cognitivist position that we left in chapter four. There is more to be said about the kind of position advocated by McDowell and Wiggins and its relationship to the issues discussed in what follows. Having done this and so fleshed-out the conclusion of the argument of the thesis, it will remain to draw together the various strands of the argument and briefly to assess the consequences of the account presented. This will be the task of the final chapter.

More specifically, in this chapter section A will consider and develop the initial account of the ‘desire to notice’. In section B we will extend the idea behind the desire to notice so that the account can be applied in cases of deliberation about and reflection on, one’s context. This extension will provide us with a final, developed account of the desire to notice idea.

It is notable that, in the context of the thesis as a whole, the strategy shifts from this point on. So far we have been concerned with the ‘regressive’ style of quasi-transcendental argument characterised in chapter one. That is, each of the arguments has

sought to support a series of premises that lead to the concluding position. This position has been reached, albeit in an as yet undeveloped form. The following chapters contain some detailed discussion of a number of examples and consideration of several possible objections — all of this is designed to further develop and elucidate the position that is suggested by the arguments of the preceding chapters. Primarily, these final chapters give body to a picture of the agent's frame of mind — in the terms of chapter one, the (partial) configuration of the agent's mental states which combine to issue in the judgements in question. In this configuration the desire to notice and its extension plays a crucial role.

A. An Initial Formulation.

The overall aim of this section is to develop the initial formulation of the desire to notice into a more precise suggestion. The plan is first to recap the initial suggestion and to develop the relationship between the desire to notice and the pressure associated with moral judgements of obligation. Second, I will suggest two considerations designed to elaborate the central desire to notice idea. In part two we will consider the sense in which the desire to notice is to be counted amongst desire-like states whilst in part three our concern will be the object of the desire to notice — 'the noticing' or the ability to notice. In the final part of this section I present four examples designed to further explicate the operation of the desire to notice and its relation to both the ability to notice and the general desire.

A word of caution: in what follows I am primarily concerned with structural issues concerning the desire to notice and its relationship to cognitive abilities and other desires. I will not be too specific about the 'object' of the desire to notice — that is, with

how best to characterise ‘p’ in ‘the desire to notice that p’. This is a problem because it makes some of the examples discussed seem, at the very least, implausible.²⁰⁷ For instance it is unlikely (though not impossible) that individuals will have specific desires to notice old ladies in need of help. Indeed it is even unlikely that such desires are desires to notice where help is needed. We can see this by considering someone who *did* only desire to notice old ladies in need of help but had no desire to notice others in need of help. This is an odd scenario whether we think of this desire as conscious or unconscious. Similarly it is unlikely that a well-trained and practiced bully will have the specific desire to notice shyness and sensitivity. We can easily imagine such a bully having, instead, the more general desire to notice those features of situations that are conducive to bullying or to antagonizing. So we might say that the ‘expert’ bully has the desire to notice shyness, sensitivity, vulnerability, insecurity, timidity, hesitancy and so on. The suggestion is that the desire to notice operates primarily as a standing desire, playing the role of a filter. As such it is more plausibly thought of as having a more general, broader content. In the following discussion, given my interest in structure, I will ignore these ambiguities. We shall return to this matter in the next section when the relation between the desire to notice, more general desires and cognitive abilities is clearer.

Our target in the last two chapters has been how an agent comes to attend to a certain context, i.e. how ‘takes’ are formed. We left things in the last chapter with the beginnings of a story about both the bully and the virtuous person in the shy and sensitive case. Both the bully and the virtuous person are good at noticing situations that contain shyness and sensitivity. In the bully’s case this is relatively straightforward to explain —

²⁰⁷ My attention was drawn to this problem by Stefan Baumrin.

he wants to antagonize. This general desire to antagonize might be a standing desire (that is, it might not always be to the fore) depending on the bully's dedication to bullying. Importantly, this desire functions to explain the bully's behaviour but it also assists in explaining why the bully notices (and is good at noticing) shyness and sensitivity (as well as other features that are characteristic of opportunities to antagonize). The additional step required in explaining the bully's selective attentiveness is, what I called, the 'desire to notice'. We can begin to see the need for positing this extra explanatory component by comparing this bully to a 'rookie bully': one who wants to be a bully, and so wants to antagonize, but, because of inexperience, is erratic when it comes to recognizing opportunities to do so. We can imagine the rookie bully making mistakes like picking on someone who is not shy and sensitive. We shall return shortly to be more precise about what the rookie bully lacks.

With the introduction of this 'desire to notice', the non-cognitivist can tell a similar story about the virtuous person. The virtuous person has the same ability to notice shyness and sensitivity as the bully. Naturally the non-cognitivist will be inclined to suggest that the virtuous person has a general desire to help (initially given at the same explanatory level as the bully's desire to antagonize).²⁰⁸ So it is the virtuous person's general desire to help that enables the 'desire to notice' shyness and sensitivity. But, as we saw in the last chapter, the virtuous person's general desire cannot be conscious (because it does not match with the phenomenology). It can, however, be unconscious — operating in the background, generally 'informing' the virtuous person's 'filter'. This unconscious general desire, then, is the key difference between the bully and the virtuous

²⁰⁸ We shall consider the cognitivist's position on this in the first section of the next chapter.

person. But how again does it enable us to explain the judgement and the pressure (recall that it is only the virtuous person who makes the judgement of obligation)? The desire to notice is the central element involved in the explanation of the idea of the filter. The filter, in turn, is the metaphorical representation of an agent's finding particular features salient. Importantly what is given through the filter is not recognized as being so given — the filter is not a 'conscious filter' and hence the desire to notice is not a conscious desire. Because of this, the features that are 'played up', emphasized or salient are not recognised as being the product of the agent's particular sensibility. Neither the bully nor the virtuous person view the shyness and sensitivity as dependent on their own viewpoint, as a product of their respective 'filters'. In this way there is an important transparency of the features that are highlighted through the filter. The demandingness of the shyness and sensitivity for the virtuous person and so, the experienced pressure, comes from the salience of these features, on the filter analogy, by way of the desire to notice.

So how, precisely, does the desire to notice assist with the explanation of the pressure associated with moral judgements of obligation? The successful 'operation' of a desire to notice invests the noticed feature with a significance that enables that feature to stand out with respect to other noticed features (and, trivially, with respect to unnoticed features). This gives the agent's take a 'depth of field' that it might otherwise not have had. Alternatively, it gives one agent's take a different depth of field to that of an agent in the same context but with a different desire to notice. This explains why features are taken by the agent to be salient — they are those that stand out in the agent's survey of the context, the agent's take. If the desire to notice some feature, *X*, is strong enough, i.e. the noticed feature is sufficiently central, that feature will, as it were, present itself as

demanding consideration — demanding to be acted or judged upon. Consider, for instance, the old lady's need of help. If an agent's desire to notice instances of *X* (where *X* is a feature that is suitably related to a person's needing help) is strong enough and satisfied in a particular case, the old lady's need of help will be front and centre in the agent's take. This is how the old lady's need of help comes to have its 'to-be-dealt-with' quality — how it presents a demand on the agent, enables the felt pressure on action or judgement.²⁰⁹ Thus, if specific features, *Y* and *Z* are 'more salient' or more in the foreground of his take then those features will present a stronger requirement or demand on the agent. The salience or centrality of certain features in an agent's take is directly related to the degree to which an agent is likely to make a judgement of obligation (find an obligation) concerning those features. On the model being suggested, this centrality is a function of the strength of the posited desire to notice: the stronger the desire to notice *X*, the more central the relevant feature and the stronger the demand on the agent to respond to that feature.

Here we have a way of incorporating moral or practical necessity.²¹⁰ We can imagine a case where the desire to notice *X* is strong enough to override other features (or considerations) so that the old lady's need of help monopolizes the agent's take — the

²⁰⁹ Recall, in all of this, that we are primarily concerned with explaining 'what goes on' when the agent judges that he ought to do such and such and not with what may go wrong or prevent such a judgement.

²¹⁰ These are an important set of cases where the agent cannot conceive of acting otherwise — 'I *must* do such and such' — but where this limitation is a moral, or perhaps prudential, limit on what the agent can do rather than a physical one. Bernard Williams' paper 'Practical Necessity' (in his *Moral Luck*. Cambridge: Cambridge University Press, 1981: 124-31) is useful on this matter, as is Peter Winch's paper 'The Universalizability of Moral Judgements' (*The Monist* 49 (1965): 196-214. Reprinted in Winch, *Ethics and Action*. London: Routledge and Kegan Paul, 1972: 151-70.).

agent ‘sees’ only that help is required: ‘there is nothing else to do but help’. Thus, in the limit, an agent’s take on a context can give such weight to a particular set of features that all other alternatives are silenced. There is an important connection here with McDowell’s idea of some considerations being silenced rather than outweighed.²¹¹ We can think of an agent’s take on a context as analogous to a three dimensional image. Those features that are in the foreground are those that stand out for the agent, whilst those in the background do not. At one limit we have moral or practical necessity, whilst at the other, those features that are ‘far in the distance’ might be those that are not registered ‘consciously’ — ‘seen but not noticed’. On the analogy, features are positioned somewhere in the ‘foreground’ through the satisfaction and strength of some relevant desire to notice.²¹²

1. Two Initial Considerations.

We can get at the issues involved with the desire to notice in another way: Why would someone *fail* to notice, say, a particular case of shyness and sensitivity? We might say that either (1) they cannot or could not notice these features or (2) that, in some sense, they did not ‘want’ to. The latter alternative takes us into the realm of desire (desire-like states). We should, however, be careful about the sense in which this is a desire. In the last chapter we referred to this desire to notice as a standing desire — a readiness or preparedness. We will return to this below. In the former case, the ‘could’ refers to an ability to notice. Such an ability might be lacking in someone who did not possess the

²¹¹ ‘Are Moral Requirements Hypothetical Imperatives?’, p.90-1.

²¹² There are some exceptions here that are the result of frequency of exposure rather than any desire to notice. We will discuss these cases in section B.

concepts of shyness and sensitivity, i.e. someone who did not know what shyness and sensitivity are and how to identify them. We can understand this ability in terms of the counterfactual: If the agent had been *looking* then he would have noticed.²¹³ Importantly, ‘looking’ here is to be understood as ‘concentrating on’ or ‘actively searching for’ the features in question. [This is attending in the strong sense — see chapter four, section B]. It looks as though this requirement begs the question against the cognitivist, like McDowell, who claims that only the virtuous person, and not the bully, understands what it means to be shy and sensitive. However McDowell does distinguish between *fully* understanding ‘what it means’ and an understanding that is “compatible with ordinary linguistic competence”. Part of the lesson of the shy and sensitive case is that both the bully and the virtuous person share an ability — an ability that is involved in the possession of this ordinary linguistic competence. So whatever ‘fully understanding’ amounts to, it is more than the basic and straightforward cognitive ability that is shared by the bully and the virtuous person. The idea is that it is this ability that is tested or picked out by the counterfactual. It is important to notice that there must be some shared ability involved here. McDowell’s initial example involves ‘getting someone to see the shyness and sensitivity’²¹⁴. It seems right to say that we do this, but there would be little point in doing so if they simply could not — and this, I take it, means: could not *if they were looking*.

²¹³ Of course it might be a regular conditional: If the agent were to look then he would notice.

²¹⁴ Cf. McDowell, John. ‘Are Moral Requirements Hypothetical Imperatives?’ *Proceedings of the Aristotelian Society* Supp. Vol. 52. (1978): 13-29. Reprinted in McDowell, *Mind Value and Reality*. Cambridge, MA: Harvard University Press, 1998: 77-94. Page references will be to this version. The relevant passage here is at pp.85-6.

Conversely, we can ask why someone *would* (initially) notice a particular case of shyness and sensitivity. Here, on the assumption that the individual is *able* to notice shyness and sensitivity, there appear to be three options: (1) the agent is actively looking for the features in question; (2) the agent has a standing readiness or preparedness to notice the relevant features; or (3), the agent could not have helped but notice the features.²¹⁵ In (1) the agent specifically *intends* to find or to try to find the relevant feature, whereas in (3), even if the agent actively attempted not to notice, he would fail. Option (2), in a way, sits between the two: it registers the idea that the agent is in some way responsible, unlike (3), without requiring an active or conscious stance (as in (1)). Naturally, the desire to notice falls into this second alternative. On the first alternative, part of what distinguishes the active looking is the phenomenology. The agent might report (or the surrounding behavioural evidence might indicate) to observers that he was ‘searching for’, ‘trying to find’ or ‘trying to see’ the features that indicate shyness and sensitivity. Again this might be the case if, as in McDowell’s example, the agent has been directed or encouraged to see the context in a certain way — “See it this way”.

The third alternative is an important one. Is there anything that someone cannot help but notice? This seems unclear. We might be tempted to say that *something* must be noticed and we might go on to specify what kinds of features must be noticed by all with proper vision and basic conceptual skills. [We might also say that *something* must be noticed but there are no particular features that any agent cannot help but notice.] The important thing is that even if there is some feature(s) of a situation that an agent (or agents generally) must notice, such a feature will not be helpful here. This is tied to

²¹⁵ Again, there might be a fourth option here involving repetition or frequent exposure (See section B.).

responsibility and voluntariness. If the agent must notice something then there will be no interesting answer to why the agent noticed this feature rather than some other or why his attention was drawn to the feature in question rather than to another. But these questions, here and in ethics more generally, *are* interesting questions.²¹⁶ (Of course we might want to explain why *this* feature *must* be noticed.) Once we are trying to give psychological explanations in ethics (of the kind with which we are concerned), an agent's having options is crucial. With the exception of very special cases of moral or practical necessity, moral judgements involve being able to judge one way or another, being able to act one way or another or, in this case, being able to notice one thing or another.

2. The Desire to Notice as a Desire.

Let us be clearer about the sense of 'desire' that is involved in the desire to notice. The desire to notice is a desire-like state, where, crucially, this is taken to be an explanatory category. We are in the business of giving a psychological explanation of the way in which an agent attends or comes to form 'takes'. The tools that are available for the explanation are the general categories of belief-like and desire-like states. Thus, in accordance with the direction of fit idea, the desire to notice does not *represent* the world — it is not belief-like — it is a way the world could be, it is a "state with which the world must fit". Understanding the distinction between belief and desire to be a primarily functional one and a product of the kind of psychological explanation that we seek, we were forced to separate the ability to notice and the desire to notice. We reached this conclusion largely because the ability to recognize shyness and sensitivity looked to be a

²¹⁶ See the discussion in chapter four, section B.

fairly ordinary one, tied to linguistic and conceptual competence, and one that was not, on the face of it, something peculiar to the virtuous person. Thus, beliefs about the nature of shyness and sensitivity and the ability to recognize instances of it could not explain why both bullies and virtuous people noticed when and what they did. The only other alternative, given our explanatory tools, is a desire-like state of some kind. But, a simple appeal to attitudes towards shyness and sensitivity did not provide a satisfactory explanation of the noticing because, assuredly, the bully and the virtuous person differed in this respect and yet still noticed. [It is here that the problems associated with the explanation of the pressure arose.] The desire to notice then, common to the bully and the virtuous person, does the work of explaining how an ability to notice gets applied in particular cases.

