

Please cite the Published Version

Hirst, David Alexander (2018) Epistemological and moral reasoning in medical ethics. Doctoral thesis (PhD), Manchester Metropolitan University.

Downloaded from: <https://e-space.mmu.ac.uk/622812/>

Usage rights:  Creative Commons: Attribution-Noncommercial-No Derivative Works 4.0

Enquiries:

If you have questions about this document, contact openresearch@mmu.ac.uk. Please include the URL of the record in e-space. If you believe that your, or a third party's rights have been compromised through this document please see our Take Down policy (available from <https://www.mmu.ac.uk/library/using-the-library/policies-and-guidelines>)

Epistemological and Moral Reasoning in Medical Ethics

David Alexander Hirst

A thesis submitted in partial fulfilment of the requirements of the Manchester
Metropolitan University for the degree of Doctor of Philosophy

Department of Interdisciplinary Studies

MMU Cheshire

January 2018

Abstract

The purpose of this paper is to go towards proving that the generalist ambitions championed in Medical Ethics discourse are both misguided and dangerous to that discourse's meaningfulness. It will be claimed that the generalist undertones which underscore the discourse are motivated by the looming threat of irrationalist accounts of moral reasoning. I first make attempts to show that generalism itself does not avoid this threat. I then aim to show that irrationalism, in turn, is predicated upon a particular and highly contestable conception of empirical epistemological ideas which are, post-analysis, simply unsustainable.

I use the work of recent particularist contributors to philosophical debates regarding meaning to show - via a reading of Wittgenstein's later work - which this empiricist epistemology fails. I then defend this reading of Wittgenstein against others, by trying to show that there is enough in the accounts of theorists such as McDowell to suggest that it can rescue meaning in ethical utterances.

Our discussion will take us through various areas of thought in claiming that a particularist account really does give the best explanation of moral reasoning. Epistemological, motivational, phenomenological and ontological concerns will all be assessed and contrasted with competing generalist claims. I also critically assess various particularist theorists' claims, making the case for a specific form of non-reductive ethics. This is cashed out by providing a commentary on the very recent debate regarding thick concepts.

As the particularist position I advocate is largely found by way of adjudication between the pre-existing positions, I cannot make much claim of originality here. But by way of applying a thoroughly worked out particularist account to Medical Ethics - one which explains what motivates generalist accounts - I aim to elucidate in an important way how the generalist underscoring is not just incorrect, but also damaging to ethical debates in this domain; as well as providing a more thorough account of how we actually can improve Medical Ethics via first-order ethical reasoning. These last points give the thesis its unique interest.

Acknowledgements

I would like to thank Michael Loughlin and Phil Hutchinson for sharing their time and wisdom with me, the department for having patience with me throughout times of ill health, and my examiners Tim Thornton and Anna Bergqvist for their fair and thought provoking feedback. This could not have been done without you.

Declaration

With the exception of any statements to the contrary, all the data presented in this report are the result of my own efforts. In addition, no parts of this thesis have been copied from other sources. I understand that any evidence of plagiarism and/or the use of unacknowledged third party data will be dealt with as a very serious matter

Signed: David Hirst

Date: 30th July 2018

Contents

Introduction	p. 7
Section 1: Practical and Theoretical Reasoning	p. 11
1.1 On the Content of Evidence	p. 15
1.2 The Form of Reasoning	p. 19
Section 2: Generalist Accounts	p. 25
2.1 Deontological and Autonomy Based Accounts	p. 26
2.2 Ross: Perception and Principles	p. 35
Section 3: Commensurability	p. 41
3.1 Utilitarian and Maximisation Accounts	p. 41
3.2 Moving Away from Value Monism	p. 51
3.3 Evaluation from Bottom to Top	p. 55
3.4 Naturalism, System and Motivation	p. 59
3.5 Motivation and Externalism	p. 74
Section 4: Moral Reasoning, Rules and Motivation	p. 78
4.1 Non-Cognitivism Revisited	p. 81
4.2 Sophisticated Non-Cognitivism	p. 85
4.3 Getting Things Right Without Foundational Rules	p. 102
4.4 Rejecting the Interpretive Regress Picture	p. 108

4.5 The Middle Position	p. 115
Section 5: Painting a Non-Reductive Moral Picture	p. 122
5.1 From Brute Movement to Norms	p. 122
5.2 Objectivity, Analogy and a Space for Morality	p. 126
5.3 Moral Motivation and the Idea of a Standard	p. 141
5.4 Relativity and Pragmatics	p. 151
Section 6: A Return to Thick Concepts	p. 166
6.1 Variability of Valence in Thick Concepts	p. 170
6.2 Objectionable Thick Concepts	p. 182
Section 7: Application to Medical Ethics	p. 187
7.1 Expertise in Medical Ethics	p. 192
7.2 Harris vs Maclean: Metaphysics and Inculcation	p. 195
7.3 Organising and Characterising the Practice	p. 211
7.4 Particularism and Practical Ethical Guidance	p. 220
7.5 Getting Moral Discourse Right – Concluding Remarks	p. 224

Epistemological and Moral Reasoning in Medical Ethics

Introduction:

This thesis aims at critiquing the academic domain of Medical Ethics, which I will define as the attempt to use methods which find their foundations in philosophy to help us tackle and supposedly solve ethical and practical problems that arise within the practice of medicine. The overarching thesis will contend that the way that the discourse there is implemented carries all the hallmarks of serious confusions in philosophy about how we can reason about ethical problems. Indeed, the majority of the thesis will be unapologetically focussed on deep philosophical assumptions which ground the sort of theory which makes itself available to be easily 'applied' to a practice.

It might make sense to give a small history, to elucidate the reasons for my choosing to not keep this thesis restricted to the more academically driven philosophical domain. Having graduated from my undergraduate, I decided to take up a course which could make the skills I had learned in philosophy more applicable to a job market. To this end I undertook an MA in Medical Ethics and Law. What I found being taught there, mostly to an audience of soon to be or already qualified medical practitioners became the motivation for this thesis. The ideas that were championed were precisely those that could be easily marshalled into systematic accounts and then applied to awaiting social systems, doing so at the cost of what I can only describe as vital first order moral goods. Everything I had come to love about ethics - studies of virtue, respect, integrity and meaning – were minimalised in a surprising way; firstly, in the name of practicality; but also in the name of a very narrowly conceived notion of rationality and objectivity. There seemed to be an unwavering belief that anything outside strictly storable, universally comprehensible and easily actionable principles and processes failed to really hit its target.

This thesis aims to undermine that picture, and thus has an eye on two separate opponents: not just the generalist but also the irrationalist. I want to show that it is a

certain conception of how we come to know the things we all feel very certain saying we do know, that when thinking philosophically, casts what I will term as ‘the unfavourable comparison’ upon ethical discourse. Furthermore, this conception does so in such a way that those who wish to defend morality are willing to create a caricature of it in order to keep its realist pretensions alive. The intuition that there must be something we can still say - in the face of an empiricist theory of mind and epistemology which almost rules the roost by fiat, given the epistemological success story of scientific practices – is admirable. Once we have bent and altered morality to fit around this epistemological conception of knowledge acquisition, however, if we no longer recognise the reasons that remain as being distinctively ours, then it seems evident we are doing a vast disservice to the whole ethical enterprise. It is surely incumbent upon us to make sure that both medical personnel and their patients understand the best reasons that are really there to be discussed are ones they already know intimately; not ones prescribed by experts guided by a very specific conception brought on by highly contestable philosophical worries.

As the targets of this thesis are wide ranging then, the thesis will necessarily be wide ranging as well. I will start by critiquing various versions of generalist theories, whose content, form, or both still have influence in the domain of practical ethics. I will show how each has limited either that form, or content, or both in order to comply with the overarching worries that are ever present in a field whose epistemological success is forever compared and contrasted to that of the scientific and empirical enterprise.

The first section of the thesis will attack the idea of overtly theoretical and generalist reasoning. By showing that no generalist theory works in a way that is intuitively acceptable, as well as diagnosing various methodological flaws in the attempt to ground such theories - either by way of overt rationalism or naturalism - I aim to steer us towards the jagged rocks of irrationalism.

Once we are face to face with this worry, and certain that we have nowhere inside reductive or definitional theoretical constructions left to hide, I will aim to show the flawed assumptions that are apparent in comparisons of morality with the scientific and empirical claims. This will require a review of reasoning itself. Epistemological, methodological and psychological factors of what it is to arrive at a good reason in

every field will hopefully start to show that the unfavourable comparison is also an unwarranted one.