Consequently, what is most important is not that the desire to notice is a 'desire' in the ordinary sense or that we classify it as unconscious, but that it belongs in the 'desire-like state' explanatory category. That it is unconscious only signifies the fact that it does not possess the phenomenology of regular desires, i.e. it is not felt. As we saw in the discussion of Michael Smith's account, many desires that we take people to have, operate in the background, primarily involved in the explanation of actions. Thus, in the simple case, my walking across the room is (partly) explained by a desire to do so even though I felt no such desire at the time. In our case, we are not explaining an action (in the ordinary sense) but a judging (or specifically here, a noticing). The desire to notice functions in the same way to explain why, say, the bully noticed the shyness and sensitivity (— together with the desire to antagonize and the ability to notice shyness and sensitivity). In our case then, the bully's desire to notice plays the explanatory role

without requiring any of the usual phenomenology of regular desires. [There is some phenomenology here: the ‘finding’ of shyness and sensitivity that is associated with these particular features being those that stand out or are salient in the context. In the bully’s case this salience is taken as an opportunity to satisfy his desire to antagonize.]

This is not to say that a desire to notice cannot be conscious (just as the desire to walk across the room might be conscious). If I want a tin of baked beans from the pantry, I scan the pantry shelves looking for it. Here the desire to notice, captured by the ‘looking’, is very much a conscious desire of mine.²¹⁷ We can also build in ‘levels’ of description here. If there are many tins in the pantry I might scan for red labelled tins first in order to most efficiently find the beans. In such a case my desire for baked beans (in conjunction with the knowledge that baked beans come in a particular kind of (red) tin) gives rise to the desire to notice the particular beans tin in the pantry which gives rise to the further desire to notice the red labelled tins. Each of the latter desires serves to help satisfy the previous one.

We should be careful not to assimilate the desire to notice too closely with the walking case. When we consider the bully or the virtuous person we tend to have in mind entrenched patterns of behaviour and explanations that cover, perhaps, a large number of similar kinds of contexts.²¹⁸ Here, I think, Smith’s example of a father’s desire that his

²¹⁷ There is an interesting but perhaps unsurprising relationship between whether I notice and my desire to notice being conscious. When I do not notice the tin of beans ‘straight away’, I have to look harder and then I feel the ‘prick’ of the desire to notice. This parallels Smith’s discussion of the directions of fit of belief and desire.

²¹⁸ Recall, again, the discussion of the duck-rabbit in chapter four. When the behaviour looked to part of a pattern the duck-rabbit case became more interesting and, more importantly, we expected an explanation of it.

children do well fits better.²¹⁹ A desire such as this underpins many of the father's actions as well as more specific desires and it allows us to give a certain kind of continuity to his behaviour. Similarly, the desire to notice, perhaps functioning more strictly as the filter metaphor would have it, often provides a base for an agent's way of construing the world. Of course, the filter or lens (as in Blackburn's initial example) is that through which we see the world. The metaphor leads us to think of it as always in place, highlighting and down-playing various features respectively. We might also think that this 'filter' is very difficult to change (alter, adjust) or, in many cases, even to notice. These ideas can be accommodated by the desire to notice and are captured in one of the initial descriptions of the desire to notice, given in the last chapter, as a standing readiness or preparedness to notice.²²⁰

Understanding the desire to notice in this way allows us to draw connections with Blackburn's discussion of values or, at least, with the account of how attitudes (that look extraordinarily like values) can come to be deeply embedded and difficult to change. [We will see the importance of these connections in the examples below.] This is not to suggest that the desire to notice is to be understood as a kind of value. Thinking of the desire to notice as a standing readiness does relate it more closely to value but this is because it is often a value that is informing/giving rise to the desire to notice. If the desire (general concern/value) that is 'responsible' for the desire to notice is one that operates in

²¹⁹ This analogy might give a slightly misleading impression also. We tend to think that of the father's desire as a completely natural one requiring little, if any, cultivation. The desire to notice, along with the general desire that it serves may well require a good deal of cultivation. Of course in some cases the father's desire may require cultivation as well.

²²⁰ This is not surprising given that part of my claim, at least to this point, is that the desire to notice is the natural (perhaps, only) way to cash out the filter metaphor.

the background or is hard to change (“relatively stable”), it will come as no surprise that the desire to notice shares the same features. Notice that ‘being hard to detect’ is a feature of the desire to notice that is not related to the features of values in this way. Our own filters are hard to detect precisely because we ‘look through them’. The desire to notice is hard to detect because it is unconscious — i.e. ‘hard to detect’ here means being able recognize it in operation.

Not surprisingly, the response involved in the desire to notice is not a normal input/output response, though this fact does not jeopardize the input/output model. This is related to the projection metaphor being replaced here by the filter metaphor. If it were the same kind of response it would require, somehow, the representation of something and then a response. [We saw the difficulties with this view in the last chapter.] The filter, and hence the desire to notice, involves how the agent comes to represent the context *as* one involving shyness and sensitivity. Importantly this desire operates by effecting the *selection* of features in the context rather than by *reacting* to them. But in order for it to be selective there must be options to be selected from as well as weightings of or preferences between those options. This is both what makes it a desire-like state and what makes it an attitude/desire to something cognitive.

A key advantage of the desire to notice is that its satisfaction allows general desires to be applied and satisfied in particular cases. As we have seen, it does this by allowing that features of contexts are highlighted or emphasized as potential ‘satisfiers’ of the general desire. It is, we might say, a desire about how our cognitive or conceptual apparatus is to function if we are to best and most often satisfy our more general desires — thus we enable a transition from the general to the particular: the satisfaction of the

desire to notice tells us 'this is a context in which a specific general desire can be satisfied'.

3. The Object of the Desire to Notice.

Aside from these points, part of what distinguishes the desire to notice from other desire-like states is its object. In the 'normal' case, 'p' in 'the desire that p' stands for a representation, a state of affairs 'of the world'. However, in the case of a desire to notice the object of the desire is not some fact about the 'world' per se, but concerns the workings of my cognitive, conceptual or representational capacities. That is, the object of the desire, the state of affairs with which the world is to conform is, say, my noticing cases of shyness and sensitivity. Of course there is a very clear sense in which my noticing a case of shyness and sensitivity is every bit a part of the world as any other object of a desire. Even if the mental does not reduce to the physical, my noticing still counts as a state of affairs of the world in some sense. What is important here is that the desire to notice, like a second order desire, has a psychological state of mine as its object (as opposed to non-psychological states and the psychological states of others).²²¹

At this point we need to distinguish more clearly the desire to notice from the ability to notice since the desire to notice is, as it were, a desire to use a certain ability. Thus, it should be emphasized that:

- (1) Being able to notice is not the same as having the desire to notice; and,
- (2) It does not follow that the possession of the desire to notice means that I will notice.

²²¹ This distinction is intended to match the distinction, made in the first chapter, between 'internal' and 'external'.

Two immediate points should be made about (2). First, the attribution of the desire to notice is not intended to be predictive, it is explanatory.²²² This is partly a reference back to the discussion of Aristotle in chapter four and partly a feature of our kind of psychological explanation.²²³ Second, as with normal desires, things can go wrong in such a way that a desire to notice is not fulfilled — other desires can get in the way.

In the last chapter we made an assumption about ordinary linguistic competence in order to ground the idea that there is some selection between conceptual or recognitional abilities. The ordinary person, for instance, can be said to possess the ‘shyness’ ability, where this is understood as part of (or coming with) an understanding of the concept of shyness. That is, someone who grasps the concept of shyness is able to recognize instances of shyness (whether or not this ability is taken to be constitutive of or necessary for grasping the concept). Clearly the ordinary person will also possess other such abilities (i.e. have a grasp of other concepts) and so we might sometimes expect an explanation of why it is one ability rather than another that is operational, both in particular cases and more generally. However the question before us is not simply about the operation of a particular ability but the privileging or weighting of one ability over another. Importantly, the satisfaction of the desire to notice enables both the noticing and

²²² In this regard though, it may go some way towards an explanation of an agent’s ‘character’ and so, it may be somewhat predictive (See section D, below).

²²³ Wiggins (in ‘Deliberation and Practical Reason’, *op.cit.*, p.234) puts the point in the following way:

Since the goals and concerns that an agent brings to a situation may be diverse and incommensurable, and may not themselves dictate any decision, they need not constitute the materials for some psychological theory ... to make a prediction of the action. ... what is needed here is not prediction, but the subject’s own decision processes, constantly redeployed on new situations or on new understandings of old ones.

the ‘weighting’ of the relevant features of the context. Although it might be automatic for someone who grasps the concept of shyness to be able to recognize instances of shyness, something more is involved in having this ability in place, as a standing desire or as a standardly important or primary ability amongst one’s other conceptual or recognitional abilities.

Here we are again touching on the considerations discussed in the analogy with the duck-rabbit example in chapter four. We tend to look for an explanation in those cases where there is a pattern of noticing — such as the bully and the virtuous person — and when this pattern is thought to involve things like value and character, over which the agent has some control. As it was in chapter one, habit is an issue here since it straddles the boundary of the voluntary and the involuntary. There habit was involved at the level of the judgement of obligation and the experienced pressure: Does someone who acts (rightly or in accord with their obligations) from habit experience the pressure? The answer was no, but then again nor do they judge there to be an obligation in each case. Here the situation is different. We are now concerned with habits of attending or, more specifically, habits associated with the noticing and weighting of certain features rather than others. Here, I think the idea of habituation is very useful. There is a sense in which particular acts done from habit might not be ‘fully’ intentional (or, say, consciously done) but I am certainly held responsible for doing them. If something that I am in the habit of doing becomes wrong, then I ought to break the habit. Interestingly, one might be given some leeway in this regard. It might take time to break a habit, to get out of the routine. This is consonant with the time it takes to get into the habit in the first place. The extent to which ‘I couldn’t help it’ because I am used to acting in a particular way is the extent

to which I am excused, as it were, this time. None of this jeopardizes the desire to notice. Instead it usefully puts habits and the desire to notice, when it operates as a standing readiness, on a par. The extent to which we are responsible for our acts done from habit is the extent to which we look for an explanation of these acts. Such an explanation will be a psychological one and, given our assumptions about psychological explanations, will involve reference to belief and desire-like states.

4. Four (Useful) Examples.

There are four ‘symmetrical’ examples that will help in understanding the desire to notice and its relationship to the ability to notice and the general desire. First, consider the rookie bully and the (merely) well-intentioned person. The rookie bully, as an enthusiastic beginner, wants to be a bully but is not very good at getting the contexts for bullying right. Perhaps he tries to bully people who are not shy and sensitive or misses opportunities to do so. He might, as it were, ‘over-bully’, prompting reprisals, or ‘under-bully’ and fail to antagonize his intended victim. In both of these cases the failure is a failure to read the situation in a way that enables proper bullying — i.e. a failure to attend to the relevant features of the context that will enable optimal satisfaction of his desire to antagonize. We might think of the (merely) ‘well-intentioned person’ as not simply good natured or ‘naturally sympathetic’ but the kind of person who just wants to help (say) but, more often than not, is mistaken about the timing of the help, the kind of help or the extent of help needed. The concern (desire) to help for such individuals is, as it were,

rudderless²²⁴ because they do not attend to their contexts in a way that effectively enables help to be given.

These examples have important ties to discussion of Aristotle in chapters three and four. For instance, Aristotle writes:

The liberal man ... will give to the right people, the right amounts, and at the right time, with all the other qualifications that accompany right giving.” (*NE* 1120a24-27)

Of note is that the way in which the rookie bully and the merely well-intentioned fail in their respective endeavours mirrors the way in which the virtuous person succeeds. Both of our characters have the ‘right’ concerns but are not very good at seeing (reading, understanding, appreciating) contexts in which they find themselves. In short, they are not very good at ‘getting the minor premise’. On the model we have been using ‘getting the minor premise’ involves both the ability to notice and the desire to notice. Thus, there are two candidate reasons for the failure: either (i) they simply lack the cultivated ability to isolate and highlight the relevant features or (ii) their general desire has not enabled/generated a corresponding desire to notice.

The point here only concerns the simplest kind of case. It is possible that an agent is simply distracted. Other ‘desires to notice’ (see below) and desires more generally can, of course, get in the way — this is the way it is for most people and most desires I suspect. However such explanations are more ‘one of a kind’ than systematic. I take it that the ‘rookie’ bully and the well-intentioned person have more systemic issues. The rookie bully and the merely well-intentioned might be systematically distracted but this, I

²²⁴ Cf. E. F. Carritt, (note in chapter three).

think, will call into question the strength of both the general desire (respectively, to antagonize and to help) as well as the desire to notice rather than the ability or the presence of the desire to notice.

The latter of the two reasons, (ii), is perhaps more unlikely given the connection between the desire to notice and the general desire. Ordinarily we would expect that if one has a certain desire then one is also concerned to notice situations in which that desire might be satisfied. Finding out about missed opportunities to help or to antagonize seems, ordinarily, to result in disappointment or frustration at not noticing. This disappointment or frustration amounts (in our terms) to the realisation that one has failed to satisfy the desire to notice when the opportunity arose rather than a lack of the desire to notice itself. Further, the conditions under which we would be inclined to attribute, say, a desire to help but *not* a desire to notice those contexts that allow for helping seem very specific and limited. What is required is some other desire or consideration that affects the desire to notice without affecting either the general desire or the ability to notice. In any case, it is the former reason for failure that is of most interest here. It seems reasonable to suggest that for the most part the ‘flaw’ in these two characters is the lack of a cultivated ability to isolate and highlight the relevant features.

When we turn to the ability that these two characters lack, it is natural to suggest that the deficiency is more complex than simply a sub-standard ability to notice shyness and sensitivity (i.e. an inadequate grasp of those particular concepts). We will address this complexity below.²²⁵ For the moment however, it is instructive to treat their lack as a

²²⁵ This complexity, I take it, involves the content of the desire to notice. Recall that it is implausible to think of a desire to notice being as specific as a desire to notice shyness and sensitivity. Here, however I am concerned with more structural issues and so, for the

simpler one, involving only shyness and sensitivity.²²⁶ Thus, while each has the desire to notice shyness and sensitivity neither fully grasps the concept. They may be able to pick out central or easy cases but as the contexts become more complex and subtle, their limited grasp leads to a failure to isolate genuine instances. [One can imagine the ‘rookie’ bully being instructed: “Although this boy has the superficial characteristics indicative of shyness and sensitivity, he is trained in martial arts and not to be trifled with.”²²⁷]

The second pair of examples is, what we might call, the ‘jaded’ bully and the reluctant saint.²²⁸ The jaded bully was once an enthusiastic and adept bully. He has now lost his urge to bully but he still recognises (notices) opportunities upon which he used to act. He may have outgrown the urge or simply not enjoy it anymore. Correspondingly, the reluctant saint no longer values or cares about doing what is right. He used to go to extraordinary lengths to do as he ought and was dedicated to the welfare of those in need. The reluctant saint differs from the jaded bully in an important respect: he, unlike the bully, actively wishes he were not a saint. In both of these cases it is the general desire/concern that is missing and yet the honed ability to notice is still operational.