My main line of attack will match that which is attributed mostly to John McDowell; but also that which has been taken up by various Neo-Aristotelian theorists. It aims at exploiting deep insights held within the unsystematic account of the later Wittgenstein in order to show once and for all that the empiricist theory of mind - the best explanation I believe we have as to why philosophically minded moralists so readily reject first order ethical concepts as valid points of discussion – rests upon a fallacy. That excursion will again put pressure on our ideas of objectivity, but from this position the comparison is not with another domain of thought. Empirical claims will be fumbling for a ground as much as moral ones are. Rather, from here it will be an intuitive ideal of objectivity which will be threatening the very idea of meaning and normativity in all forms of reasoning. Again, this follows the sorts of issues that McDowell was dealing with in the late 1970s and early 1980s. If we can salvage reasoning from this position without giving too much up, as I believe we can, we will emerge with empirical claims and moral claims both properly situated and understandable. Their distinctive qualities will still be recognisable, but neither one will need to be understood as holding primacy over the other. It is crucial to note that much of this attack on generalism and non-cognitivism must be attributed to McDowell.

The originality of the thesis finds its ground after these lengthy excursions into McDowell's work; excursions which at most act so as to bring together ideas which can be found in various articles into a clearer image. The originality of the thesis lies in both assessing the ways in which McDowell's own account might be bettered - here taking ideas from such theorists as Garfield (2001), Wiggins (2000, 2006), Thomas (2006) and Crary (2009) - to see which parts of each rendering we might keep and which we might lose to help give us a better grasp of the scope of the Wittgensteinian inspiration McDowell uses to defend ethics.

Section 6 includes what I claim is a crucial distinction between the McDowellian non-reductive particularism we develop and another particularist theory – namely that of Jonathan Dancy (2001). This comparison is inspired by an analysis of some tricky

issues regarding thick concepts which have garnered much recent attention and aim to undermine the non-reductive picture I am aiming to build.

It is only once all this is done that we will turn to the domain of Medical Ethics and Clinical Practice, to contemplate how the sorts of moral reasons we have earned the right to can help to remedy the issues which are prevalent in that academic domain, and to see if the sorts of generalist tendencies we hoped to resist at the outset can indeed be resisted. We will also look to see if the non-reductive particularist picture that we have built can have anything meaningful to contribute to the sorts of problems that arise in medicine. I will state now that my aim in this thesis is to discuss the moral, metaphysical and epistemological assumptions that underscore applied moral discourse. As such I will be aiming to stay out of the fray of any specific conflicts that might occur within that discourse.

Section 1: Practical and Theoretical Reasoning

Some Preliminary Comments:

Moral reasoning, if there really is such a thing, is about producing practical outcomes which are in some way the 'right' outcomes. When one is faced with a position in which one needs to work out 'what one ought to do' one is faced with a practical problem. Not all practical problems are, at first blush, anything like moral problems. When we decide what we should have for breakfast, which holiday destination would be suitable for our family, or whether to take out health insurance, we are faced with a practical problem. Whilst planning a family holiday I will be plausibly constrained by various factors: the time I can take off work, the expense, making sure the destination suits all the persons, that it has appropriate facilities for children or the elderly and so on. While such decision making is evaluative, the values involved are not immediately evident as being akin to moral values. Deciding whether grandma would value seeing the museums of Florence more than the natural beauty of the Gorge du Verdon and whether this preference should be seen as holding any importance to my final decision of where to go, is not the same as, say, deciding whether I value her autonomy enough to not extensively persuade her to try one more bout of chemotherapy, despite her previous protests that she would prefer not to go through the ordeal.

What we have for breakfast in the morning seems an even more obviously non-moral decision than choosing the holiday. Still, what I 'ought to do' seems to have some moral sense to it. If I wish my daughter to have her favourite breakfast of bacon sandwiches as a treat and there's only enough bacon for one, or if I should have a low cholesterol breakfast on my doctor's say-so, I may still here have an evaluative choice to make that, at least *prima facie*, moves beyond merely attending to my more selfish desires, non-controversially.

Practical choices, by hinging on which of several possible things one ought to choose are characteristically normative; and norms reflect values. Deciding to keep to a low cholesterol breakfast because it makes my spouse happy to see me following my doctor's orders – even though I'd much prefer the 'full English' - shows

that I value my spouse's happiness over my own immediate preference in this matter. Framed as such, the decision seems a somewhat moral one: underlying the choice of greasy meat or cornflakes is a decision between values, such as my own life expectancy, the happiness of my wife, my evaluation of long term health versus immediate pleasure. As nearly all practical decisions which are not a choice between two alternatives which might best be categorised as a value-equal preference ('cornflakes or branflakes?') have the potential to be decisions which exhibit values I hold about what kind of life to live, I will not yet make any major distinction between practical choices and moral choices. Non value-equal choices all exhibit some sense of evaluation of 'what matters the most to me' which will take the form of reasons in favour of my decisions, and the content and outcomes of these decisions seem apt for obviously moral assessments. If I choose the bacon over the cornflakes I might be open to criticism for weakness of will in not heeding my doctor, callousness towards my spouse, or greediness in denying my daughter. Not all theorists see the distinction this way, but drawing any borders around what constitutes a moral decision at this stage would beg the question against those who allow that nearly all practical decisions exhibit values that speak to moral evaluation, or those who believe none do.

Even for those who do take a reduced view, most would accept that moral decisions are all practical decisions, (even if that equation cannot be reversed). They are normative. They exhibit values; or at least our evaluations. The output of moral reasoning is practical; that is, it leads to actions. In contrast to this type of reasoning stands another type: theoretical reasoning. As stated in the introduction, above, when one takes part in moral philosophy one always has an eye upon objectivity. We want to say that we are in some way getting things right, that the end result (the action of booking the holiday to Florence as opposed to the south of France) is something we did because it really was the 'better' option. This implies that we can give reasons for our action which show its superiority to other options on the basis of the decision making process we went through. We justify our actions with reasons. If these reasons are to mean anything then they are set against a standard, the purpose of which is to show that the decision was a better one. The postulation of a standard aims at 'grounding' the decision as something that was not just an arbitrary choice; rather it shows how this choice is of greater value than another option. But

what standard we might invoke when asked by an interlocutor who disagrees is never instantly apparent in the practical case. It is not, in the face of disagreement, something we can simply point to in order to settle the issue once and for all. The very fact that the interlocutor disagrees implies her own standard is different than our own. Her standard says that alternative factors should have been given greater value when the decision was being made (“why just focus on Grandma’s wants – the kids would be much happier in France!”)

At first glance, most pieces of theoretical reasoning have none of the problems that are associated with practical reasoning in this regard. Rather than the output of such reasoning being an action, or a motivating belief that justifies an action, it is instead just a belief – something which, if the result of good reasoning, reflects the way that the world actually is. And unlike practical reasoning, where one can only attempt to ground one’s claims via invoking a standard which is itself open to critical assessment, most examples of theoretical reasoning have an obvious ground: the world as it is anyway. The output is thus an ‘is’ claim. The distinction between the two types of reasoning thus brings out the logical distinction famously made by Hume: the ‘is/ought’ distinction. ‘Is’ claims are logically distinct from ‘ought’ claims in various ways. What ‘is’ the case seems non-controversial in its essence. A good piece of theoretical reasoning ends with a belief that reflects the real world, the world as it is absent of our own biases, values and preferences and exists in its own right as a set of causes and effects. A belief can be tested and shown to be a belief that is in some very real sense ‘right or wrong’. Crucial to such claims are modes of verification, and in the natural sciences we have a method of investigation that gives us insight into how accurate our ‘is’ claims are. A belief is a simple proposition that thus can be established as being true or false, or at least a close approximation dependent upon the modes of verification or falsification available, against a standard that is not in any sense observer dependent; what keeps making a belief true or false is the picture of the natural world against which it is cast. Moreover, the features of the natural world are subject to constant forces in the form of constant causal physical laws which help us to predict effects from causes, or work out what caused a certain phenomenon. These laws are the very stuff of more complex examples of theoretical reasoning. These combined features all contribute to us feeling that when we reason to a conclusion of the ‘is’ form, assuming we have the

right sort of evidence at hand, then the output is a belief that can be called knowledge.

The contrast just drawn between theoretical and practical reasoning leaves the latter in a bad light in epistemic terms. The features which distinguish the two are unfavourable in such a way as to suggest that even to address the mental processes involved in practical reasoning as though they constitute an actual form of reasoning is no more than a charitable misnomer. It will be the purpose of this section to alleviate practical reasoning from the burden of this contrast in two ways. Firstly I will press the idea that the evidence in theoretical reasoning is non-controversial. Secondly, I will attack the idea that the deductive form that leads to a theoretical conclusion in which such evidence is utilised is a necessary element for the output of a piece of reasoning to be described as knowledge.

The suggested threat to practical reason rests on a contrast of the content and form that constitutes a piece of such reasoning to that exhibited in theoretical reasoning. By showing that theoretical reasoning has similar problems to the ones which give us pause when considering its practical cousin, we will be left with two options: take up a sceptical position on reasoning as a whole, or accept that the prima facie advantages of theoretical reasoning discussed above are not necessary markers for attributing knowledge to the output of a piece of reasoning. I will opt for the latter option, with further arguments for this position coming later in the thesis. The final part of this section will show how lessening the distinction between the two sorts of reasoning is crucial to getting on in the world. While what 'is the case' and what one 'ought to do' will remain logically distinct, I will address the idea that reasoning always has a purpose, and that in using reasons in the characteristic way that we do involves a mesh of both practical judgement and theoretical judgement. This will involve canvassing various options offered by moral theories for addressing the distinctively practical content of moral issues. As we will see there, the epistemic form of theoretical reasoning still has a serious draw which in turn shapes the content of most moral theories.

1.1 On the content of evidence:

As stated above, practical reasoning seems to be at a deficit to theoretical reasoning in an epistemic sense. The intuition that guides this thought is that the various features available to theoretical belief claims which are not recognisable in ought claims (that the natural world grounds their rightness or wrongness; that they are straight forwardly factive; that they have an accepted form of verification...) act as markers for the correct application of the concept of 'knowledge' by a competent user. As has already been mentioned, if we wish for our moral claims to be taken seriously - as well as the majority of practical claims, given their conspicuously evaluative nature - then there must be a way to show that the discourse, if it is to be a meaningful one, is in some sense 'truth-oriented'. For us to make an assertion that the conclusion of a piece of reasoning is true we need an epistemic 'hook'; something which shows why we believe that this claim, and not other claims which oppose it, is truthful (and by a matter of inference that the opposing claims are false). Given the above analysis of the differences between the two forms of reasoning, however, if we accept that the features that theoretical reasoning has which practical reasoning lacks constitute necessary markers for a piece of reasoning to lead to a conclusion that can be claimed to be truthful then the idea that practical reasoning can lead to knowledge seems under serious threat. Is the moral discourse a bogus one then; one which is given an unjustified air of objectivity by the complex language we have developed around moral and evaluative concepts which underwrites our moves within it? (Gibbard, 2003)

It should first be noted that there is a form of practical reasoning that is acceptable even to those who believe that the natural world, as science reveals it to us, contains everything that 'is' the case and therefore the totality of true propositions. When one makes an 'ought' claim, one is expressing a statement that says not how the world is, but how it 'ought' to be. We form an intention to act so as to fulfil this conception. If I would prefer Florence to Marseille, or think I ought to have a word with Grandma about her attitude to various cancer treatments, then it is rational that I set about to act so as to bring the world into line with my end (whatever that may be) by assessing the various means to achieving it. The various accounts which can accept this form of instrumental (ends-means) reasoning, whilst still denying the truthfulness or universality of the 'ought' which is its output, are labelled 'non-cognitivist' theories. We might say that such theorists would claim that finding the means to what we think

we 'ought to do' is rational, but that the proposition 'I ought to do...' is not apt to be called truthful or a piece of knowledge, for it is hedged in terms of a relative form of rationality.

If I believe that Grandma should try chemotherapy then I ought to think of the best means by which to persuade her to try chemotherapy. But the 'ought' is not universal. It reflects only my own intentions, values, desires and ends. And these values themselves have no place inside the conception of the natural world as science reveals it to us. The desires which drive me are non-cognitive; a projection of my own ideals which, as they remain devoid of an epistemic hook to anything that exists in the natural world, cannot be the output of a truth-oriented piece of reasoning and thus can be declared to be neither right nor wrong; they just are.

What *can* be assessed, once my ends are known, is the effectiveness of the means which I choose to achieve my end. If one's end is to get to Florence as quickly as possible, one acts irrationally if one chooses the boat instead of the available aeroplane as the best available means. In this diminished sense, even by the non-cognitivist's lights, practical reasoning is still possible and its output is rational. Reason speaks only to the relations between the physical causes in the world, though, and never to values which drive the practical reasoning; its end or aim. Whether one should have a particular aim is not a matter of reasoning; how one achieves one's aim once it is non-cognitively given is. Once again this highlights the idea that theoretical reasoning has a distinct advantage over practical reasoning. The underlying idea is that the norms that are in place for the acceptance of certain aims or goals as being better than others cannot be grounded by anything but our contingent social practices or mental constitution; how we were raised to value X over Y, or how we are psychologically constituted so as to value X over Y. But social practices and the makeup of our mental states seems an arbitrary and purely relative matter. The presence of norms in our decision making thus threatens the notions of knowledge, rationality and truth in any end-norm. Those things are present, however, in finding the means to that end.

But the naturalistic conception of theoretical reasoning also has a normative slant, just deeper down. If we can show that this is indeed the case then the contrast between theoretical and practical reasoning used to cast aspersions upon the latter

form of reasoning will be shown to be an unfair contrast. Let us take a simple case of a theoretical assertion: 'My mother is wearing a green dress today'. At first blush this statement seems to be one which could be seen as a proposition for which the evidence is so apparent that it can be seen as immediately rationally compelling and apt for being a belief that one can assert in the appropriate sense – the sense that we just 'know' it is true. But if this is the case, how can this be reasoning at all? Surely one just looks and on the immediate sensory evidence one knows. As any epistemology undergraduate can tell us though, sensory data is not so kind as to leave no room for doubt.

P1: If I can trust my senses then my direct sensory experience reveals facts about the world

P2: My direct sensory experience indicates my mother to be wearing a green dress

P3: I can trust my senses

∴ Therefore my mother is wearing a green dress

Premise three is always contestable. While the argument is valid, the contestability of the third premise means it comes up short of being sound and as such its conclusion need not be seen as rationally compelling. A fault in my eyesight, the unexpected arrival of my mother's long lost twin, bad lighting, a matching top-and-skirt combination which only appears to be a dress from afar; all of these options could explain how the third premise denies a deductive transition from sensory experience to a factual conclusion. One has to make a judgement that in this case one can truly trust the sensory data that is presented to them. Waking up in bed after seeing my mother in a green dress with no memories in between could indicate I have awoken from a dream with this content or perhaps that I fainted and was put to bed after seeing my mother in a green dress. As either option explains the circumstances, neither compels rational assent as to its truth. Again a judgement is called for.

Now take a scientific claim:

P1: Washing cotton garments in excess of 50 degrees Celsius will shrink the garment

P2: This garment is made of cotton

P3: I have washed it in excess of 50 degrees Celsius

∴The garment will be shrunk

Premise one states a scientific theory which predicts the effect of heated water of a certain temperature on cotton. While this proposition is the end of a scientific theory, there is a logical gap which stops it from being one that compels rational acceptance. One can doubt that the methods of the studies which have produced the theory are conclusive; or that the paradigm in which the theory is based is correct (“it’s an outcome of God’s anger at our using cotton for such frivolous garments and washing them at such expense of natural resources that explains it being ruined”). While these options seem unreasonable, they are not illogical. As such their reasonableness seems to hang not on certain truths but heavily engrained norms.

These simple examples show that there is a normative bent to theoretical reasoning also, and it occurs due to the fact that the evidence that we use when reasoning that it’s better to believe one thing over the next never logically entails the assertion. Those who argue that the naturalistic conception of the world as science reveals it to us contains all the ‘is’ propositions there can ever be, all the truths, have already begged the question against those who disagree with this statement. When we decide that a certain feature or certain features act as evidence for our forming a belief, we are not just assessing that the belief follows from the evidence, but that these features constitute the *right sort* of evidence for forming this belief. Thus it is appropriate to say: ‘one *ought* to believe that cotton garments shrink when washed in excess of 50 degrees Celsius’ or ‘one *ought* to believe my mother is wearing green dress today’. In making such a claim one is also evaluating the evidence on which the belief claim is made as being apt for supporting the claim. But this only elucidates just how normative even empirical belief claims can be. And if theoretical reasoning fails to be non-normative then the fact that practical reasoning contains norms more conspicuously does not pose such a threat to it as we first might have thought. It is not just the natural world that grounds such claims, but the notion that the natural world *ought* to ground these types of claims; that the way we investigate this feature of the natural world *ought* to ground this specific claim; and that the way the natural world impinges upon my senses *ought* to allow me to make this specific claim (though at other times I may be less certain in my assertion of said claim).

That we are forced, even in the theoretical model of reasoning, to make a judgement call regarding the evidence we are using, should likewise lessen the pull of the non-cognitivist assessment of what is happening when we reason practically. That the natural world does not act as a ground for our particular ends may not be such an issue. After all, neither is it a sufficient ground for our belief that our sensory data, or a scientific theory, provide ample data for making a theoretical assertion, or coming to believe that assertion. We will come back to non-cognitivist theories in later sections. (4 in particular) For now let's be content in seeing the distinction between the content of theoretical and practical claims as a lessened threat to those of us who believe that practical, and therefore moral claims, might be apt for knowledge.