With both of these examples we run into a problem (although it is more exacerbated in the case of the reluctant saint). The story has been that the general desire to antagonize gives rise to the desire to notice cases of shyness and sensitivity, so how

moment, the simplified version is to be preferred (cf. the beginning of this chapter).

²²⁶ This is consistent with the way in which both McDowell and Blackburn use the example.

²²⁷ This is a useful example of the kind of practical knowledge that has an important role to play in what is to follow.

²²⁸ The latter of these two examples was suggested to me by Brad Armour-Garb.

can the desire to notice persist once the general desire has gone? Indeed, what reason do we have for supposing that either the jaded bully or the reluctant saint still have the desire to notice at all? The reluctant saint, in particular, does not want to do ‘good’ or be good and, we might suspect, does not even want to notice opportunities to do so — and yet he does. The urgency of this problem should not be underestimated. If there is no reason to posit a desire to notice in these two cases then our reason for positing such a desire at all is undermined.

Blackburn’s discussion of values is pertinent here.²²⁹ Part of the relevance is that Blackburn considers a set of cases in which

It seems natural to interpret some agents as regretting having the values they do: a person of unswerving integrity, for instance, might at some level wish that he was a little more relaxed, if those who cut corners or bend the rules just a little seem to be having more fun or getting on better in the world. (*RP*, p.66)

Both the reluctant saint and the jaded bully, to varying degrees, regret having the values that they do. Thus our understanding of the reluctant saint example can piggy-back on the relevant features of Blackburn’s discussion of values. The most relevant passage here is:

Because we are talking about the way an agent is set, there is a speed limit on change of values. ...Over sufficient time, of course, a person’s concerns and interests and priorities can change, although it is also worth noticing that values one thinks one has suppressed can be surprisingly resilient:

²²⁹ There are also close parallels with Blackburn’s specific example of Satan in Milton’s *Paradise Lost*.

habits of valuing, like habits of deference or old superstitions, tend to come back to haunt us. It is difficult to shake off their grip ... (*RP*, p.67)

If a desire to notice and its satisfaction have become sufficiently entrenched and have been cultivated by similarly entrenched values, the ways in which and the speed with which those desires to notice might change is similarly limited. A standing desire to notice may well be as entrenched or as resilient as values. One of the signs that a value is resilient in this way can be through the operation of a desire to notice. Indeed what people notice and what, in this sense, they desire to notice is part of the evidence we use to ascertain their values. Particularly in a case like the reluctant saint, the desire to notice is on a par with values. Like values they can be relatively stable “ways of being set”. So just as values can be “surprisingly resilient”, ‘old’ desires to notice might also linger.

It is an interesting question whether we should think of the reluctant saint as continuing to make the judgements of obligation that he used to. Presumably whether he does or not depends on the strength of his residual desire to notice. That is, on the model being suggested, how compelling he finds shyness and sensitivity will depend on the strength of his desire to notice these features and so, the weighting that these features are given in his take on the context. If his desire to notice is still strong then the shyness and sensitivity will still present him with requirements to act, i.e. will present him with an obligation — perhaps he judges that he ought to do such and such and then remembers that he has disowned his former life or perhaps he judges that he must do such and such but loathes himself for so judging (as in Blackburn’s example).

For the ‘good’ bully and the virtuous person, then, the general desire and the ability to notice are connected via the desire to notice. Rookie bullies and the (merely)

well-intentioned provide us with examples where the ability is not developed but where the general desire and the desire to notice is strong. The jaded bully and the reluctant saint provide examples where the general desire has faded but the ability and the standing readiness, the desire to notice, are still in place (albeit, perhaps, on the way out).

There are two further remarks to be made here. First, as with ordinary desires, it makes sense to think of an agent as having multiple ‘desires to notice’²³⁰. This idea is, I think, implicit in the analogy between an agent’s take and a three dimensional image. At any given time we might think there is a ‘resultant’ filter, understood as the ‘sum’ of standing desires to notice. Thus, we might imagine a teenage boy who, while being ‘well brought up’, is fascinated by the opposite sex. Ordinarily he would notice the elderly lady who needs a seat but in the presence of an attractive woman his attention is drawn away and he does not notice what he ordinarily would. The boy’s mother might be disappointed to find this and wish, in our language, that her son’s desire to notice those who need help were stronger than his teenage fascinations.²³¹ [The ‘resultant filter’ need not only be thought of as one desire overcoming another. We might also imagine cases, say, where the desire to notice feature *X* and the desire to notice feature *Y* combine to yield particular attentiveness to contexts involving *X* and *Y*.]

²³⁰ This was pointed out by Stefan Baumrin.

²³¹ Perhaps the more usual story might be that the mother hopes that her son’s concern to help overrides his fascination with attractive women. I do not intend that talk of desires to notice replace or be preferable to the ordinary way of putting things. The positing of ‘desires to notice’ serves our explanatory purposes here. For the reasons discussed in the last chapter leaving the explanation at the level of these broad concerns will not adequately explain judgements of obligation (— such as the one that the boy would have made had his concerns been to his mother’s liking). As we shall see in the next two sections the role of general concerns is important and I do not wish to suggest that the view that underpins the mothers’ hopes is not a correct one.

Second, it also seems reasonable to expect that one might have desires not to notice (or desires to ignore). Perhaps the best example here is noticing racial features. We can easily imagine someone thinking that the colour of a person's skin is not something that matters one iota, indeed, skin colour is something that ought not to be noticed.²³² After some practice such a person might become good at not noticing skin colour (or learn to ignore). This is not to suggest that such an individual would never notice racially-related skin colour. Again, like ordinary desires, other desires might get in the way or override the standing desire not to notice skin colour. Being annoyed at someone may well resurrect negative racial stereotypes that link certain behaviour with racial groups and thus highlight skin colour — and this even for an individual who is otherwise good at ignoring such features. There are many nuances involved here. We might suggest that there are 'ways of noticing' skin colour, some of which might be taken to be innocuous and others not. Whether a particular case or set of cases is innocuous or not will depend on the rest of the agent's values and the place within these of the desire not to notice. Similarly we might suppose that there are times when, in the agent's view, it is right to notice a person's skin colour — these might be cases involving affirmative action and the like.

B. Extensions.

There is little doubt that the 'desire to notice' is the lynch pin of the account being developed here. However, it is equally clear that the story told to this point (about both

²³² This is clearly a judgement of obligation however it is one that, in the language of chapter two, is to be classed as a judgement of value. As such it does not cause any problems here since judgements of value are not the target of the discussion. Again, see chapter two, section D.

the shy and sensitive case and the desire to notice) is the simplest version of the story. We had a glimpse of this when considering the rookie bully and the (merely) well-intentioned person. There are three specific reasons for complicating the desire to notice story to this point. Each of them serves to focus the overall aims of the section. We begin by considering these.

First, as I cautioned at the beginning of the chapter, my initial concern (and the concern of section A) was with the structure of the desire to notice idea in relation to the ‘general desire’ and to the ability to notice. This concern meant that I ignored an ambiguity involved in the object of the desire to notice. This ambiguity manifested itself primarily in the implausibility of some of the examples of in which the desire to notice figured. By extending the desire to notice it is hoped that this ambiguity will be redressed.

Second, it is not entirely clear that we have been treating McDowell’s ‘shy and sensitive’ example in the way that it was intended. When we considered the shy and sensitive case we interpreted the sense of ‘see’ in McDowell’s example of getting “someone to see” in a very ‘thin’ way: someone could see shyness and sensitivity if they possessed a “competence, by all ordinary tests, with the language used to describe the circumstance”. McDowell, I think, is concerned with a different kind of failure to see, one that is tied more closely to justification. It is not that the person is (literally) unable or ‘struggling’ to see this case as a case of shyness and sensitivity, rather that they are unwilling to do so (or to do so using the ‘full’ sense of shyness and sensitivity). Thus, the issue is whether the elements of the situation add up to a warranted or justified attribution of shyness and sensitivity (again, in its ‘full’ sense). Putting the matter this way makes a

big difference but, of course, it takes us back to the questions about justification (chapter four, section D). This reading makes the exchange a matter of which view of the context one *ought* to adopt. But of course our question here is not whether ascriptions of shyness and sensitivity are justified but how we are to explain an agent's attention to some features rather than others.

This is not the end of the matter however. It does not bode well for McDowell's view if this is the end of the story about his example of shyness and sensitivity. The example is intended to illuminate the sense in which "a conception of how things are suffices on its own to show us the favourable light in which the action appeared".²³³ That is, it is meant to be an example of the way in which the normative can be built into a cognitive sensibility, so that the virtuous person can be said to see matters in the right way and that this, by itself, is motivating. This is why Blackburn's bully is a concern for McDowell's position. The bully provides us with an example where the conception of how things are does not lead to virtuous action (or better, to the judgement of obligation). It is perhaps more appropriate (and more in line with McDowell's view), then, to think that the example is only partly about justification. The other part involves the relation between attention and reflection or deliberation (in the sense of actually considering or reconsidering the situation). By (re)considering one's context, the initial 'take' might be adjusted in the light of other features present but not initially noticed or given weight. Thus, we can understand the call to see the situation differently as a call to rethink one's initial view ('take') — to perhaps adjust the weighting given to various features or to register other features not incorporated or emphasized in the initial view. Putting things

²³³ 'Are Moral Requirements Hypothetical Imperatives?', p.86.

this way enables us to connect with the justification issues. These would be involved in a case where the call to rethink or reconsider involves a call to see the context not simply in a *different* way, but in a different *particular* way, namely, what is claimed to be the *right* way. Now the demand is not one about a simple or initial noticing but about dwelling or reflecting on how one initially construed the context. This still causes a mismatch between McDowell's example and our interests but it is now one that is within the scope of the present inquiry.

The 'desire to notice' story to this point best fits how we might explain an agent's immediate take — their first glance, as it were. But reconsidering or reflecting, and the changes in an agent's take due to these, needs to be given a role. McDowell's story combines this reflection with a cognitivist story about justification into an example in which the agent was encouraged not simply to reconsider but to reconsider in a particular way — the right way. While we need not be concerned here, with the right way of seeing (and so, with the surrounding issues of justification), we do need to say something regarding the reconsidering of or reflection about one's contexts. This (probably correct) reading of the shy and sensitive example shows, I think that the issue is not simply the agent's *initial* understanding of ('take' on) the situation. We need to be sure that the idea behind the desire to notice can accommodate reconsidering and reflection. Moreover, the desire to notice idea needs to be able to accommodate the influences and considerations that can be brought to bear on such considerations.

The third reason that we must extend the desire to notice idea comes from what is involved in an agent's take on a context. In chapter three, the account of a 'take' involved the selection of a subset of features from a vast number (perhaps, an indefinite number).

Amongst the totality of features were ‘historical’ connections, interpretive features and probabilistic or inductive considerations. The agent’s take was not limited to the first, initial impression of the immediate context — the agent was certainly permitted to revise an initial take to formulate and reformulate a view of the context. What is missing then, from the discussion of the desire to notice are explicit links to the way in which an agent’s view of a context can involve features that are more complex than those apparently involved in shyness and sensitivity and to the way in which the agent’s ‘initial’ take can be revised by reflection or consideration. Given the breadth of possible components in an agent’s ‘take’, the possibility of revision of this ‘take’ through things like reflection and reconsideration, the idea of a desire to notice needs to be extended. This extension can give us a clearer view of the more general idea behind and informing the suggestion of the desire to notice. Importantly, it will involve an expansion, consistent with the form of the desire to notice, of the kinds of things that may be taken as the object of such a desire.

We can get an initial grasp of this extension of the ‘desire to notice’ idea by considering the following example. Suppose that Ed and Ned are about to move to New York from a small country town. They know relatively little about the city and about life there. They speak with people who have been and who know what it is like. They also begin to take note of news pieces, stories and films that concern New York. Suppose further that the vast majority of the information, testimony and stories that they hear focuses on the size, pace and, perhaps, the ‘brutality’ or harshness of life in the big city. Their picture of New York is a harsh and threatening one and their impression of New Yorkers is that they too are harsh and threatening — largely because they have become

accustomed to living in such a city. [It is no doubt possible that all of the information that Ed and Ned gather about New York is true, although their respective impressions and overall picture of the city prior to their arrival may be different. It is, I think, equally possible that their respective impressions and overall picture be largely the same. How they 'deal' with it might be the major difference.] Suppose further that Ed and Ned differ in their general attitudes to new experiences or, perhaps more specifically, in their attitudes to the paradigmatic 'big city'. Ed prefers the quiet and the greenery of the country or the space of the suburbs. He is more generally conservative in his tastes and his lifestyle. Ned is adventurous and lives a little more on the edge. Perhaps he is a frustrated bohemian or sees himself as the high-finance type. There are myriad ways of characterising these differences and indeed the differences themselves can vary in indefinitely many ways. What is important is that Ed and Ned have different kinds of antecedent responses and attitudes that are pertinent here.

Now, it is likely (perhaps more than likely) that Ed's experience of New York, at least initially, will bear out the picture that he had developed before coming to the city.²³⁴ He will notice homeless people sleeping on the streets, the noise and the dirt. He will also notice the harshness and rudeness of the people. Ned, on the other hand, might have a different experience. Instead of encountering reckless cab drivers and hundreds of faceless commuters he finds chatty, rambunctious cabbies and a melting pot of people and cultures leading their lives in a vibrant metropolis. This is not simply a point about descriptions; it is how the people and the city appear to ('strikes') the two newcomers.

²³⁴ It is important that the considerations that follow for both Ed and Ned concern only the initial part of their stay. Over time one might expect that their direct and personal experiences of the city will outweigh the information with which they began.

Only part of this difference can be explained by differences in their actual experiences. Ned may well have happened upon as many homeless people, as much dirt and noise and as much rudeness as Ed but for Ned, both in his lasting impressions and, more importantly, his immediate recognition, these are downplayed.

What this example shows, I think, is that it is not just the organization or orientation of ‘concepts’ and specific features that can be operating in a person’s filter but, as it were, bodies of knowledge or belief. We can take up stances with respect to what we know, allowing ‘sets of knowledge’ to affect, in different ways, how we view contexts in which we find ourselves. The body of knowledge that Ed and Ned gather before coming to New York is an example. Importantly, Ed and Ned are examples of the way in which roughly the same body of knowledge, their prior knowledge, together with their attitudes can combine to affect the way in which agents come to see contexts that confront them.²³⁵

As we have seen, part of what is important about the desire to notice is the object of the desire — a noticing. The example of Ed and Ned illustrates that the way in which what we notice is affected by what we know in more complex ways than the discussion in the previous section suggests. It is not simply a matter of ‘conceptual orientation’ but can involve, as Ed and Ned’s did, whole bodies of knowledge. If the ‘desire to notice’ idea is to prove satisfactory it must be able to accommodate this kind of complexity. It must be able to deal with knowledge and bodies of knowledge. It must remain explanatorily

²³⁵ Notice that Ned’s adventurousness (for example) might be the product of his own attempts to change himself and his values. If he is now to be considered ‘adventurous’ then we would expect him to conduct his life in this way — i.e. we would expect him to view the world as an adventurous person would. [Here, character, desire and knowledge come together.]

useful in cases where agent's revise, reconsider, reflect on and deliberate about their take on a context. The extension of the desire to notice, then, fundamentally amounts to a broadening of the possible objects of the desire to notice.