1.2 The form of reasoning:

While it is clear then that our first worry for practical reasoning can be assuaged by showing there are more parallels between practical and theoretical reasoning than might have first been apparent - given that there is no simple justification in expounding either beliefs or intentions on the grounds of simple access to the natural world - there are still credible differences between them which might cause further worries. While there is room to doubt the soundness of syllogisms of the theoretical type, they still seem to be conducive to producing valid conclusions; something not readily available in the practical realm at first blush. The normativity present in a case of theoretical reasoning causes trouble for their content, but not their form. Whether or not we agree that the evidence is the correct evidence for the first premise, once we have put it into place in the syllogism then the rest follows as a matter of necessity. Deductive reasoning here seems perfectly valid as a means to arriving at a conclusion. But practical reasoning does not start with a belief; rather it starts with a desire, intention or end. This takes up the first premise. The minor premise then states something factual, which suggests a means of achieving this end, a belief, which by our analysis in 1.1 would be required to be the conclusion of a piece of theoretical reasoning or something that hits us in a flash (which indicates a normative aspect of how the agent sees the world). This form of argument aims to suggest that this is the right means to achieving that end, leading to a practical conclusion as to how one ought to proceed.

This shows how normativity appears as a double dose of doubt in practical reasoning. Not only does the second premise, the factual premise, suffer from the normativity described above which leads to possibly unsound arguments, the conclusion itself (by virtue of relating to the end or desire in the major premise) describes something patently normative; an ought statement, which appears to render it unsuitable for being a valid (and therefore necessary) conclusion. A piece of deductive logic used in theoretical reasoning can be criticised for being unsound when the 'if \rightarrow then' relationship which makes up the major premise is open for criticism; but this has no bar to the validity of the argument. The 'if \rightarrow then' still retains its necessity in the conclusion of the reasoning, relative to the reasoning, regardless of whether the first premise holds:

P1: $X \rightarrow Y$

P2: X

\therefore Y (necessarily Y given the acceptance of P1)

Whereas in practical reasoning the fact that $X \rightarrow Y$ serves as the minor premise renders the practical conclusion as not being a matter of necessity at all.

P1: I desire Y/ Y is my end

P2: $X \rightarrow Y$

\therefore I ought to X so as to achieve Y

The introduction of a desire/end strips the ' \rightarrow ' of the logical necessity it exhibited when it applied to the conclusion of the piece of theoretical reasoning. For the premises together do not necessitate that the conclusion holds. Say that ' $X \rightarrow Y$ ' represents, 'if I make this mince into burger patties and fry them **then** I can have burgers'. Theoretically this is true. But in the practical syllogism the desire Y (to have burgers) can be obtained by ordering them in, by having someone else make them for me, or I might disregard my desire to have burgers (perhaps due to my growing gut). There is nothing necessary here in the conclusion of the practical case (I ought to X to achieve Y) as there is in the theoretical case (if X then Y). In this regard, while the reasoning seems an accurate representation of a thought process we might go through, it is not deductively valid.

This is essentially Anscombe's (1989) complaint against von Wright's (1983) attempt to create valid practical inferences by making the means necessary. Anscombe complains, quite rightly, that necessity goes hand in hand with validity and, as shown above, we lose our claim to validity by adding an intention into the process of deduction. Rather, whether theoretical or practical, the *aim* of the reasoning stands outside of the reasoning itself. Take the following two statements:

'Flying from Manchester to Florence is the only way to arrive in Florence by tomorrow'

'Being in Florence tomorrow is the only way to see this year's Florentine New Year celebrations'

∴ As I am in Manchester I need to book a flight if I am to see the Florentine New Year's celebrations

This piece of valid deduction can be used for various aims. In the practical sense it can guide us in which transport to take, and in the theoretical sense it can be used to predict the transport someone might take whom we know wishes to go to Florence only because of the Florentine New Year celebrations. According to Anscombe, whether one is conducting a piece of practical or theoretical reasoning, one is starting from a position in which one has an end – to predict an effect, to explain a cause, to intend to achieve that end, or to intend to achieve its opposite. The former two are theoretical versions of reasoning and will thus end in beliefs which have what have come to be called a 'world to mind fit'.

The latter two are practical versions of reasoning and thus end in intentions which have a 'mind to world fit'. Crucially, however, in both cases, the reason for which one wishes to utilise a deductive chain such as that above stands logically distinct from the chain itself. The chain will be the evidence that is utilised to make the assertion, be it practical or theoretical; and the decision to utilise the chain in either case will be a judgement call that this chain gives us reason to suppose the assertion is true. Given the normativity that is present in both of these cases that were canvassed in 1.1 we already have strong grounds for assuming that there is nothing that beliefs have which intentions do not (or vice versa) which precludes the latter from being seen, out of court, as somehow lesser than the former.

Without an analysis which shows how normativity is also present in our most obvious theoretical attempts at reasoning towards knowledge there is a temptation to believe, however, that beliefs are more readily open to being simple cases of knowledge, underpinned as they are by something that exists in a seemingly observer independent manner. We have already shown that even assertions which end in beliefs that are guided by our most obvious means of access to observer dependent data - namely our senses and our paradigms of scientific verification via the postulation and testing of theories - make use of normative criteria about the evidence used being the 'best evidence' for making these epistemic claims, and that these norms cannot themselves be grounded by the natural world alone. From this, it seems to follow that deduction, being relevant to reasoning only once one has made a series of normative judgements, cannot tell the whole tale of reasoning towards knowledge. We need more focus upon the normative nature of discernment about what constitutes evidence which leads us into the deductive process if we are to better understand either type of reasoning. I will conclude here that despite the seemingly vast differences between practical and theoretical reasoning with which we began, both types of reasoning have very similar issues. Both show normativity in the content of the claims and neither can fall back upon a form in which knowledge is a simple case of attaining the right sort of evidence and applying deduction as a way of attaining outputs which are rationally compelling as indisputable cases of knowledge; any piece of deduction stands logically distinct from the purpose it is being used for. What remains true of the differences between them is that theoretical reasoning leads to beliefs about the way the world is, while practical reasoning ends in an attempt to align the world with how one wishes it to be, by forming intentions to act. This relates to issues of psychology between intention and belief formation, the former requiring some level of motivational content. We will discuss motivational issues in later sections. (**3.5**, **4.2.2** and **6.3** in particular)

That the two types of reasoning turn out to be as interlinked as they are should, perhaps, be of little surprise. When one embarks upon a piece of theoretical reasoning one does so for a purpose; more often than not for it to be utilised as the minor premise in a piece of practical reasoning as that is usually construed in syllogistic form. We can always be seen as aiming at knowledge for a reason, and

that reason is more often than not a practical aim. Although even if this is not so, we can assume that 'finding out what is the case' has some normative value; even if one does so simply for the sake of furthering one's knowledge in an area that may be of no more practical purport than just 'knowing more', one exhibits a normative evaluation that such knowledge is of interest, that they value knowing it. Likewise, the very idea that practical reasoning is of philosophical interest is in some very real sense theoretical. It is under the guise of wondering about the objectivity of values that we come to wonder about whether *this* practical decision is the right decision, or whether *any* practical decision has any claim to being referred to as the right decision, and if so, how. Answering either way to either case constitutes one adopting a theoretical stance on the matter.

It will be the aim of the coming sections to entertain theories which have attempted to show that moral claims, a species of practical claim, really can be shown to be either right or wrong in their prescriptions. I will utilise theories which are commonplace in normative medical ethics text books which attempt to make the 'ought' which acts as the conclusion for a piece of practical reasoning binding, not just to those with a given end, relative to a particular agent, but binding upon *all* agents. Despite the fact that we have shown norms are present in both cases, the touchstone or grounds that the real world gives to theoretical belief claims are always thought to be assessable as right or wrong; if not now, then at some future time when our ability to agree upon and gather evidence is so compelling as to require one to make a huge leap to be able to deny it. Who is making the claim is irrelevant to the correctness of that claim; whereas who is making the practical claim seems to be absolutely crucial to the practical claim; the conclusion requires reference to the intentions of an individual agent.

As we will see, this universal ought is hard to pin down. Given the above analysis of the logical gap between evidence and assertion we already have strong reason to believe that this will indeed be the case. Attempting to fill that gap with a theory seems futile given that accepting the theory will count as making a judgement that it is theory X and not theory Y which should be followed when making practical decisions. The desire to formalise and universalise our judgements in this manner seems to be an attempt to align our reasoning towards *prescriptions* with the now discredited and distorted version of our reasoning towards *descriptions*; one which,

given the normative bent of theoretical reasoning, seems to be a pre-philosophical illusion. But the sheer abundance of available differing moral positions towards any subject matter seems to require something like the postulation of a standard by which we measure them, while at the same time the postulation of a standard seems to be just as 'intuitive' as the postulation that any particular judgement is right or wrong. As Loughlin (2002: 56) puts it, this is a 'methodological flaw' with the whole debate: either our 'intuitions are sound' and a theory is not needed, or our intuitions are unsound and we have thereby called into question the only ground we might have for judging any postulated theory.

However, going over this well-worn terrain will give us insights into several positions and more importantly the assumptions which underpin them. After seeing why each fails to attain rational assent we will turn to the intuitions which seem the crux of the matter in both the theoretical and practical cases of reasoning to see if we can rescue the idea of having knowledge in either area.

Section 2: Generalist Accounts

This section and the next will address three separate attempts at giving moral knowledge the kind of ground we wish it to have if we are to keep the sense of 'ought' that seems implicit in our moral assessments. In purely practical reasoning we have a sense that there is seemingly some relativity as to how one ought to rationally proceed. If one desires to get to the destination quickly one will fly. If one prefers the anticipation of the journey and the idea of seeing the changing landscape leading up to their destination one might take a ferry or drive. Once the end is decided upon, instrumental reasoning can easily provide something closer to an objective account of realising that end with the appropriate means, but the end itself seems something relevant only to the person deciding.

In the language of moral assessments, however, no such relativity seems to pertain in the way we make our judgements. As MacIntyre (1981: 11-35) points out in his assessment of emotivism, the emotivist idea that when I make the claim 'A is morally wrong', what I really *mean* by this is actually 'I do not like A, and nor should you' or 'boo to A', seems utterly misguided. When people make moral statements what they profess to mean is that by some *universal standard* A is *definitely* wrong. There is no relativity implicit here. The problem occurs in the fact that no standard has (or by way of the analysis in 1.1 & 1.2, *could*) elicit universal acceptance. MacIntyre believes that this puts us in an awkward position by way of which emotivism may well be the best theory to describe our current moral position, though not as a 'theory of meaning' but as a 'theory of use'. While we mean to evoke a standard by which our moral prescriptions are measured, all we are really doing is appealing to our own emotional responses to certain states of affairs; responses which have no ground in reason or reasoning.

From this assessment it follows that moral claims are also relative in regards to our individual psychological tendency to find some states of affairs or actions to be conducive to feelings of approbation and others to feelings of disapprobation (though it is likely that MacIntyre would dislike the use of such Humean expressions in an analysis of his own work). It is these psychological concerns that guided the theorists who created, modified and defended the theories we will now canvass. If our moral

claims are to have any authority, there *must* be some ground upon which they sit; an 'epistemic hook' which warrants our saying some claims are right and others not so. Without this, the threat of non-cognitivist theories looks problematic.

As well as the need for an epistemic hook, the realist will need to find a way to overcome the Humean driven idea of direction of fit, which so nicely fits into both fact value and the belief attitude distinctions. These points relate to psychology. Given that these suggestions for grounding morality are theoretical, I will also need to make methodological points about each theory.

Three theoretical attempts to ground morality will now follow. The first, which will make up the bulk of this section, will be deontological. I will focus mostly on Kant but closely related theorists will also be looked at. The second of the three theories (starting at **3.1**) will be Utilitarianism. Looking at this type of theory will allow for analysis regarding the form of arguments and commensurability. The third type of theory (starting at **3.4.2**) to attempt to ground moral claims will differ in two ways. Firstly, it is not as action guiding or principle based as the first two, and secondly it attempts to ground claims via a naturalism that often underscores virtue accounts, rather than the rationalism that underscores the principled accounts. In the epistemological and methodological arguments against these three positions we will have time to assess other theorists in passing as well.

2.1 Deontological Principles and Autonomy Based Accounts:

The first attempt at grounded moral reasons that I will analyse are Deontological principles. There are various ways in which theorists have pictured these as coming to guide us in our decision making processes. The most famous defender of such a position was Kant, however, and it is he more than most others who truly attempted to give these principles a rational hold on us that could not be simply argued away by complaining that the principle being evoked was simply not shared by a disagreeing party. Since his attempts there have been various other forms, some of which are highly constructivist Korsgaard (1996) and some of which are aligned more strongly with a virtue account (O'Neill, 1996). I will have time to

mention the constructivist position in passing and directly look at a challenge set by O'Neill towards the final sections of this thesis (see 7.4)

Kant attempts to duck the charges that we laid out above, regarding practical reasoning necessarily embodying the normative features of deciding whether the evidence we take to be good evidence for a practical decision is the right sort of evidence, by postulating, as Dancy puts it, 'a bifurcated conception of practical rationality' (Dancy, p. 68, 2004). For Kant, in most cases of practical rationality, one simply follows the given instrumental model of using reason to find a means to one's ends. Crucially, however, in the moral case this is not good enough. When one has a moral decision to make, one must test their intention against the Categorical Imperative which is given by 'pure practical rationality'. Before we explore what the Categorical Imperative states, let us first see why Kant determines pure practical rationality as necessary for making the right moral choice, and thus sees the bifurcation of practical rationality into two distinct realms - the moral and the practical - as being necessary to understanding the distinct claims that the merely 'practical ought' and 'moral ought' make.

For Kant, his theory is a third way to securing the idea that moral judgements are categorically right or wrong. He believed that it was needed because all the others failed to capture the importance of practical reasoning to moral theory (1948: 441-444). The two types of theory he mainly argued against both fall into scepticism by his strictly rationalist lights. He denies that a sort of intuitive naive realism such as that proffered by Leibniz can give us categorical 'oughts' and thus properly ground morality. If there are simply a set of facts like empirical facts which exist anyway and are recognised by us then they have no more authority over us than a sign that reads, keep off the grass (Dancy, 2004). Knowing that this is a law does not compel action, for this we would need another element, namely a desire to follow the law. For Kant the only way that moral reasons can have moral authority over us is if those reasons are self-legislated and vindicated by reason (in the form of pure practical reason) as being *our* laws to follow. For this, no such desire is needed. For Kant, a law to keep off the grass is only rationally compelling and categorical if it is a self-legislated law. In the same breath, he disregards Hume and Smith's 'Sentimentalism' on the grounds that they openly deny the idea of reason generating laws or critiquing anything other than finding means to ends. Those ends are already set by one's

anthropological make-up, learning the practices of a society, enculturation into moral practices and norms, psychological constitution and various other things that might be seen as establishing in us a second and acquired nature. Such concerns - our wants, likes, values, passions and so forth - are, for Kant, only *contingent* matters and thus cannot be the ground for a definitive moral reason. As we will see then, the Kantian notion of certainty about moral claims will necessitate us accepting a very different conception of rationality than that which we normally utilise in practical judgements.

Having abandoned quite openly all the substantial data human beings do have when making moral decisions as being contingent - including the ends that we happen to have - Kant claims that the only way of coming to know our decision is right is if it is one that is suggested to us by *pure practical reason*. It is practical, in that it suggests an 'ought', but also pure, in that it suggests it is arrived at *prior* to any particular knowledge we may have about our own situation. This *a priorism* (Kant, 1948) that Kant adheres to vehemently, is a crucial and distinctive claim of his account which sets it apart. Any theory which lacks it has the issue of proving just how one is stating a properly universal principle and not merely their contingent biases. If one is to adopt a principle that is truly rational, one must will it freely and not be constrained by exterior alien forces, as Kant conceives those to be. If these exterior forces act as a cause for our choice we have not, according to Kant, chosen freely. An otherwise quite reasonable desire or end, such as that to not see family members harmed, would be contingent for Kant in this regard. One might will this end in the *ordinary* sense, but the cause for this willing would be set upon contingent and conditional criteria. To truly will something *universally* that is not contingent (*Wille*) means that thing be freely chosen and thus truly proposed by rationality itself, prior to any of our particular inclinations or concerns (Kant, 1948: 398). Thus, the principle one wills must be self-legislated upon grounds that pure practical reason (and not an everyday notion of reason) can endorse for itself as being a law that makes no reference to the agent's contingent circumstances, but to the act the agent is intending to commit.