However, before examining the ways that these requirements may be met we must return to the basic desire-like state at the heart of the issue. To do this we will consider, very generally, the possible ways that these desire-like states may enter the fray. Recall what we mean when we talk of desires and attitudes: desires are one among many manifestations of 'Dionysus'. There are also concerns, commitments, 'favourings', etc. This, perhaps, is why 'attitudes' or 'desire-like states' is preferable to 'desires'. Now the attitudes we can have might be distinguished in the following two ways, according to their objects — i.e. that to which they are attitudes:

- (i) Attitudes to the world, or perhaps better, to representations of the world (in short, attitudes to representations). Examples of this kind are what we usually consider when thinking of desires: the desire for the weather to be nice, the desire to eat dinner and the desire to go for a walk. Here the object of the attitude is a representation.
- (ii) Attitudes to attitudes, including attitudes to sets of attitudes and, naturally, 'higher order' attitudes. Here, we include desires to desire, like wanting not to want the last portion of ice-cream. The object of the attitude is another attitude.

However, depending, in part, on how we understand the first of these categories, we may need to suggest another broad kind of attitude, again, distinguished by its object.²³⁶ This kind will include the desire to notice.

(iii) Attitudes to *ways* of representing, attitudes to recognising, attitudes to deliberating, attitudes towards reasoning, attitudes to investigating and, perhaps more generally, attitudes to how what we know influences (or is to influence) what we see.

The suggestion is that the 'desire to notice', as it was discussed in the last section, is a special case of this third, more general attitude-type: an attitude towards how one sees the world or an attitude to allow what one knows to influence the way one sees ('takes') the world to be. It should be stressed however, that there is an important difference between having an attitude towards what is known and having an attitude towards how what one knows is applied or how what one knows informs (influences) one's understanding of the world. This third attitude-type involves, not caring about what we know but caring about how our knowledge affects the way that we see the world. In this respect it follows the 'structure' of the desire to notice. Just as the bully could be said to be on the lookout for or inclined to notice shyness and sensitivity or other features indicative of bullying opportunities, so Ed (for example) might be looking for or inclined to notice features indicative of the harshness and brutality of the city.

²³⁶ One might be tempted to think that the 'process' of noticing should be thought of in the same way as the 'process' of walking. This would not be unreasonable. It does, however, presume a good deal (perhaps too much) about the reduction of the mental to the physical. The thought is that noticing is a 'psychological' phenomenon in a way that walking is not and since we are concerned with psychological explanations, I need not be concerned here about the reduction. Prichard's paper 'Acting, Willing and Desiring' (*Moral Obligation*. Oxford: Clarendon Press, 1949: 187-98.) is useful on this matter.

Part of what underpins this third attitude-type is the fact that in most normal cases there are various ways of representing a given context (forming a take). Moreover I am normally *able* to represent given contexts in at least some of these various ways. These ways of representing can be crucially dependent on knowledge or, as it were, the processes of representing. Now if I possess this knowledge and so can represent contexts in various ways, one might suspect there to be an explanation of why I, perhaps consistently, represent contexts in one kind of way rather than another. It is in the service of these explanations that this third type of attitude is relevant.

On a more concrete level, it is not hard to imagine examples that fit this type. We might notice a tendency in ourselves to think about certain kinds of problems or situations in particular and perhaps, unsatisfactory ways. We might end up wanting to think differently about these problems or situations. Similarly, Ed might want to change his outlook on life. If it doesn't come naturally, he might actively try to see things in the way that Ned does. In endeavouring to explain his trying we would suggest, I think, that he wants to see the world differently.

Perhaps the most useful way of proceeding here is by example. This is partly because I do not think that there is any one way of characterising this extension of the desire to notice idea. After discussing two examples I will say something more general about the kinds of knowledge that might be involved. In each of these examples and in the ensuing discussion, it should be remembered that the focus is on instances of attitudes of the third attitude-type above.

Example One.²³⁷

When Fred was an undergraduate he studied Thucydides' history of the Peloponnesian War. He learnt that Thucydides was concerned to represent and to trace out the cycles of history. Thucydides mapped not only the cycle of fortune for 'empires' — in this case, the Spartans and the Athenians — but the cyclical fortunes of individuals like Pericles. Fred was particularly taken by this and by the ways in which the various cycles were interrelated. Now, what Fred learnt from Thucydides affects the way that he thinks about problems that confront him. Indeed it is not only this particular way of viewing his circumstances that Fred learnt from Thucydides but also the more general idea that the patterns that are exhibited throughout history are useful in understanding the events of the present. Fred has learnt that history can provide one with a deeper appreciation of the course of human lives. Fred cares about history. This concern manifests itself in many ways — he values historic landmarks, he values attempts to preserve them and he values the projects that historians undertake in the attempt to understand history more thoroughly. Additionally, Fred likes the fact that he cares about history. He enjoys and is proud of being a champion of history-related causes. More importantly for us though, is the fact that Fred thinks it is important to allow history or an understanding of history to affect the way in which one thinks about the present.

Now the former set of Fred's concerns illustrate, what we might call, the normal projectivist story. They involve the first two kinds of attitude: attitudes to 'things' and attitudes to attitudes. He cares about landmarks etc. and likes the fact that he cares about

²³⁷ A similar, but perhaps more involved story could be told about literature. There are a number of connections here to some of the work of Martha Nussbaum and Cora Diamond.

history. However the last of Fred's concerns does not straightforwardly fit into either of these categories. It is a concern, a preparedness or a standing readiness to allow a particular body of knowledge, in this case, a knowledge of history, to influence the way in which he understands circumstances with which he must deal. Moreover, this concern explains why Fred understands (reads, conceives of, forms 'takes' on) contexts in the way that he does. It is the value that he places on historical knowledge and particularly, the role that he gives it, which explains why he attends to contexts in the way that he does.

Fred's knowledge of history and his concern to 'read' contexts in the light of this knowledge might help to explain at least some of the judgements of obligation that Fred makes — those in which his knowledge of history figures in his understanding of the situation or better, those in which his knowledge helps to explain why certain features of the context present themselves to him as particularly salient and important (here). The reading of a context in the light of history will play up certain features and, as it were, locate others²³⁸ and thus, lead to judgements about what is called for — i.e. about what ought to be done.

Fred himself of course, might not make the connection between his knowledge of history and the world that confronts him. That is, he might only think of himself as (simply) caring about history. He might never recognize or notice the way his knowledge of history affects what he sees. Alternatively, he might sometimes (say, when asked) affirm his commitment to looking for the historical perspective or he might stress the value of the lessons of history for understanding the present.

²³⁸ Say, within historical narratives or sequences of events, perhaps with a particular point, a lesson or, simply, an outcome.

It should be pointed out that Fred's understanding of contexts in which he finds himself can only be influenced by the historical knowledge that he already possesses — i.e. obviously if he knows nothing of the history of China, he cannot be influenced by it. As a result his knowing more history means that there is more to influence the 'takes' that he forms. That is, someone who knows more (either about specific matters or more generally) is better placed to allow what they know to affect their view of the world. We will return to this below.

This however, is only one form that this third attitude-type may take in Fred's case. Fred may, more simply, be 'oriented' towards historical considerations. He might (consciously or unconsciously) seek out the historical bases of problems or situations. He might, as we say, have an historical 'bent' without knowing a great deal of history. Interestingly, this historical 'bent' affects not only his initial take on contexts but how his reflections about his situation proceed. Seeking out historical bases of problems and situations is a 'process' that can influence his view of what others do and the various relationships between those with which he interacts. Alternatively, Fred might have focussed primarily on the lesson of Thucydides' exploration of the cycles of human history (and not be particularly interested in history more generally). Thus his view of others generally, his relationships with fellow workers and his view of both his and others' fortunes might be influenced by the way in which these cycles are traced by Thucydides. Again these cycles not only might be read into Fred's view of the world but may affect his ongoing reflections about his surrounds.

Example Two:

Suppose that Martha considers herself a Christian or, at least, tries hard to be a good Christian. (Let us suppose, anyway, that she is fairly successful.) Her commitment to being a Christian is one of her most cherished values and it is a key feature of the way in which she identifies herself — when she thinks of who she is or of the defining features of her life, it is this commitment that features most prominently. However we cash out this complex (identifying) value, it undoubtedly involves, at some level, desire and attitude. Because of her desire to be a good Christian, we can assume that she also values doing (and hence, noticing) all of the things that a good Christian would. Alongside these attitudes is her knowledge of what it is to be a good Christian. ‘Good Christian’ here means good at being a Christian, where what it is to be a Christian is understood (for the purposes here) to be relatively clearly defined. It is not implausible to think that the set of the Church’s teachings, the Bible, stories of the lives of saints and practical guidance given in sermons and the like make up a relatively coherent whole that can be understood as a conception of what it is to be a good Christian. So Martha has, as a fundamental, identifying value: her commitment to being a Christian. She also knows, very well and in great detail, what it is to be a Good Christian.

Now in her daily life Martha does many good things for others. Indeed the vast majority of her good deeds are things that the rest of us think are well beyond anyone’s obligations. She does them as we might say ‘out of the goodness of her heart’. The trouble is that this is *not* how Martha sees what she does. She sees her actions as strictly things that she is obliged to do — the things that any good Christian would and ought to do. She does not want to work at the soup kitchen on Friday and Saturday nights nor is it

exactly correct to say that generally she wants to help people. She is happy to do these things but that is because she's generally content with her life. If we are to understand her making the judgements of obligation that she does, then her desires associated with being a good Christian enter into the explanation only at the very general level. When pressed for an explanation of why she volunteers to work at the soup kitchen, in the end, she might agree that it is because she wants to be a good Christian. But this is not how things strike her in her day-to-day thoughts and judgements about what she does and ought to do. What she knows about being a good Christian influences both how she understands herself and how she understands the world that she inhabits.

There are a number of general remarks to be made regarding these examples and the extension of the desire to notice idea. First, let us be clearer about the object of this extended version of the desire to notice — i.e. the object of the third attitude type. The important parts of the initial (un-extended) account of the desire to notice were its object, the 'noticing', and the related 'ability to notice'. Both, as we stressed, were most definitely classed amongst the cognitive states and were closely related to a general conceptual or linguistic knowledge — a part of which is the knowledge of how to apply the relevant concepts. As I think is clear from the above examples, the kind of knowledge involved in the extended idea of the desire to notice is much broader than a 'basic' conceptual/linguistic grasp. How might we characterize the broader involvements of knowledge?

1. The involvement of knowledge.

We might first distinguish the role played by practical knowledge.²³⁹ Practical knowledge, we might say, is closely related to matters of practice. It is often likely to be specific or particular.²⁴⁰ The rookie bully learns, for instance, that some of his potential victims only appear shy and sensitive. It is no doubt an important lesson for him to learn how to distinguish the more finely grained signs of good and bad bullying opportunities. In some cases the only way for this to be achieved is by trial and error — through experience. That is, he might simply have to find out that Billy is not to be trifled with. Of course experience is only one part of the equation here. There might be more or less general rules of thumb concerning the appropriate ways and extent of antagonizing in addition to those rules about who are the best victims. These ‘practical rules of thumb’ may be imparted by more experienced bullies or may simply be observed by the rookie bully in operation.²⁴¹ The lack of practical knowledge is also evident in the case of the (merely) well-intentioned person. Such an individual might not know, say, that homeless people sometimes spend the money given to them on alcohol and drugs, thus perpetuating their problems.²⁴²

²³⁹ I do not wish, here, to imply any hard and fast distinction between practical and theoretical knowledge.

²⁴⁰ For the distinction here between specific, particular, general and universal see chapter three, section D.

²⁴¹ Unlike the discussion in chapter three, these practical rules of thumb are not normative rules. They are rules about ‘the way the world goes’ not how it should go.

²⁴² Of course knowing this may not affect the overall judgement about whether to give or not.

Second, we might distinguish the involvement of ‘sets’ of knowledge from that of particular knowledge. Here we might think of Martha’s body of knowledge of what it means to be a Christian and the way, as a whole, it influences the way that she sees the world. Similarly, Ed and Ned’s picture of New York affects the way in which they see the city and its inhabitants. Here, however, the scope of the influence is less than in Martha’s case since it primarily applies only when they are in New York. The differing attitudes and so the differing priorities that Ed and Ned have with respect to what they know of New York affects how their knowledge of the city is applied, not that it is applied. For Martha, presumably, there is no real ‘time off’ from being a good Christian except perhaps when she is alone.

Finally we may speak of, as it were, ‘operational’ or ‘procedural’ states being the object of this attitude-type. We might have stances or approaches to how we represent the world. Here, the prime example might be Fred’s historical ‘bent’. In looking for the historical perspective on situations, Fred has an ‘investigative procedure’ in place. In the business of understanding situations that confront him this procedure gives priority to related past events. Similarly, it would not be too surprising to find that Martha ‘looks for the good in people’ — that is, she looks for and interprets people as having largely good (i.e. nice, charitable, well-meaning) motives, qualities and intentions. This again, might reflect a stance that she has adopted as a result of her commitment to being a good Christian.

There is a sense in which these categories are somewhat artificial as applied to the desire to notice. Whether we think of the object of the desire to notice as a body of practical knowledge, a process or a stance depends largely on other explanatory factors.

In some cases we might think that talking in terms of bodies of knowledge is appropriate whereas in others it may be a general level of practical knowledge and experience might be important. The respective categories are clearly related. Indeed, all of the examples we have discussed involve each category to some extent. Each of them involves the agent being 'on the lookout' for different kinds of things, sets of things, features or relationships. This can be characterized in terms of a 'bent' or stance, an operational procedure of looking for certain features or as having certain features at the forefront of one's mind.

2. Ranked knowledge.

One of the things that this kind of attitude achieves is the application of an ordering or ranking of knowledge. The desire to notice, as we saw, allowed that an agent could rank his 'conceptual apparatus' in such a way that a take on a context could have a kind of depth: playing up some features and ignoring others. The bully found features related to antagonizing particularly salient in his view of contexts. Similarly, in the extension, what the agent knows can have the same kind of ranking. Although Martha knows a good deal about the design and manufacture of clothes (say) this is only a side interest. She is dedicated, first and foremost, to the business of being a good Christian. The considerations and features involved in doing so always precede her interest in fashion. She notices when help is required before she notices how others are dressed. Of course the 'conceptual ranking' and this 'knowledge ranking' are related. Ed knows that New York is a hard place and the various ways in which this 'hardness' is manifest. If

this knowledge is to the fore, then he will be on the lookout for those features that are related to the hardness of New York.²⁴³

3. The object of the desire to notice and its extension.

The extension of the desire to notice helps with the problem surrounding the object of the desire that we confronted earlier. The examples in the first section of this chapter seemed to suggest that agents have very particular desires to notice — a desire to notice shyness, a desire to notice sensitivity or a desire to notice, specifically, old ladies in need of help. This, as I pointed out, is somewhat implausible. Instead, we should think of the object of the desire to notice in much more general terms. We should think, for instance, of the bully's desire to notice as being more along the lines of a desire to notice those features of contexts that are conducive to bullying. Initially this meant that the bully had a standing readiness to notice such features as shyness, sensitivity, vulnerability, etc. With the extended idea we can see the role that is naturally played by each of the 'categories' of knowledge discussed above — i.e. the bully needs to know what features of contexts are conducive to bullying (i.e. (say) that Billy is not to be trifled with or that isolation is often associated with vulnerability). What is most important in the extension is that the desire to notice is seen as one manifestation of a more general attitude-type, one that serves to organize or prioritise the agent's 'cognitive orientation'. So, caring about bullying (wanting to antagonize) here leads the bully to prioritise and utilise his bullying-related knowledge over other 'sets' of knowledge.