From these assumptions Kant creates the first formulation of the categorical imperative:

Act only according to that *maxim* by which you can at the same time *will* that it should become a universal *law*. (402) (my emphasis)

Without any substantive content to supply his principle, a principle which will act as a test for whether an intended action is morally right or wrong, Kant's formula has only the added highlighted features to work from. Beset by the ideas that: no other way could truly mean we are acting rationally, by way of his own definition and commitment to his idealism; that morality requires pure rationality to have any authority; and that this authority must be expressed in laws that are relevant to all rational agents, this is what he is left with. '... [either] that or think of a simpler or better formulation!' as Wiggins decisively puts it. (Wiggins, 2006: 92)

The first thing to note, however, is that Kant's first formulation seems to intuitively fail to secure whether an action is right or wrong. It would seem that absolutely anything can be willed into being a universal law on the basis of the Categorical Imperative as it stands here. If we are a strong man, or a nihilistic man, we might have no qualms about willing into existence the law 'let all disagreements, no matter how minor, be resolved by a duel to the death'. Or, if we are a selfish man, willing, 'let everyone have freedom to their own possessions, except when I (and only I) demand use of them' (for other examples see MacIntyre, 1981: 46). There seems to be no rational bar to these sorts of flagrant misuses of the Categorical Imperative being used with impunity. One might note that the fact we see these as 'flagrant misuses' shows how dependent upon our ongoing moral intuitions an analysis of any postulated moral theory must be. Is it not odd that Kant claims to disown all such inclinations? Kant's setting of the moral law is specifically meant to not invoke laws imposed from outside, but to instead allow us to create our own laws, as this is what makes us what we are as rational, autonomous and moral beings. But this feature also seems to allow for us to will the most intuitively immoral things if we wish. Indeed, it seems to allow us to will anything; and a rule which allows for everything cannot hope to guide action. Kant attempts to avoid this charge by invoking the idea that we cannot truly *rationally* will just anything without contradictions. If he can show this to be the case then Kant may well have given us food for thought.

Kant gives us four examples (1948: 422-424) from which to work from. Kant's examples are: of a man considering suicide; of one considering whether to borrow

money despite doubting that he could pay back the debt; of one considering whether to neglect his natural talents; and of one who believes that, while he should not negatively affect another man's standing, there should be no sense of duty to aid his welfare were he to ask for it. Kant's problem here is to show that these examples cannot be rationally willed and to do so without helping himself to any teleological, anthropological, or in any other way contingent notions about human moral life or nature. To do so would be to subvert the *a priori* framework which he has so vehemently claimed is the only framework which can truly justify moral claims.

Sadly, for Kant, his examples all make use of either teleological or anthropological notions. In the suicide case he refers us to the contradiction of willing a 'system of nature whose law would be to destroy life by the very feeling whose *special office* is to impel the improvement of life'. This teleological characterisation of natural elements having end goals, such as to 'impel the improvement of life' obviously undermines the strict *a priori* account he had so far adhered to. Likewise, the other three examples have similar issues. The idea that borrowing without being able to pay back would undermine the entirety of the financial system seems a truth unlikely to be proven *a priori*. The example regarding the man neglecting his talents, again, requires a teleological notion to not become a logical contradiction. 'For as a rational being he necessarily wills that all his faculties should be developed...and serve him for all sorts of purposes'. This additional teleological notion of what the rational ends of man might be is again unsuitable to the *a priori* framework. Finally, the man who refuses to be benevolent to the other in need is derided by Kant for a contradiction of the will. One cannot, Kant believes, will that it become a universal law that people do not help others in need, as 'instances arise in which he would need the love and sympathy of others' (Kant, 1948). Again, Kant's insistence on a completely pure *a priorism* works against him. While we can reasonably say from an *a posteriori* position - one that might happily take advantage of all the local knowledge we have about human needs and human lives - that the majority of persons do indeed rely upon the benevolence of others, such a point is still a contingent fact about *human* nature and not a truth discoverable *a priori*, one which could be applied to *all possible rational beings*, as Kant needs it to be for his Categorical Imperative to work. Indeed, hermits who eschew contact with others from a personal preference for solitude are testament to this.

But Kant later does supply those testing their intentions by way of the Categorical Imperative with some serious substantive content by way of his second formulation. For Kant, 'rational nature exists as an end in itself' (ibid: 420). This is not to say that we *must* always act so as to become more rational, but that we can never rationally act in such a way that harms our rational nature. Likewise each rational agent sees his rational nature in this way: he necessarily acts (if he is acting rationally at all) in such a way that sees his rational nature as an end in itself. He might forgo safety or dignity or health for other ends he sees as worthy, but he cannot ever rationally will that he becomes merely a means to the end of another will. What we have reverence for in ourselves and other persons is not a kinship by way of being members of the same species, but of sharing this unique feature of rationality, of being able to legislate laws freely by way of our *Wille*, that rises above the animal reasoning of just being pressed by one's contingent desires into wanting certain things and reasoning about the means to get them. It is in light of this feature that has 'absolute worth' and is 'an end in itself' that we must likewise never treat other rational beings as means to our own ends. He calls this feature that we must respect 'humanity', or the capacity to rationally will laws for oneself on the basis of pure practical reason.

However, by way of empirical evidence we know that rational beings do indeed act so as to constrain and manipulate the wills of others by lying, cheating, stealing, killing and so on. They choose to act not from formulations of the Categorical Imperative, which acts as a means to understanding what the moral law demands of us prior to our personal circumstances, but from hypothetical imperatives which state conditional pieces of reasoning such as, 'to achieve Y do X', where the end (Y) indicates only reasoning which *fully* takes into account the particular conditions and personal concerns, ends, wishes and desires of the agent doing the reasoning. Kant does not see this as a worry similar to that with which he charges intuitive realists such as Leibniz (ibid: 443); the worry that one requires an extra desire to be moral, a desire to follow the moral law, alongside knowing the moral law. When we test an intended action by way of the Categorical Imperative, we are not just finding out what the moral law says we must do, as though we had just discovered a list of commandments that were in no way ours. Instead, we are ourselves making our own laws; and to fail to live by one's own laws is to give up one's freedom as an autonomous agent. One needs not a desire to be good to accord their actions to the

moral law; rather, not doing so shows that one has abandoned the prescriptions of their own *Wille* and is thus acting irrationally.

2.1.1: Critiquing Kant:

What becomes apparent in this summation is the notion that we started with: the notion of a bifurcation of practical rationality. In making this first step, Kant makes various claims which seem to lead to his conclusions in ways that will appear question begging to those who disagree with him. In the first sense, the bifurcation of practical reason leaves us with two sets of methodologies which apply to two distinct sets of imperatives. In normal practical rationality we create laws for ourselves by willing the means to our antecedently willed ends, which give us hypothetical imperatives to act. In pure practical rationality we *Wille* in a way which acts only in cahoots with the conception of law as determined by pure practical reason. Having rejected any ends as worthy for us to pursue as distinctively moral agents, given that any such ends would themselves be contingent upon anthropological experience, Kant gives up on the ability of claiming this or that ends as being worthy direct ends for moral behaviour. All we are left with is the notion of the moral law with which to judge individual actions. This strategy alone, forced upon Kant by his rejection of purportedly contingent beliefs of what is good, forces him to reject the evaluation of moral ends as being worthy or not and find a separate means of assessing behaviour. To do so he has to abandon our normal understanding of practical reasoning ('what ends should I have and what means are suitable to obtaining such ends?') and focus instead upon the idea of the moral law, which is only understandable outside of the normal scope of practical reasoning, which he takes to be a distinct methodology of finding an imperative by which we can ascertain the moral law *a priori*.

This in turn gives Kant both a license and a need to adopt the strategy of separating the 'intelligible' and the 'sensible' selves. The sensible-self is subject to all the desires regarding the contingencies of the world we live in. But the intelligible-self steps outside of these determined causes, and, by means of rationality, can adopt laws which condemn or condone these desires as irrational or rationally necessary. Reason suggests something to us, say, the desire to feed our family by means of

stealing, and pure practical reason either ratifies this thing that we will as being something it is moral or immoral to do by assessing it in terms of the categorical imperative. It is only in this sense of the two selves, and of the two realms which each self-inhabits, that the Kantian notion - that freedom and morality are really one and the same thing (defined as supplying the laws of reason to oneself on the grounds of reason and nothing else) - is not viciously circular. Otherwise we can be free only if we are moral, but we are moral only if we are free; each vindicates the other. Given the two realms, however, self-legislating is to be moral in the intelligible world, to see others with the same rational will as equal, while it also insures our freedom from the sensible world. In this way the very notion of willing is also bifurcated. I might sensibly 'will' my family not starve, but I cannot intelligibly 'Wille' it without contradiction if that would guide me to steal to feed them.

Those who intuitively feel that the types of laws which the Categorical Imperative produces are too stringent may well come to question the bifurcation all the way down. Is it not the ever-so-early assumptions in Kantian morality that lead to us separating the two types of practical reasoning; begging the question somewhat against others who do not believe that the judgements they come to when making a moral decision need be laid out in such in such law-like terms which focus never on ends but only on actions? Can it really be the case that we can be perfectly moral without an understanding of what kinds of ends are suitable for us as moral beings?