²⁴³ This ordering, as it was in the unextended case, is not necessarily fixed, nor is necessarily stable or clearly defined (cf. Wiggins, 'Deliberation and Practical Reason', *op.cit.*, p. 233.). Having said this, typically we might expect some stability with respect to 'moral' features. So although an individual's attention may vary with context — in the context of a baseball game, baseball related features are to the fore — we might expect that 'moral' features are more constantly to the fore.

It is in this way that general values can be applied. In various ways in chapters two, three and four we have seen that there is a gap between general values and what is noticed — the way in which the agent sees, describes, construes their contexts. Through these chapters and the later ones there has been a distinction between the agent's noticing/ability to notice and the agent's values. The desire to notice (extended) enables this gap to be bridged. It depends upon the agent's values (or more general desires) but involves a stance concerning what is noticed or what the agent is ready to notice.

4. Reflection.

The extension of the desire to notice idea allows more than the agent's 'first glance' to be included in our explanatory story. Agents clearly can and do reflect on their initial 'takes' and consider 'deeper' or, simply, other features of a particular context. Because the extended desire to notice incorporates knowledge and bodies of knowledge more generally, it allows more scope for such reflection, deliberation and reconsidering of one's context.

5. Judgements of obligation.

The way in which the extended idea of the desire to notice assists with the explanation of judgements of obligation and with the experienced pressure associated with them is much the same as for the simpler account of the desire to notice. The successful 'operation' of a desire to notice invests the noticed features with a significance that enables those features to stand out with respect to other noticed features (and, trivially, with respect to unnoticed features). The difference is that now the object of the desire to notice can be more complex and can involve more than simply features or properties like shyness or sensitivity. Of course there remains a strong connection to

these features. The application of the rookie bully's knowledge that Billy is not to be trifled with might be that he pays more attention to Billy and factors Billy's proximity into the decisions that he makes about what to do. Similarly Martha's knowledge of what it is to be a good Christian informs her desire to notice those features that make up the content of that knowledge. These developments still allow that the agent's (final, initial or intermediate) take has a 'depth of field' that it might otherwise not have had. The contours of an agent's values and desires are reflected in a more complex way, not only in what is noticed, but in the agent's concentration and deliberation about the circumstance at hand. Again this explains why certain features are taken by the agent to be salient or important for a 'proper' understanding (according to the agent) of the context — they are those that stand out in the agent's survey of the context, the agent's take. In those cases where the noticed features are sufficiently central — i.e. the desire to notice is sufficiently strong —, they will present the agent with a demand for consideration — demanding to be acted or judged upon. This, again, is how features and considerations come to have the 'to-be-dealt-with' quality that we observed and have been calling the experienced pressure accompanying judgements of obligation.

6. 'Accidental' noticing.

It might be complained (particularly with respect to the history example) that much of this noticing can be put down to mere frequency, an accident of repetition. The historian (the lawyer, the geologist or the philosopher) notices what he notices not because of a desire to notice but because historical (legal, geological or 'philosophical') facts and features of situations are the things he deals with day in and day out. The main claim to be defended against this kind of objection is that whenever there is a judgement

of obligation made, some desire to notice is operating (can be understood to be operating) in the background. This is *not* to claim that whenever there is a noticing (or attending) there is a judgement of obligation or whenever there is a noticing there is a desire to notice. That is, it is not to claim that the possession of a desire to notice or a particular 'noticing' is sufficient for a judgement of obligation nor that a noticing is sufficient for the possession of a desire to notice. So how then, are we to distinguish cases of attending that are accidents of repetition from those that are the 'product' of a desire to notice?

To press the point and before answering the question, let us consider the following example that, I think, encapsulates a more pointed version of the objection. Albert is not a cycling enthusiast but each year for two weeks he follows the Tour de France — he watches the highlights on the television and reads the newspaper reports each day. As a result he spends more time than usual (namely some time) thinking about cycling. Out on the street he notices cyclists more often during these two weeks than at other times — not just any cyclists but those dressed in professional cycling gear. Suppose that on one particular occasion a cyclist's momentarily unattended bike has fallen onto the road and is likely to be damaged by a passing vehicle. Albert judges that he ought to warn the cyclist of what has and is likely to happen. Here, what explains Albert's noticing the cyclist's predicament does not look to be any kind of desire to notice. He notices the cyclist and his unattended bike primarily because he has been following the Tour de France. Albert's judgement, based on this, looks to be merely an accident.

This is a useful example because it illustrates the complexity involved in noticing. The important point though is that the cyclist makes up only one part of Albert's view of

the situation. Frequency or repetition might draw Albert's attention to the cyclist but his construal of the context is not thereby fully explained. Crucially, Albert is also alert to the possible danger of the busy road and the passing traffic. We can contrast Albert with an individual who has a similar (albeit passing) interest in the Tour de France but does not have the same desire to notice as we are taking Albert to have. This individual notices the cyclist but does not construe his predicament as one that requires action (i.e. does not judge there to be an obligation). Perhaps he finds the falling bike amusing, reminiscent of slapstick comedy. The difference between Albert and this individual lies partly in a more general stance or preparedness on Albert's part. Albert engages with the world in a way that the other does not. Certainly there is some engagement on the part of the other individual, but it is only for Albert that this engagement amounts to an obligation. For Albert, the unfolding scene presents itself as requiring something.

The lesson here is that cases of noticing that are not tied in some way to values or settled desires — i.e. accidents of repetition — are not capable, *by themselves*, of 'generating' judgements of obligation. What distinguishes accidental noticing or noticing tied to repetition from noticing that is the product of a desire to notice is their connection to more general or settled desires. One who, through frequency or accident, notices cyclists is to be distinguished from someone who wants to notice cyclists. Similarly, one who notices historical features (facts, etc.) because of frequency of exposure or accident is to be distinguished from one who is on the lookout for them.

Having said this, these 'accidental noticings' are not quite as problematic as they might first appear. Sometimes it is just lucky that an individual is in a position to notice what they do, just as it is sometimes lucky that an individual is in a position to, say, offer

help to a person in need. The fact that Albert might not have judged as he did were it not for the Tour de France does not tell against the explanation. The desire to notice is still necessary for an adequate explanation as the example above shows.

The main point of the above discussion of the involvement of knowledge in the extension of the desire to notice and the examples of Fred and Martha is not to play down the complex role of desires and attitudes, nor still to simplify the complexity of these attitudes. Instead it is to make room for and to insist upon the kind of attitude that relates, as it were, a body of knowledge to particular situations — i.e. to enable a way of valuing (or set of concerns) to be put into effect — to influence the way an individual sees the world. Perhaps the most important aspect of both the desire to notice and its extension (particularly with respect to judgements of obligation) is the way in which the relevant desire-like states operate *indirectly*. Since these attitudes have belief-like states or ‘cognitive processes’ as their object they are not direct responses to an agent’s context. Instead they play the role of the filter. They can thus explain why it is that an agent attends to the world (comes to form a particular take). Again, the central idea is that various bodies of knowledge are given priority in the agent’s ‘field of view’. Thus, practical knowledge, specific concentrations of knowledge or ways of approaching the world can be positioned so as to influence the way an agent attends to particular contexts. At a minimum this knowledge is a basic conceptual ability — the desire to notice explains the way in which one’s ordinary grasp of the linguistic tools are organized. However as we saw in the rookie bully case, it is difficult to limit the relevant knowledge to this grasp of concepts. Agents, guided by their concerns, gather and give significance

to both practical and non-practical knowledge that they learn particularly in the attempts to fulfil their concerns. Thus, given the right kind of background (or unconscious) concerns, what agents know is applied or related to contexts in which they find themselves in such a way that they judge themselves to have obligations. The example of Fred illustrates how a body of knowledge and concerns related to that body of knowledge can affect the way an agent understands contexts that confront him. The desire to notice explains how Fred's valuing of history is manifest in particular contexts without requiring a direct appeal to the value he gives to historical considerations. In the example of Martha this indirectness is more important because it allows us to explain the judgements of obligation that she makes. Martha's central concern is to be a good Christian where what it is to be a good Christian is given 'descriptively'. She judges that she has an obligation because she takes her context in a certain way — noticing those features that a good Christian would.

Now I do not wish to suggest that this discussion is anywhere near enough to adequately explicate the kind of attitude-type that I think forms, what I am calling, the extension of the desire to notice. My point in this section has been to sketch those parts of the extension that are particularly pertinent to the discussion of cognitivism and non-cognitivism in the context of judgements of obligation. The central idea and point of the desire to notice remains the key and this I hope to have adequately explained.

CHAPTER 7: COGNITIVISM AND AN AGENT'S CONCERNS.

In this chapter we consider the relationship between the developed account of the desire to notice of the last chapter and the cognitivist position sketched in chapter four. Whilst the desire to notice plays the key role in the overall account of judgements of obligation offered here, it is by no means to be taken in isolation.

Since the first section of chapter five the discussion has taken on a decidedly non-cognitivist flavour. This, of course, was a result of the arguments, concerning shyness and sensitivity, of section A of that chapter. The connection between the considerations of the last chapter and the version of cognitivism sketched in chapter four has not yet been addressed. Given this, it is useful to consider in some detail what I take to be the key moves in McDowell's paper 'Virtue and Reason'²⁴⁴. The paper as a whole can be taken to represent the beginnings of his development of the position which was discussed in chapter four. It is closely tied to an understanding of Aristotle's practical syllogism and it incorporates the kind of cognitivist position with which we have been concerned. As such it draws together the various strands of the discussion to this point. Of particular note here are the consequences of the discussion of the desire to notice for the McDowell/Wiggins position as well as the role of, what we have been calling, the general desire or concern that informs the desire to notice.

²⁴⁴ McDowell, John. 'Virtues and Reason' *The Monist* 62 (1979): 331-350. Reprinted in McDowell, *Mind Value and Reality*. Cambridge, MA: Harvard University Press, 1998: 50-73. Unless otherwise stated references will be to this version.

In what follows, I shall first present a sketch of the key part of McDowell's argument in 'Virtue and Reason'. We will then consider the three issues that connect with the discussion of chapter six. First, we will consider the relationship between the examples and explanations in McDowell's argument and our concern with the pressure and judgements of obligation. Second, we will examine the consequences of McDowell's argument for the involvement of the general desire (the agent's 'concerns'). Finally, we will consider the role that the desire to notice plays with respect to McDowell's argument and particularly, for his conclusions about cognitivism. Of special note is the second of these issues. It is here that we will be more precise about the relationship between the general desire and the desire to notice. This, it will be recalled, is the second question that needed to be addressed at the end of chapter five.

A. The Argument in 'Virtue and Reason'.

Recall from chapter four that, with respect to the practical syllogism, the thesis of uncodifiability means that the envisaged major premise, in a virtue syllogism, cannot be definitively written down. Any attempt to capture it in words will recapitulate the character of the teaching whereby it might be instilled: generalizations will be approximate at best, and examples will need to be taken with the sort of "and so on" which appeals to the cooperation of a hearer who has cottoned on. (p.67)

This is a consequence of Wiggins' discussion and rejection of the rule-case interpretation of the practical syllogism. It is also justified by the discussion in Nussbaum, McDowell and Wiggins of Aristotle's view of the role of generalizations in ethics. As such it

parallels the discussion of rules in chapter four. In two papers particularly, ‘Deliberation and Moral Development in Aristotle’s Ethics’ and ‘Some Issues in Aristotle’s Moral Psychology’²⁴⁵, McDowell develops a way of reading both the idea of *eudiamonia* and moral development which does not require an externally sanctioned account. Once these accounts are in place the need for a ‘blueprint’ of what counts as the best way to live dissipates. Instead McDowell suggests a “Neurathian” account of ethical reflection. This account enables an agent to reflect on their own conception of living and doing well (tied, by Aristotle, to *eudiamonia*) from within that conception.

Neurathian reflection on an inherited scheme of values takes place at a standpoint within the scheme; the scheme can be altered piecemeal, but not suspended in its entirety, with a view to rebuilding from the ground up. (‘Some Issues’, p.37)

Notice here that an ‘inherited scheme of values’ can be taken to be that which is the product of an upbringing or, more simply, that scheme which one already possesses.²⁴⁶

²⁴⁵ ‘Deliberation and Moral Development in Aristotle’s Ethics’ in Engstrom, Stephen and Whiting, Jennifer. *Aristotle, Kant and the Stoics: Rethinking Happiness and Duty*. New York: Cambridge University Press, 1996: p.19-35. ‘Some Issues in Aristotle’s Moral Psychology’ in *Mind, Value and Reality. op.cit.*, pp. 23-49.

²⁴⁶ Of some significance here is the relation between this and the kind of view suggested by Blackburn in *Ruling Passions*. Blackburn utilizes precisely the same analogy when speaking of his quasi-realism and the way in which agents are able to be self-critical and reflective concerning their own moral perspective. He writes:

The criteria [for membership in the best set of attitudes and for improving one’s set of attitudes] are moral criteria, judged from within a moral perspective. That is, when I wonder how I might improve, I have to think about it deploying my current attitudes—there is no standing aside and apart from my present sensibility. But that does not mean that I have to deem myself perfect, or incapable of improvement. It just means that whatever flaws I suspect are judged as flaws in the light of other concerns (Neurath’s boat). But I might certainly suspect flaws other than sheer

One of the key points in McDowell's position is the claim that the virtuous person's "perception might suffice, on its own, to constitute a reason for acting in a certain way". This is how McDowell gets his brand of cognitivism. Part of the idea is that the demand or the call to action (i.e. the pressure) that is captured in the judgement of obligation is understood if we can understand how the agent saw the situation. That is, the agent's perception of the situation, alone, is enough to constitute a reason for action. Even though it involves a prudential rather than a moral matter, the example McDowell gives in 'Are Moral Requirements Hypothetical Imperatives?' is useful here.

Suppose, for instance, that we explain a person's performance of a certain action by crediting him with awareness of some fact that makes it likely (in his view) that acting in that way will be conducive to his interest. Adverting to his view of the facts may suffice, on its own, to show us the favourable light in which his action appeared to him. No doubt we credit him with an appropriate desire, perhaps for his own future happiness. But

inconsistency: immaturity, lack of imagination, bias, coarseness, and so on. (Appendix, *RP*, p.313.)