The only way to see this as being so is if one agrees from the outset, with Kant, that there is no room for rational action if we start from an agglomeration of anthropological data - such as feelings and wants about our particular circumstances (as contingent as they might be upon those circumstances) - and practical reason, described as a faculty which discerns how to either take the means to realising these wants, or to disregard them as unworthy of us or too complicated to achieve and so on. For Kant, from the very outset, to accept this framework is to accept something which will deny the claim of the moral ought as being objective. It will be subjective all the way down as our ends will not be provided by pure practical reasoning but what he conceives of as 'alien' (Kant, 1948: 404). But these are only alien to us *if* we are already inclined to accept his separation of two realms: one in which freedom is possible (the intelligible) and one in which it simply isn't (the sensible) and to assign our 'real' selves to the former. If we accept this notion - that our practices,

inclinations, experience of what constitutes good, bad, permissible and impermissible behaviour constitutes an alien force - then perhaps we will want to bifurcate practical rationality so as to supply moral judgements with a sense of will (Wille) quite distinct from that used in other practical judgements. As Wiggins points out however, there is something destructively creative in the way that Kant defines “Wille as the causality of living beings in *so far as they are rational*” (2006: 126). The italicised words beg the question against those who believe that rationality requires no such thing as a faculty which supplies laws from a standpoint metaphysically distinct from the seemingly natural standpoint in which we practice morality as part of our specific form of life. It seems to beg the question against other theories to rule out of court that from this more ‘normalised’ stand point, one which does not separate realms of being and existence, we have no reason to disavow the actions of an immoral man and our actions cannot be rational. Either we accept the bifurcation or we must reject the *a priori* structure that Kant attempts to instil. As I have argued, the former possibility is mired with problems, the most evident being accepting what I can only see as a philosophical metaphysical fiction.

Most of this analysis is at the metaphysical level, but at the normative level alone, as mentioned, the principles that Kant’s theory produces cause problems. The most famous of these is the idea of clashing principles. Judgement still has a great deal of work to do when we come across two such principles. This is of course an infamous problem for those who wish to retain the sanctity of a principled morality, even if they are not Kantians at the metaphysical level. There is an obvious pull towards being able to have moral norms that hold universally. On a purely epistemological level it makes the intelligibility of the domain instantly understandable. There are actions that we are just expected to never or always take part in. These actions are then laid down in principles which act as an arsenal for the moral man. As we approach a situation, we test the features of the situation to see which fit and when a principle fits it simply must be heeded as having universal authority over us.

As we saw with the Kantian analysis however, how we discover such principles is of key concern. Kant believed pure practical reasoning could uncover an imperative that would act as an universalisation test for any candidate maxim for action. Once it had passed the test the candidate became a working universal principle. Part of the normative issue for Kantian philosophy is how distinct some of its ideas are from

everyday morality as we seem to practice it. Other concerns than that of a sense of duty to law seem to pervade our moral reasoning. Yes, we have a duty to never act in certain ways (just think of the most heinous crime a man might commit and add 'one must never...' as a prefix to it for clarification of this point), but we also seem to have further duties. The notion raised by Constant (Korsgaard, 1986: 325-349) of an axe murderer asking for the location of his intended victim while you shelter him in your home from danger brings this to the fore.

2.2.: Ross: Perception and Principles:

Kant's rigor for the moral law, as he decreed it, blocks his ready acceptance of ideas based on everyday moral notions such as beneficence, care, consequence and the rest. This objection was raised by W.D. Ross and was the cornerstone of his positive account. In such a situation, claims Ross, what we have is a clash of principles. In coming to choose which principle to act upon a man shows his moral worth. There is no unresolved clash of universal and necessary principles in Ross's account because each of his principles (he names seven but admits this account need not be complete) is *prima facie*. By this he means that they are not absolute duties, as these cannot be discovered by us through *a priorism* or any other means. For Ross, his principles are best described as responsibilities which have certain features which should impel us to action when they fit the particular situation an agent is in when deciding what to do. I prefer the term *pro tanto*, following (Kagan, 1989: 17)

Pro tanto principles, then, act as normal principles except on occasions where one or more principles are in play and state different possible actions for the agent. In this case, one must take precedence over the other. Ross's view has the advantage of thus making our appeal to principles more in line with our intuitions; something he thought was a positive aspect of his theory (Ross, 1930). There is not, by these lights, one single source of morality, such as the moral law as Kant sees it, but multiple ones, often with varying and distinct concerns regarding the right and the good. In the case of the axe murderer at our door then, Ross would be able to simply answer that two norms apply, the norm to be honest and the norm to preserve life. In this case the former is ousted by the latter; but on what grounds? Well, Ross

believed there was no real ground for these principles except the sort of ‘...moral reflection of many generations, which has developed an extremely delicate power of appreciation of moral distinctions.’ (Ross, 1930: 41) This sort of knowledge of choosing between general concerns can only be, in Ross’s eyes, a sort of ‘perception’ which we develop by being actively involved in the moral practices we are necessarily involved in. Given the complexity of this sort of knowledge that Ross expounds as so central to morality in comparison to that which fits the broad concerns of principles and concerns he did state - such as fidelity, non-maleficence, reparation (Ross, 1930: 21-25) and so forth which are simple enough to write down and recognise as and when they occur - it seems that most of the work in Ross’s account is being done by the intuitive part of his theory. Principles, in Ross’s case, seem to be doing very little work, except picking out major concerns that the moral man who can choose between the principles in the correct manner, might already have. Richardson points out (Richardson, p. 42: 1994) that practical (and therefore, moral) dilemmas ‘may arise from a single principle’, such as that of promising. In deciding in either case, what one ought to do ‘all things told’, it is more likely to be the intuitive knowledge and not knowledge of principles doing all the work.

This being the case, we have something of an issue regarding deontological principles which have no basis in a Kantian sort of *a priori* - which I have argued is too fantastical in its transcendentalism to seriously proffer as a valid account of how we should decide what to do. In the non-*a priori* case, we select a principle for ourselves which, if it is to guide our actions in any meaningful sense, must hold on each occasion in which it is relevant: For all circumstances, C, do action A. But it is clear that the world being as it is, and that the principles themselves are set by our own judgements, that there will be cases where the same ground which suggested setting or accepting a given principle, namely our judgements, will wish to override this same principle as it seems to not suggest the right sort of action. ‘When my students are disrupting the lesson plan by talking loudly in class it is best to first give them a withering look to make them remember they are disrespecting my place as teacher’. But now imagine the discussion is academically useful even though I, as the teacher, did not start it. Perhaps here I will let them talk and disregard my principle. Those attuned to the idea of the consistency that principles afford us would claim that, if this judgement is not to be a subjective matter, what we must be

appealing to here is another principle, perhaps, 'do not interrupt my students when they are making academic headway', which has superseded the first of my teaching-based principles. I have, *a la* Ross, used my intuitive judgment to pick the one over the other. Ross's theory, we have already decided though, does not itself seem to explain our moral and practical reasoning in a principled manner, as it simply imports the faculty of intuition into the mix. There will be further commentary on intuitive judgements later. In the above example a further way of explaining what is happening in deontological terms is by way of *expanding the principle*.

By the 'expansionist' method, a term I take from Dancy (2004: 11), rather than intuiting that I should choose one of my principles over the next, what moral principles come down to is actually specifying the true content of each principle in more detail. According to expansionists, the educational non-moral case, which is represented as having two conflicting principles under Ross's analysis, really has but one principle: 'Do not allow students to interrupt my lesson plan *unless* they are actively advancing their academic knowledge by doing so'. The benefit of this analysis over Ross's, that those such as Scanlon who employ it enjoy (*ibid*), is that the Intuitionism inherent in Ross's account is removed. One is never picking between clashing principles in such cases, but rather giving an instance of a single principle which is specified according to the needs of the interlocutor with whom they are conversing. The problem is the myriad of variables which can affect any situation seem nearly infinite. Perhaps we would still agree with the more fully specified principle unless we feel that the students took over in a highly disrespectful way, which would add a further specification. These concerns go on almost infinitely given the complex nature of relations between the tutor, the student, and the material, incorporating all the expectations, power relations and other factors relevant to the learning situation. A fully specified principle according to Scanlon's account would thus be a nearly infinite thing with equally infinite 'unless' clauses which took one's decision on how to judge the situation between behaviour that was or was not to be tolerated. The idea that such a principle could underwrite action seems plainly false. One could not carry around this principle for use in deliberation itself, as one could not surely have, in front of one's mind, a principle with infinite clauses, or set of principles with some pertaining to differing contexts, against which one measured behaviour.