It is here that McDowell's and Blackburn's positions are closest. The kind of objectivity that is arrived at via McDowell's 'Neurathian' reflection is not the kind of objectivity that is sought by those who look for an "externally validated" conception of doing well. McDowell points out (and hence distinguishes his view from those who would look for such an externally validated conception) that the disappointment in Neurathian reflection is the result of thinking that "something as historically contingent as a conceptual scheme" could not yield the appropriate standard of objectivity — it could not give us "undistorted contact with the real" ('Some Issues', p.37) The differences between the two (McDowell and Blackburn), of course, remain (in the same way that they did so earlier). Blackburn speaks only of attitudes and stances that are given by attitudes and then he uses the tools of quasi-realism to account for our ethical talk ('reconstructing' the apparent objectivity). McDowell, on the other hand, speaks of, as it were, an internally validated objectivity — a kind of objectivity that is given by the rest of the argument of 'Virtue and Reason'. These remarks support the suggestions made in chapter four, section D about the status of justification that are, I take it, the major issue between Blackburn and McDowell/Wiggins.

the commitment to ascribe such a desire is consequential on our taking him to act as he does for the reason we cite; the desire does not function as an independent extra component in a full specification of his reason, hitherto omitted by an understandable ellipsis of the obvious, but strictly necessary in order to show how it is that the reason can motivate him.

Properly understood, his belief does that on its own.²⁴⁷

Importantly, the ascription of the desire is secondary to the explanation of why the agent's 'view of the facts' casts the action in a favourable light. In other words, the agent's take on the context shows the action in a favourable light — the weightings and saliences form a description of the context in which the demand to act in a certain way, felt by the agent, is clear. This does not mean that the desire in question (the general desire, for us) is irrelevant, only that it plays a lesser role than might otherwise be supposed.

In order to demonstrate this claim McDowell distinguishes between the 'core explanation' and the 'supplementary explanation'. The core explanation is that in which the desire figures, along with the relevant fact, to explain the action (judgement). The supplementary explanation details why the relevant fact/description is relevant or salient for the agent.

The former of these, "the core of a satisfying explanation of an action which is in fact virtuous", is perhaps better illustrated by the pair of psychological states in the following example:

²⁴⁷ 'Are Moral Requirements Hypothetical Imperatives?', *op.cit.*, p.79. Notice that it is of no concern to us that this is an example of a prudential rather than a moral act (cf. chapter one). So long as it involves a judgement of obligation (which this presumably does) it is of relevance here.

A concern can mesh with a noticed fact about a situation, so as to account for an action: as, for instance, a concern for the welfare of one's friends, together with awareness that a friend is in trouble and open to being comforted, can explain missing a pleasant party in order to talk to the friend. (p.67)

But this explanation does not relate the agent's action to his conception of how to live — that, which for both McDowell and Aristotle, distinguishes the virtuous from the non-virtuous and so, makes this act a virtuous one. Recall that for McDowell, it is one's conception of how to live, when adequately put into practice, which makes an agent virtuous (and makes an action right).

The core explanation must at least be seen against the background of the agent's conception of how to live; and if the situation is one of those on which any of several concerns might impinge, the conception of how to live must be capable of actually entering our understanding of the action, explaining why it was this concern rather than any other which was drawn into operation. (p.68)

Quite clearly the example above, the core explanation, fails to achieve this: it would be a satisfactory explanation of someone acting "in spite of thinking it was not, in the circumstances the thing to do" just as much as someone doing as they thought they ought (p.67).

McDowell's suggestion is that we can make this distinction by seeing how or by allowing the agent's conception of how to live to enter into the explanation of the action. This is where the supplementary explanation enters. It is at this level

that knowledge of how to live interacts with particular knowledge: knowledge, namely, of all the particular facts capable of engaging with concerns whose fulfilment would, on occasion, be virtuous. This interaction yields, in a way essentially dependent in appreciation of the particular case, a view of the situation with one such fact, as it were, in the foreground. (p.69)

Having thus given a sketch of his position McDowell turns to the now familiar non-cognitivist objection:

To talk of virtue — a propensity to act in certain ways for certain reasons — as consisting in a sensitivity, a perceptual capacity, is to amalgamate the required appetitive component into the putative sensitivity. (p.56)

The target is, of course, the *perception* that takes the place of the minor premise in the core explanation. [Recall that McDowell's claim is that an adequate account of the virtuous person's view of the situation will be enough to give the agent reason to act and, hence, that the concern involved in the major premise is insignificant.] Thus, the non-cognitivist claims that McDowell is wrong to classify the 'perception of a salience' (in this case, the recognition that one's friend is in trouble and open to being comforted) as a cognitive state. Instead the so-called perception is to be properly understood as a genuine cognition along with some additional appetitive state. McDowell's response to this objection is as follows:

But what appetitive state? Concern for one's friend yields only the core explanation, not the explanation in which the "perception" of salience was to figure. Perhaps the conception of how to live? That is certainly an

orectic state. But, given the thesis of uncodifiability, it is not intelligible independently of just such appreciation of particular situations as is involved in the present “perception” of salience. (p.70)

Thus, both options, the ‘general concern’ and the (more general) conception of how to live, fail as candidates for the relevant appetitive state. This allows McDowell to conclude that the virtuous person has a special cognitive ability that is conditioned by the complex of concerns and knowledge involved in a conception of how to live.

B. Three Issues.

There are three matters that need to be addressed concerning this argument and its bearing on the suggestions that have been made in the last chapter. First, we need to be clear about the relationship between McDowell’s argument — in particular, the split between the core and the supplementary explanations — and the our task of explaining judgements of obligation with special reference to ‘the pressure’. Second, we need to examine the consequences of McDowell’s argument for the relationship between the ‘general desire’ and the desire to notice. And third, we need to consider the potential of the desire to notice to answer McDowell’s query about the appetitive state involved in the non-cognitivist’s attempt “to represent the ‘perception’ of a salience as an amalgam of the purified awareness with an additional appetitive state.” (p.70)

Issue 1:

For McDowell, as we have seen, there are two levels of explanation: the core and the supplementary explanations. The core explanation is modelled on, as it were, the

straightforward applications of the practical syllogism schema — i.e. where the major premise registers the relevant concern or desire²⁴⁸ and the minor premise captures the relevant or salient features of the particular situation. The supplementary explanation primarily involves the minor premise and endeavours to relate what the agent takes to be the relevant features with that agent's conception of how to live.

Our understanding of such a conception [of how to live] enters into our understanding of actions — the supplementation that the core explanation needs — by enabling us to share, or at least comprehend, the agent's perception of saliences. (p.69)

Importantly, if we de-emphasize the heavy involvement of the agent's conception of how to live in this [and instead allow, more generally, the involvement of general concerns or values of some sort], the supplementary explanation becomes an explanation of why it is that the agent perceives the saliences that he does. That is, the supplementary explanation enables us to see how an agent's general concerns enter into the agent's perception of saliences. Clearly, this is the task that we have had since chapter two: to explain how an agent comes to attend in such a way that results in the experienced pressure and the judgement of obligation.

Despite this parallel McDowell's discussion and examples are not explicitly focussed on judgements of obligation. Similarly there is no reference made to the phenomenology associated with judgements of obligation — the experienced pressure. As we saw in chapter five, there can be no conscious desire-like state to help if the agent

²⁴⁸ Cf. "The role played by the major premise, in these straightforward applications of the schema, is to give the content of an orectic psychological state: something we conceive as providing the motivating energy for the actions explained." (p.66)

experiences the pressure (and perhaps makes a judgement of obligation). Those considerations equally well apply here. The concern to help one's friend is ruled out not because it yields only the core explanation but because it does not permit the appropriate phenomenology — that is, unless the concern in question is unconscious. This makes all the difference because it now seems as though part of McDowell's point about the core explanation aligns with our points about the phenomenology.

Part of the point here is that when we are explaining judgements of obligation or acts that are consequent upon such judgements there is a gap between our explanation and the agent's concerns. That is, whilst we would be quick to attribute a desire in non-obligation cases (say, wanting more ice-cream and having another serving) we cannot move immediately to the desire-like state in the obligation case. This is part of McDowell's point about the conception of the situation providing reason enough to act. When it is clear that an agent is doing as they ought (say, helping the old lady across the road), the natural explanation is *not* the attribution of a desire-like state. Instead, (and this is what seems right about McDowell's remarks) our explanation is itself couched in normative terms, like 'obliged' or 'required' and in terms of the relevant features of the situation. Thus, we might say that the agent judged that he ought to do as he did because he noticed (recognized) that the old lady needed help. Our explanation will show how "the agent's conception of the situation, properly understood, suffices to show us the favourable light in which his action appeared to him" ('Are Moral Requirements Hypothetical Imperatives?', p.80).

Issue 2:

A consequence of McDowell's separation of the core and supplementary explanations is the potential for explaining the virtuous person's judgement. In short while the core explanation applies to the virtuous person *at some level* it will not do all of the required explanatory work. It will not explain why the virtuous person judged that such and such ought to be done. Now in chapters five and six the 'general desire' has been taken to be the catalyst for the desire to notice. The general desire most definitely is a part of the core explanation. [Cf. Martha might eventually (when pressed) acknowledge that the her desire to be a good Christian underwrote what she did.] But if the separation of the core and supplementary explanations is required for an adequate explanation of the virtuous person's judgement then this general desire will be in the background (as part of the core explanation) and the supplementary explanation will be to the fore.²⁴⁹ This gels well with our discussion. Since chapter three the key to explaining judgements of obligation has been the agent's take — i.e. how the agent views contexts in which he finds himself. Thus the way in which McDowell down-plays the core explanation, for the reasons that he does so, means that we similarly can down-play the significance of the general desire in the explanation of judgements of obligation. Of course this is not to deny the 'existence' of the general desire, it is to remove it from the immediate explanatory 'vicinity'.

The way in which this general desire is downplayed can be illustrated by the following example:

Scenario 1 (Yesterday):

²⁴⁹ Cf. Petit, Phillip and Smith, Michael. 'Backgrounding Desire' *Philosophical Review*. Vol. XCIX, No.4. (October, 1990): 565-92.

Barney gets into the train, sees a seat, sits and begins reading his book. As he gets off at his stop he notices a pregnant woman standing, with some discomfort, nearby. He feels badly, thinking that he ought to have offered her his seat or that he ought to have noticed the woman. He resolves to try to be more observant in the future.

Scenario 2A (Today):

Barney gets onto the train, notices the seat and a pregnant woman. He asks her if she would like the seat. She declines saying that she is fine. Barney takes the seat.

Scenario 2B:

Same as 2A but when the woman declines Barney notices something — maybe something about her tone of voice or about her stance — which suggests that she is being polite. As a result, Barney thinks that she indeed does want to sit, so he offers again, perhaps adding some excuse suggesting that he was not particularly committed to sitting: “I’ve been sitting all day.”

For our purposes here scenarios 2A and 2B are the important ones: Scenario 1 sets the scene. We might say that after scenario 1 Barney reaffirms his commitment to giving up his seat for pregnant women. In our terms he either gains or strengthens his general desire to be helpful and his desire to be more attentive (his desire to notice). [We might add that he has recognized a failure in his attentiveness — his resolve is to be more attentive, say, particularly in ‘train’ situations.] So scenarios 2A and 2B are the test cases. The differences between these two are differences in the content of Barney’s desire to notice.²⁵⁰ In the former case he notices the pregnant woman. In the latter, he also notices

²⁵⁰ Notice that we could add another variant of 2B, namely:

features of her demeanour that indicate politeness (etc.). Thus in the latter case he was also attentive to features other than her being pregnant. Otherwise put, his standing readiness to notice was broader in the latter case. The desire to notice thus provides an account of why Barney was *looking* or, to use McDowell's kind of terminology, receptive to the details and nuances.

Let us be clear about the consequences of this example for the role of the general desire. In both 2A and 2B Barney's concern to give up his seat for pregnant women remains the same. The variation, as we have seen lies with his attentiveness. Thus our explanation of Barney's different judgements in the two scenarios will not include reference to his general concern. If we also assume that, in both cases, Barney is able to notice as he does in 2B then, on the model proposed here, the difference lies with Barney's desire to notice.

Overall we are torn between two tendencies: (1) there is a sense in which it is plausible to say that the virtuous person who helps the old lady, in some way, wants to help (has a general concern to help people etc.) but we must be careful about concluding that this concern is straightforwardly a desire; (2) the pressure inclines us away from attributing (at the very least) a conscious desire. The argument and the example of Barney incline us away from attributing a general desire to help *in explaining the*

Scenario 2C (alternative to 2B):

Same as 2B except that the cues that Barney detects or notices do not suggest politeness but pride. Accordingly he decides not to press the matter, thinking that to do so would risk offending the woman.

The difference between this scenario and 2B might primarily be a difference in the context itself. Barney's attentiveness to the nuances of the woman's response would presumably allow him to be alert to different kinds of nuances. Of course it might be that detecting the 'pride clues' requires a still more subtle ability than is required to detect the 'politeness clues'.

judgement. Why? Because the general desire does not cut the distinctions in the particular case finely enough to adequately explain the judgements made. The above example is a case where the 'general desire' remains fixed but where the resulting judgement differs. That is, an agent can have the general desire but fail to notice (even though they could notice). What does the bulk of the explanatory work here is the desire to notice and what is noticed. This shows that the general desire (conscious or not) is not required to explain the specific judgement. Again the general desire is a key part of the core explanation but does not feature in the supplementary one. When we are concerned to explain a judgement of obligation, we are concerned with the supplementary explanation.

In cases of bullying and the like, the 'general desire' may be, as it were, close to the surface. That is, perhaps only a few moments of reflection might reveal to the bully that his noticing (his being on the lookout) and his judging are closely related to his desire to antagonize. There is no need for a supplementary explanation here, partly because the desire is 'close, but more importantly because there is no judgement of obligation. However, as the case of Martha illustrates, the general desire at work may be very deeply embedded. It may indeed take Martha a good deal of reflection to realize the connection between her 'noticings'/judgements of obligation and her desire to be a good Christian. She might recognize that she tends to notice certain kinds of things and because of this takes herself to have related obligations. But for her to see a connection between the kinds of things she notices, her obligations and her concern to be a good Christian there would need to be a good deal of prompting. This 'distance' between her judgements of obligation and her general, identifying desire to be a good Christian is reflected in McDowell's distinction between the core and supplementary explanations.

Issue 3:

To this point in the discussion we have not directly considered the consequences of the desire to notice model for McDowell's argument. This is largely because I have been concerned to draw out the relationship between our discussion and the specific arguments given by McDowell in 'Virtue and Reason'. The third issue then is how the positing and development of the desire to notice and its extensions affect McDowell's argument.

The final step in the argument is the one that allows McDowell to conclude in favour of cognitivism. Recall that the non-cognitivist's objection is to McDowell's refusal to split the 'perception of salience' — in our terms, the agent's attention to his context — into a genuine cognition and an additional appetitive state. McDowell's response, again, is:

But what appetitive state? Concern for one's friend yields only the core explanation, not the explanation in which the "perception" of salience was to figure. Perhaps the conception of how to live? That is certainly an orectic state. But, given the thesis of uncodifiability, it is not intelligible independently of just such appreciation of particular situations as is involved in the present "perception" of salience. (p.70)

Quite clearly, I think, the desire to notice is just such an additional appetitive state. In precisely the required way, the desire to notice is informed by the agent's values/general concerns or, for McDowell, his conception of how to live. It provides the link between the agent's 'perception of saliences' and these more general concerns. The immediate

lesson here is that *this* argument for cognitivism does not go through. On the desire to notice model the agent's attention (rather than perception) is to be thought of as the combination of an ability to notice and a desire to notice. These two elements make up the agent's filter. It is by understanding this filter that we are in a position to explain the agent's judgements of obligation along with the associated pressure.

This is not the end of the story for this kind of cognitivism. The version of cognitivism that we have been considering (here and in chapter four) is 'desire friendly'. Wiggins' story about the evolution of evaluative concepts and the co-involvement of property and response is not inconsistent with the suggestions here (and in particular with the suggestions involving the desire to notice). The decision about cognitivism rests in this case with the justificatory strategies associated with or akin to the idea of 'there being nothing else to think'. What I take the arguments here to show is that truth, as is could be applied to moral judgements of obligation, cannot be secured through the analysis of the psychological precursors to the act of making a judgement of obligation. The cognitivist must take a route other than through the kind of states responsible for judgements of obligation in the effort to show that such judgements are capable of being true or false.

CHAPTER 8: THE FRAME OF MIND.

The aim of this thesis was to come to some conclusions about the nature of moral judgements of obligation. Crucially what is meant by 'nature' here are the necessary psychological preconditions of such judgements. Thus the aim is to provide the necessary psychological preconditions for moral judgements of obligation. Throughout the argument I have been concerned with the act of judging, thus 'judgement' refers to this *act* rather than to the content, what is judged.

In our ordinary practices of explaining and understanding the behaviour of others and of explaining our behaviour to others we can attribute/ascribe a whole range of mental states to agents. The juxtaposition of various of these states explains what people do and say by reconstructing, in the terms and according to the rules of the practice, the (partial) configuration of the agent's mental states which combined (or would combine) to issue in the behaviour being explained. The task then was to come to some understanding of the preconditions of moral judgements of obligation — namely, that configuration of mental states which leads to (explains) the making of the judgement. If these preconditions are to be necessary preconditions then the account was to provide a largely structural or schematic description of the relevant portion of the configuration of an agent's mental states just prior (in this explanatory sense) to the making of the judgement. I have called this configuration the agent's 'frame of mind'.

Within this practice of providing psychological explanations, the primary evidence is, not surprisingly, what the agent says and does. In addition *some* weight is to

be accorded to the agent's experience of the relevant 'processes' or states — their phenomenology²⁵¹. This is particularly evident in how we understand our own behaviour and in checking the explanations of our behaviour given by others.²⁵²

The general strategy has been to construct a quasi-transcendental argument which begins with a general feature of our experience of making moral judgements and proceeds by endeavouring to uncover those preconditions which could both explain our judgements and do justice to this feature of our experience. It remains a 'quasi-transcendental' argument because I have not argued (though I think it is plausible to think) that the general experiential feature is a necessary feature accompanying moral judgements of obligation.²⁵³

An important upshot of the concentration on the explanation of the act of judging rather than on the content is that the account aims to avoid making normative claims about, in particular, morality.²⁵⁴ That is, I have not been concerned to say which

²⁵¹ We must be mindful, of course, of the problems associated with an over-emphasis on introspection.

²⁵² Cf. Butler, Joseph, *op.cit.*, Sermon 1, §6, note.

²⁵³ The steps in the argument can be broadly characterized as arguments for one of the disjuncts in each of the following disjunctions: (i) external or internal source, (ii) general or particular internal source, (iii) belief-like or desire-like particular internal source and (iv) desire-like particular internal source with representations, other desire-like states or ways of representing as its object. These disjunctions can be thought of as nested — each being a division within the preferred disjunct of the previous step. Hence the argument is regressive. Each of the disjunctions, I will claim, exhaust the possibilities at that point, thus the necessity of each step is preserved. These are the two characteristics (besides the claim of necessity for the first, experience-based premise) that are required for an argument to be counted as transcendental. (See chapter one.)

²⁵⁴ There is/might be some normativity involved in the practice of psychological explanations with which we are concerned. (See Blackburn, *RP*, ch.3.) Whatever norms are involved will be those associated with this practice.

judgements are right, preferable or justified. A further consequence involves the possibility of saying precisely which judgements of obligation are moral judgements of obligation. Given this avoidance of normativity, it is not possible to say with any degree of specificity (or independence) which judgements are judgements of moral obligation. Any attempt to carve out those judgements that involve a 'moral matter' will itself involve a normative assessment. The central claim involved here is that what makes a judgment of obligation a judgement of moral obligation is determined by the content of the judgement.

A relatively standard distinction in the literature is between judgements of value and judgements of obligation (or evaluative and practical).²⁵⁵ By and large this distinction is between judgements involving what ought to be done (by myself or others) and what is good, preferable, and so on (evaluative judgements tend to involve the predication of a value term to some object or event). It is this distinction that enables my concentration on judgements of obligation. In spite of this distinction it is natural to suspect that an account of judgements of obligation will, at some point, run into an agent's values. Thus part of the project has been to demonstrate/facilitate an explanation of how values interact with and are manifest in an agent's judgements of obligation (without presuming an account of value).

The argument has proceeded with the help of three assumptions. The first of these is that is simply that we make moral judgements of obligation.

²⁵⁵ See, for instance: Frankena, W.K. *Ethics* Englewood Cliffs, NJ: Prentice hall, 1963: 9; Wiggins, David. 'Truth, Invention and the Meaning of Life' in *Needs, Values, Truth* (*op.cit.*), p.95; and, Darwall, S, Gibbard, A and Railton, P. 'Toward *Fin de siecle* Ethics' in *Moral Discourse and Practice* (*op.cit.*).

In chapter one I claimed that when we make moral judgements of obligation, it is consequent to feeling a pressure which seems to originate independently of us. There are two features that are characteristic of the pressure and which must be accounted for in the explanation of moral judgements of obligation: (1) the motivational character (the conative force) of the pressure; and, (2) the 'external' or 'objective' character of the pressure. Further, we can distinguish between two kinds of moral judgements of obligation: direct and indirect judgements. Direct judgements involve what I or we ought to do, whereas indirect judgements involve what he, she or they ought to do. I have primarily considered the former kind since the latter can be understood as a complex version of direct judgements.

Now either this pressure originates independently of us or it does not. If it does not originate independently of us then it is primarily the result of a configuration of our frame of mind. I claimed, at the beginning of chapter two that it does not originate independently of us. Thus, whenever we feel this pressure (which seems to originate independently of us), it is a result of our 'frame of mind' being configured in a certain way. Despite this, it might be claimed that although the pressure does not originate independently of us, our obligations do. That is, an agent might in fact have an obligation even though he does not judge that he does. Hence, an argument more powerful than one which only focuses on the pressure is required. This is provided by Prichard in 'Duty and Ignorance of Fact'. Prichard's argument was useful because it is not specific about the kind of account that should be given of the grounds of obligation (i.e. necessary and sufficient conditions for an agent's having an obligation) — it is purely negative in that it only seeks *to rule out* one of the options. A second reason that this stronger argument was

required, rather than something like Mackie's arguments from relativity and queerness, is that, unlike Mackie we are not in a position to rule out cognitivist or realist views like intuitionism. The argument from queerness, in particular, presumes that intuitionism is untenable. This was not a luxury that we could afford at that stage of the argument. There is a strong sense also in which the cognitivist positions put forward by McDowell and Wiggins in particular, take the claims of intuitionism to be precisely what requires explanation. Being clear about Prichard's argument also enabled us to say more about the distinction between judgements of obligation and judgements of value. In particular, there are some judgements that look to be judgments of obligation but which actually function more like judgements of value and thus should be taken as falling on the 'value' side of the distinction.

In chapter three we began to consider the 'shape' that the frame of mind might take when we make judgements of obligation. Again we began with the following disjunction: either the pressure depends on general features of the configured frame of mind, such as rules, or it depends on how the agent attends to particular features of the context. I argued that it does not depend on general features. Thus, whenever our 'frame of mind' is configured in this certain way (i.e. such as to result in the judgement and the pressure), we have attended to particular features of the context in a certain way.

This is not to say that there is no role for rules or generalized sentiment. Both can and do play a role in how we understand the judgements that an agent makes. The key point however is that an agent can make such judgements without possessing any specific, predetermined set of rules or sentiments. Thus the possession of any given set of rules or general sentiments is not necessary for making moral judgements of obligation.

Again, this is not to say that the possession of any rules or general sentiments whatsoever is not necessary for moral judgements of obligation. What was in question is the form of these rules or sentiments and the role that they play in the making of moral judgements of obligation. The points involved here were largely related to those surrounding particularism. I argued, not that particularism is correct simpliciter, but that it is not necessary to give the set of rules that the agent takes himself to be following in order to explain the judgements that he makes.

At the end of chapter three we could make the following claim: the particular configuration of the 'frame of mind' that results from attending to the particular features of the context is a necessary condition for the experience of the pressure (and hence for making a moral judgement of obligation). The focus was on attention and thus we required an explanation of the kinds of states that are 'operating' when an agent attends in a particular way. It was noted that why an agent attends in the way that he does might be the result of habit, training or experience. Thus the kind of explanation we sought needed to have something to say about these seemingly more 'settled' characteristics.²⁵⁶

At this point, we confronted the second, more significant assumption: the tools that are available to explain both the judgement and the two characteristics of the pressure are states that are broadly characterized as belief-like or desire-like. This assumption was the product of the kind of explanation that we sought and, broadly speaking, the rules of the practice of giving psychological explanations of an individual's actions (including, I take, their judgements).

²⁵⁶ It is worth noticing that from this point the order of this argument sketch diverges from the order in the text. This is only a matter of aesthetics both here and in the text.

The debate between cognitivists and non-cognitivists can be taken to be primarily about whether or not the contents of moral judgements are capable of having truth-values. One way of attempting to decide this question is to consider the kind of mental state involved in or responsible for the moral judgement. The idea is that if the moral judgement is explained by primarily cognitive or representational states then it is capable of having a truth-value whereas if the judgement is explained by primarily non-cognitive states — that is, states understood as responses to the world — then it is incapable of having a truth value. These issues frame the debate (and have done so roughly since Hobbes) and gave us the tools that we could use in our account. This is not to say that these distinctions are final. A failure to find an adequate account of the kind we required may well have led to a rejection of this way of thinking about the problem²⁵⁷. Corresponding to this debate, we made the following distinction. A mental state that is belief-like is one whose role is to represent (or help in representing) the world: beliefs, knowledge and perception. A desire-like mental state is one which plays the role of responding to the world (responses to the world): desires, attitudes and sentiments.²⁵⁸ Further, corresponding to this distinction there are various accounts that tend to privilege one of these states over the other in their explanations of moral judgements. Typically

²⁵⁷ It is here that we run into the issues raised by McDowell's attack on the 'eighteenth-century philosophy of mind'.

²⁵⁸ This distinction is similar to the distinction suggested by Anscombe and discussed by Michael Smith and Mark Platts. Platts writes: "The distinction is in terms of the *direction of fit* of mental states to the world. Beliefs aim at being true, and their being true is their fitting the world; falsity is a decisive failing in a belief, and false beliefs should be discarded; beliefs should be changed to fit with the world, not vice versa. Desire aim at realization, and their realization is the world fitting with them; the fact that the indicative content of a desire is not realized in the world is not yet a failing *in the desire*; and not yet any reason to discard the desire; the world, crudely, should be changed to fit with our desires, not vice versa." (Platts, *Ways of Meaning*, p.257.)

those accounts that privilege belief-like states are cognitivist accounts whilst those that privilege desire-like states are non-cognitivist accounts.

In the first section of chapter five I argued that an account that privileges belief-like states in the explanation of moral judgements of obligation (typically, cognitivist accounts) could not accommodate the motivational character of the pressure.²⁵⁹ With a number of qualifications, the main claim was that for any given account of the cognitive capacity/state (the possession of which is claimed to be responsible for moral judgements of obligation), it is possible for someone to possess that capacity/state and not make the corresponding judgement.²⁶⁰

Turning to the second alternative, I suggested that, in general, desire-like states could have three kinds of objects: (1) representations (things, states of affairs, etc.), (2) other desire-like states or (3) ways or modes of representing. Importantly this is a very general assumption about the possible *contents* of desire-like states. The simplest example of a desire-like state with an object of kind (3) is what may be called a ‘desire to notice’. Noticing is presumably a cognitive activity and it seems possible for an agent to want to notice or, more colloquially, be looking to notice certain features rather than others. A more complex example might be an attitude to seeing the world in a certain way. Here one may realise that there are various ways of seeing the world — ways which

²⁵⁹ It is specifically the motivational character of the pressure that is problematic, rather than motivation simpliciter. This is a point about the typically cognitivist and externalist route: judgement as a cognition plus a desire, say, the desire to be moral or the desire to do one’s duty.

²⁶⁰ This is a development and extension of the argument and counterexamples given by Blackburn and Smith against McDowell. (See particularly, *Ruling Passions*, pp.92-104.)

emphasize certain features or sets of features rather than others. A preference for one of these 'ways of seeing' over the others would count as a desire-like state.

The third section of chapter five argued that an account which relied on desire-like states with objects of kind (1) or (2) cannot accommodate the second, 'objective' characteristic of the pressure. The main thrust of the argument for this claim was a mismatch between the phenomenology of desire and the second characteristic of the experienced pressure.²⁶¹ The 'feel' of normal desiring is, as it were, 'subjective' relative to that of the pressure. Given this mismatch it is hard to see how desire-like states with objects of kind (1) and (2) can be used to explain both features of the pressure. The appeal to unconscious desires (i.e. desires with no phenomenology) will not help. The presence of an unconscious desire can only explain why there is no phenomenology, not why there is a phenomenology of a wholly different sort. In short, either the relevant desire-like states (i.e. those desire-like states that are to explain the judgement) have a 'feel' or they do not. If they do then we have the relevant mismatch. If not, then the pressure is not explained.

I have argued, particularly in chapter six, that it is the third kind of desire-like state that can provide us with a satisfactory solution. That is, an account which centred on desire-like states which have 'ways of representing' as their object is well placed to account for both characteristics of the pressure that need to be accommodated. In virtue of these states being desire-like they are well placed to explain the motivational character of the pressure (1). It should be noted that the ability to satisfy desire-like states of this kind is tied to habit and training in a distinctive way and thus they are likely to be without

²⁶¹ I take it that any decent account of desire will have something to say about the phenomenology of desire.

the usual desire-like 'feel'. They are perhaps better understood in terms of a standing readiness or preparedness to notice or consider one's context in a particular way. If one of these desires is satisfied then I will attend to a given context in a certain way. Whatever pressure I experience will be a result of the way that I have attended. Thus the pressure will feel as though it 'comes from the world' rather than from myself — satisfying the second characteristic. A useful analogy here is that of a 'filter' or 'lens'. These states operate to sift and colour the way one views a context. "The emotional and moral lens is not readily visible to the person who sees through it." ²⁶²

Thus, the configuration of the 'frame of mind' that is necessary for making a moral judgement of obligation has as its centrepiece a desire-like state, the object of which is a way of representing the world. This central desire-like state can be thought of as a standing readiness or preparedness to allow what one knows (a particular body of knowledge) to influence how one sees the world or how one attends to contexts. Habit, training and experience play a distinctly important role in the ability of the agent to satisfy this state (hence the importance of education and experience). The ability itself — the ability to represent the world in that certain way — is a cognitive one and is also a required component of the frame of mind. Invariably, there are many other desire-like states associated with the central one however these may take many forms, no particular one of which is required.

There is a sense in which this account sits (or attempts to sit) between the cognitivist and non-cognitivist positions. Although it is granted that the 'centrepiece' of the frame of mind is a desire-like state, knowledge and other belief-like states are not

²⁶² Blackburn, *RP*, p.7.

involved in an ‘inert’ way. Instead, this centrepiece trades on the idea of a desire-like state that is indirect, fixing not on objects or events ‘in the world’ but on the way the agent attends to the world, i.e. involving the way the cognitive apparatus of the agent functions in particular contexts. If we take the cognitivist/non-cognitivist debate to be one about moral truth, then this account definitely tends towards the non-cognitivist side. The central desire-like state ensures this. However, McDowell, Wiggins and Winch²⁶³ have all made suggestions about the way in which it might make sense to speak in terms of the truth of moral judgements but in a highly particularized and context sensitive way — i.e. a long way from a universal normative theory. Moves like these rely, to some extent, on shifting the emphasis of the debate away from truth as it applies to the content of moral judgements and onto, as it were, the psychological precursors of the judging. In short, there is perhaps still room for some kind of (perhaps, weak) cognitivism — a cognitivism that can link the justificatory strategies associated with moral discourse and practice with truth without requiring that judgements of obligation are psychologically explained strictly by belief-like states.

²⁶³ Winch, ‘The Universalizability of Moral Judgements’, *Ethics and Action*. This paper and the matter at hand is discussed by Wiggins in *Needs, Values, Truth*, ch.4, ‘Truth, and Truth as Predicated of Moral Judgements’.

BIBLIOGRAPHY OF WORKS CITED

- Altham, J.E.J. 'The Legacy of Emotivism' in MacDonald, Graham and Wright, Crispin, eds. *Fact, Science and Morality: Essays on A. J. Ayer's Language Truth and Logic*. Oxford: Blackwell, 1986.
- Ameriks, Karl. 'Kant's Transcendental Deduction as a Regressive Argument', *Kant Studien*, vol. 69, 1978, pp.273-287.
- Anscombe, G.E.M. *Intention*. Oxford: Blackwell, 1957.
- Aristotle, *Nicomachean Ethics* (trans W.D. Ross) in *The Complete Works of Aristotle* (ed. Jonathan Barnes) Princeton: Princeton University Press, 1984.
- Black, M. *Philosophical Analysis: A Collection of Essays*. Englewood Cliffs, NJ: Prentice Hall, 1963.
- Blackburn, Simon. 'Reply: Rule-Following and Moral Realism' in Holtzman, S. and Leich, C. (eds.), *Wittgenstein: To Follow a Rule*. London: Routledge and Kegan Paul, 1981: 141-187.
- Blackburn, Simon. *Spreading The Word*. Oxford: Clarendon Press, 1984.
- Blackburn, Simon. *The Oxford Dictionary of Philosophy*. Oxford: Oxford University Press, 1996.
- Blackburn, Simon. *Ruling Passions*. Oxford: Clarendon Press, 1998.
- Broad, C. D. *Five Types of Ethical Theory*. New York: Harcourt, Brace & Co., 1930.
- Brown, S. C. (ed.) *Philosophy of Psychology*. London: Macmillan, 1974.
- Butler, Joseph. *Fifteen Sermons Preached at Rolls Chapel*, (1729) (ed. Revd. D. Matthews). London: Bell & Sons, 1958.
- Carritt, E.F. *The Theory of Morals*. Oxford: Oxford University Press, 1928.
- Cavell, Stanley. *Must We Mean What We Say?* New York: Charles Scribner's Sons, 1969.
- Cavell, Stanley. *The Claim of Reason*. Oxford: Clarendon Press, 1979.
- Cooper, John M. *Reason and Human Good in Aristotle*. Cambridge, MA: Harvard University Press, 1975.
- Dancy, Jonathan 'Ethical Particularism and Morally Relevant Properties', *Mind*, Vol. 92, 1983: 530-47.

- Darwall, S, Gibbard, A and Railton, P. 'Toward *Fin de siècle* Ethics' *The Philosophical Review*, Vol.101, No.1 (January, 1992): 115-189. Reprinted in Darwall, S, Gibbard, A and Railton, P. *Moral Discourse and Practice: Some Philosophical Approaches*. New York: Oxford University Press, 1996: 3-47.
- Darwall, S, Gibbard, A and Railton, P. *Moral Discourse and Practice: Some Philosophical Approaches*. New York: Oxford University Press, 1996.
- Davidson, Donald. *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press, 1984.
- Dennett, Daniel. *The Intentional Stance*. Cambridge, Mass.: MIT Press, 1987.
- DesAutels, Peggy. 'Gestalt Shifts in Moral Perception' in May, L., Friedman, M. and Clark, A. (eds.), *Mind and Morals: Essays on Ethics and Cognitive Science*. Cambridge, MA: MIT Press, 1996.
- Engstrom, Stephen and Whiting, Jennifer. *Aristotle, Kant and the Stoics: Rethinking Happiness and Duty*. New York: Cambridge University Press, 1996.
- Evans Pritchard, E. E. *Witchcraft, Oracles and Magic among the Azande*. Oxford: Clarendon Press, 1937.
- Ewing, A. C. *The Definition of Good*. New York: Macmillan, 1947.
- Falk, W. D. 'Ought' and Motivation.' *Proceedings of the Aristotelian Society* 48 (1947-48): 492-510. Reprinted in Falk, W. D. *Ought, Reasons and Morality*. Ithaca: Cornell University Press, 1986: 21-41.
- Falk, W. D. *Ought, Reasons and Morality*. Ithaca: Cornell University Press, 1986.
- Frankena, W. K. 'Obligation and Ability' in Black, M. *Philosophical Analysis: A Collection of Essays*. Englewood Cliffs, NJ: Prentice Hall, 1963: pp.148-65.
- Frankena, W.K. *Ethics*. Englewood Cliffs, NJ: Prentice Hall, 1963.
- Frankfurt, H. *The Importance of What We Care About*. Cambridge: Cambridge University Press, 1988.
- Gaita, Raimond. *Good and Evil: An Absolute Conception*. London: Macmillan, 1991.
- Gibbard, Alan. *Wise Choices, Apt Feelings: A Theory of Normative Judgement*. Cambridge, MA.: Harvard University Press, 1990.
- Gregory, R. L. 'Perceptions as Hypotheses' in Brown, S. C. (ed.) *Philosophy of Psychology*. London: Macmillan, 1974: 195-210.
- Guyer, Paul. *Kant and the Experience of Freedom*. Cambridge: Cambridge University Press, 1996.

- Haldane, John. and Wright, Crispin. *Reality, Representation and Projection*. New York: Oxford University Press, 1993.
- Hobbes, Thomas. *Leviathan* (1651) (ed. Edwin Curley). Indianapolis: Hackett, 1984.
- Holtzman, S. and Leich, C. (eds.), *Wittgenstein: To Follow a Rule*. London: Routledge and Kegan Paul, 1981.
- Humberstone, I. L. 'Directions of Fit' *Mind* (1992): 59-83.
- Kant, Immanuel. *The Critique of Judgement* (trans. J. C. Meredith). Oxford: Clarendon Press, 1952.
- Kripke, Saul. *Wittgenstein on Rules and Private Language*. Oxford: Blackwell, 1982.
- Lovibond, S and Williams, S.G. (eds.) *Identity, Truth and Value: Essays for David Wiggins*. Oxford: Blackwell, 1996.
- May, L., Friedman, M. and Clark, A. (eds.) *Mind and Morals: Essays on Ethics and Cognitive Science*. Cambridge, MA: MIT Press, 1996.
- MacDonald, Graham and Wright, Crispin, eds. *Fact, Science and Morality: Essays on A. J. Ayer's Language Truth and Logic*. Oxford: Blackwell, 1986.
- MacIntyre, Alasdair. *After Virtue* (2nd Ed.). London: Duckworth, 1985.
- Mackie, J. L. *Ethics: Inventing Right and Wrong*. (1977) London: Penguin, 1990.
- Mandelbaum, Maurice. *The Phenomenology of Moral Experience*. Baltimore: Johns Hopkins Press, 1955.
- McDowell, John. 'Are Moral Requirements Hypothetical Imperatives?' *Proceedings of the Aristotelian Society* Supp. Vol. 52. (1978): 13-29. Reprinted in McDowell, *Mind Value and Reality*. Cambridge, MA: Harvard University Press, 1998.
- McDowell, John. 'Virtue and Reason' *The Monist* 62 (1979): 331-50. Reprinted in McDowell, *Mind, Value and Reality*. Cambridge, MA: Harvard University Press, 1998.
- McDowell, J. 'Aesthetic Value, Objectivity and the Fabric of the World' in Schaper, Eva. *Pleasure, Preference and Value*. Cambridge: Cambridge University Press, 1983.
- McDowell, John. 'Values and Secondary Properties' in Honderich, T. (ed.). *Morality and Objectivity*. London: Routledge & Kegan Paul, 1985.
- McDowell, John. 'Deliberation and Moral Development in Aristotle's Ethics' in Engstrom, Stephen and Whiting, Jennifer. *Aristotle, Kant and the Stoics: Rethinking Happiness and Duty*. New York: Cambridge University Press, 1996.

- McDowell, J. 'Some Issues in Aristotle's Moral Psychology' in *Mind, Value and Reality*. Cambridge, MA: Harvard University Press, 1998.
- McDowell, J. *Mind, Value and Reality*. Cambridge, MA: Harvard University Press, 1998.
- McNaughton, David. *Moral Vision*. Oxford: Blackwell, 1988.
- Millikan, Ruth Garrett. 'Pushmi-pullyu Representations' in May, L., Friedman, M. and Clark, A. (eds.), *Mind and Morals: Essays on Ethics and Cognitive Science*. Cambridge, MA: MIT Press, 1996.
- Moore, A. W. 'On There Being Nothing Else to Think, or Want or Do' in Lovibond, S and Williams, S.G. (eds.) *Identity, Truth and Value: Essays for David Wiggins*. Oxford: Blackwell, 1996.
- Moore, G. E. *Principia Ethica* (1903). Cambridge: Cambridge University Press, 1962.
- Nowell-Smith, P. H. Review of 'The Phenomenology of Moral Experience', *Philosophy*, 32 (1957): pp.170-3.
- Nussbaum, Martha. *The Fragility of Goodness*. Cambridge: Cambridge University Press, 1986.
- Nussbaum, Martha. 'The Discernment of Perception: An Aristotelian Conception of Private and Public Rationality' in *Love's Knowledge*. New York: Oxford University Press, 1990.
- Nussbaum, Martha. *Love's Knowledge: Essays on Philosophy and Literature*. New York: Oxford University Press, 1990.
- Petit, Phillip and Smith, Michael. 'Backgrounding Desire' *Philosophical Review*. Vol. XCIX, No.4. (October, 1990): 565-92.
- Platts, Mark. *Ways of Meaning*. London: Routledge and Kegan Paul, 1979.
- Prichard, H. A. 'Acting, Willing and Desiring' in *Moral Obligation*. Oxford: Clarendon Press, 1949.
- Prichard, H. A. 'Duty and Ignorance of Fact' in *Moral Obligation*. Oxford: Clarendon Press, 1949.
- Prichard, H. A. *Moral Obligation*. Oxford: Clarendon Press, 1949.
- Prior, A. N. *Logic and the Basis of Ethics*. Oxford: Clarendon Press, 1949.
- Raphael, D. D. (ed.) *British Moralists: 1650-1800* (2 Vols.). Indianapolis: Hackett, 1991.
- Ross, W. D. *Foundations of Ethics*. Oxford: Clarendon Press, 1939.

- Ross, W. D. *The Right and the Good* (1930). Indianapolis: Hackett, 1988.
- Smart, J. J. C. and Williams, B. *Utilitarianism: For and Against*. Cambridge: Cambridge University Press, 1973.
- Smith, Michael. *The Moral Problem*. Oxford: Blackwell, 1994.
- Taylor, Charles 'The Validity of Transcendental Arguments', *Proceedings of the Aristotelian Society*, Vol. 79, 1978-9, pp.151-65.
- Timmons, Mark. *Morality Without Foundations*. New York: Oxford University Press, 1999.
- Wiggins, David. 'Deliberation and Practical Reason' *Proceedings of the Aristotelian Society* vol.76 (1975-6). Reprinted, with revisions, in *Needs, Values, Truth* (1987) 3rd ed. Oxford: Clarendon Press, 1998: 215-37.
- Wiggins, D. 'A Sensible Subjectivism' in *Needs, Values, Truth*. Oxford: Clarendon Press, 1998.
- Wiggins, D. 'Truth, and Truth as Predicated of Moral Judgements' in *Needs, Values, Truth*. Oxford: Clarendon Press, 1998.
- Wiggins, David. 'Moral Cognitivism, Moral Relativism and Motivating Moral Beliefs'. *Proceedings of the Aristotelian Society* (vol. 91), 1990: 61-85.
- Wiggins, David. 'Cognitivism, Naturalism, and Normativity: A Reply to Peter Railton.' in Haldane, John. and Wright, Crispin. *Reality, Representation and Projection*. New York: Oxford University Press, 1993.
- Wiggins, D. 'Objective and Subjective in Ethics, with Two Postscripts about Truth.' *Ratio* N.S.8, 1995: 243-58.
- Wiggins, David. 'Truth, Invention and the Meaning of Life' in *Needs, Values, Truth*. Oxford: Clarendon Press, 1998.
- Wiggins, David. *Needs, Values, Truth*. (3rd Ed.) Oxford: Clarendon Press, 1998.
- Williams, Bernard. 'Practical Necessity' in his *Moral Luck*. Cambridge: Cambridge University Press, 1981.
- Williams, Bernard. *Moral Luck*. Cambridge: Cambridge University Press, 1981.
- Williams, Bernard. *Ethics and the Limits of Philosophy*. Cambridge, MA.: Harvard University Press, 1985: 129.
- Winch, Peter. 'The Universalizability of Moral Judgements' *The Monist* 49 (1965): 196-214. Reprinted in Winch, *Ethics and Action*. London: Routledge and Kegan Paul, 1972: 151-70.).

Winch, Peter. 'Understanding a Primitive Society' in his *Ethics and Action*. London: Routledge and Kegan Paul, 1972: 8-49.

Winch, Peter. *Ethics and Action*. London: Routledge and Kegan Paul, 1972.

Wittgenstein, Ludwig. *Philosophical Investigations*. Oxford: Blackwell, 1958.

Wittgenstein, Ludwig. *On Certainty* (1969). New York: Harper and Row, 1972.