Such a way of discerning what is right and wrong seems to fall back on the same sort of intuitionism that is there in Ross's account, just less obviously so. Rather than choosing between principles, one is deciding whether an 'unless clause' is valid in deciding what one's overall judgement of the situation will be. When one thus makes that call it cannot be the principle itself that is used as a ground for deciding; instead it is one's decision which shows how one would go around specifying the principle if one's interlocutor were to ask. As there is no way to expound the principle prior to the deliberating process it seems it plays no such role in underpinning the deliberation and thus has no epistemological authority over us. It is more a process of self-discovery; discovering at which point the 'unless clauses' seem to run out for us, or of at which point a new variable does not override the last.

As we cannot link every possible variable to every other in such a way as to pre-deliberatively know which combinations will result in which sort of action (or reaction) from us, it seems that each 'unless clause' has the same sort of judgement in it as does Ross's principles, of which there may be an infinite amount once we see the same principle can lead to being forced to judge between two accounts of it and given the various cases in which such individual instances can clash seem infinite in and of themselves. If we imagine the way in which two promises might force a Rossian into an intuitive call, the deciding of which one to honour may appear to contain the same sort of judgement inherent to it as the 'unless clauses' in keeping a promise where another promise is involved as a reason not to uphold the first, in Scanlon's account. Rather than choosing between two principles intuitively, Scanlon's account would have us believe that there was only one principle which needed further specification. It might run: 'one must always keep one's promise, unless another promise is involved, which should then be honoured unless the first has greater necessity, in which case it should be honoured...(and so on, to an indefinite length based on the possible variables)'. One might well say that the sort of concerns expressed in the 'unless clauses' are what a Rossian might be bothered with in making his 'intuitive' call, except that Ross sees them as not apt for being a part of the principle itself. In this I would agree with Ross, as it seems crucial for a principle based account that the principles involved count towards deliberation, not obtained post-deliberation. The parallel issue in Ross's account is that although the

principle makes some import to the deliberative process, when one needs to choose between them they are simply not relevant to deliberation.

It seems, then, from the three attempts to ground moral thinking upon non-consequentialist principles, that such principles do not afford us the universal ought we are looking for if we are to attain the objectivity implicit in the language of our moral assessments that MacIntyre alluded to. The principles we have either match the phenomenology of our thinking more clearly, as Ross's do, whilst failing to provide anything close to a universal ought that is based upon the reasoned part of the account, or they do manage to attain a universal ought, as in Kant's case, but only by supplying a rendering of practical reasoning that is distinct from our usual phenomenology and thus remains suspect. Still, what we have here is our first hallmark of the orthodox account of moral theory: it is principled and must be principled in order to achieve objectivity (principles should be universal), epistemic validity (principles must have an evident ground) and rational (principles allow for discursiveness). What eludes each of these accounts is just what the ground of such principles is to be if they are to fit each of these criteria. Kantian *a priorism* came closest to giving a satisfying answer to such a question but failed to convince us. Whilst Ross's account seems to have a more sensible account of the seeming flexibility of principles, it seems we require intuition more than we would like without further analysis on this issue. I will come back to intuition later in the thesis.

Such an analysis leaves us in a tricky position regarding our search for universal principles to underpin moral claims. Any principle also seems to come with possible exceptions built into it, unless we take the Kantian route of *a priorism*, which rejects out of hand any such exceptions at the cost of a sensible account of practical reasoning for arriving at such principles. Such accounts have, as it were, a framework that would allow for a fully objective moral account based on principles, one which will remain empty until we can find exactly what those principles might be. But this always seems a matter for dispute on the normative level without any possibility for resolution in a way that commands rational acceptance. As such, choosing any theory over the next seems to be as intuitive a matter as deciding between competing actions in an individual scenario.

In the domain of medical ethics this is a serious issue; what we are left with is the ability to find principles which we all agree upon but in so doing we are not actualising the purpose of those principles, which is to *guide action*. Principles which are proposed thus could guide the action of people who disagree if only we could find suitable grounds for them which would compel acceptance, or they would already be accepted but would do no work in guiding action for this very reason and are thus impotently tautological. As deontological principles are duty based, and rights are derivable from duties, this analysis puts the entire framework of moral rights into the above quandary. There is a sense in which we like to think that having a right to something protects us. But if someone disagrees that they have a duty to uphold that right and we have no recourse to explaining just why it is our right in some universal sense, it seems the addition of this linguistic nuance does nothing at all to protect us. Rather, as principles seem to be more intuitive than this in their application, you have a right to be told the truth by your doctor only if your doctor does not think that his duty to protect your welfare is more important than this, and that in this case the truth may well be harmful, irrelevant of how much you might disagree. If you wish to disagree by arguing that his ordering of principles is incorrect then the appeal to principles simply breaks down; we are left once again looking for a serviceable hook.

Section 3: Commensurability

3.1: Utilitarian and Maximisation accounts:

The previous section covered deontological principles, noting that the best grounds for such principles were heavily rationalist. Those accounts, such as Ross's, which did not appeal to *a priori* accounts of practical rational thought seemed to rely upon the idea of the principles in principled accounts being *suis generis*. This, in turn, relied heavily on the postulation of an intuitive perception in order to choose between clashing principles and even in discerning what the right principles might be to carry around with us and when they were appropriate for use. Both these attempts to ground principles failed in various ways to provide what we were looking for: an epistemic hook upon which to hang our moral judgements to provide them with a level of clarity and certainty which would align them with descriptive judgements.

We have already seen that descriptive judgements themselves rely upon some level of normativity (1.1 & 1.2); however, the distinction remains valid to a certain degree in various ways. As mentioned previously, one of these is a matter of the difference in psychology between the formation of belief and intention. Secondly, the norms upon which we choose to apply descriptive knowledge seem to have smaller variation and are more widely held than those regarding evaluative knowledge. It takes philosophical analysis to see that the epistemology of the descriptive is normative in a way that is not necessary for the epistemology of evaluations. The former has a hook, which is the natural world itself, as revealed to us by science and our sense data. "When I see a bird that walks like a duck and swims like a duck and quacks like a duck, I call that bird a duck", said James Whitcomb Riley. Gettier (1963: 121-123) examples and evil demons aside, we have solid reason to follow this advice on the descriptive side more than we do if we transformed it into an evaluative moral version, such as: When I see an act that appears in every way to be morally repugnant I call that act morally repugnant. In the latter case everything comes down to the way in which normativity is so apparent to appear instantly and readily open for critique. And what, we may ask, are the defining features of moral repugnancy *for you*?

This small diversion is meant to draw us back to the idea that if there is moral knowledge then there must be some set of moral facts to which they relate, some content which would make a case of recognising the features in an everyday setting as an instance of knowledge, in the same way a beak, webbed feet, a waddling walk, feather patterns, size, the tendency to quack and the abilities to swim and fly all act as features for us to correctly apply the term duck to a duck. One way such knowledge could exist would be if there were natural features of the world, easily recognisable, to which terms such as good, bad, right and wrong related. Then seeing how such natural features were affected by certain acts could determine just how good, bad, right and wrong those acts were. What would be needed is a solid argument that such a descriptive feature or features truly did ground these evaluative terms. Such an attempt would give moral knowledge a hook into the natural world. This section and the next will evaluate two ways in which theorists have attempted to give us such a hook for our moral judgements. This section will consider Utilitarian theories which identify right actions as those which maximise happiness. This analysis has the benefit of giving us an insight also into a popular idea regarding the way we make practical judgements in terms of an ends/means reasoning as well as the important notion of commensurability. While not all Utilitarian theories incorporate a naturalistic element, some will.

As with all moral theories, those which employ the notion of maximisation of a single good start with the notion of just how we should go about deliberating in what we ought to do. In most cases we seem to follow basic axioms for deciding, be it avoiding harm, promoting autonomy, being kind, generous or charitable and so on. As we have seen, however, in regards to having such Rossian principles at the ready, such axioms are likely to clash, be it between respecting varying principles in one case, or one and the same principle in two separate cases where both can't be respected; as in the case of two promises to two people, one of which one will be forced to break if either is to be carried out. Whereas Ross was happy enough here to leave the case up to a sort of intuitive perception, many see this as an unsatisfactory answer to the issue.

There is always the sense that when one decides to respect one principle over another, or one case of the same principle over the next, if one is to be deciding rationally then there must be a *reason* for this decision that goes beyond a simple

appeal to perception. When one decides that promise X matters more than promise Y, if the decision is not to be a completely arbitrary decision, then case X must exhibit features which are absent in Y which the deciding party must view as being of greater value than those present in Y. These features will then act as a reason, which can be formulated as a principle, for deciding between such cases. Indeed, one may say that in promising one could form principles for deciding between promises; principles such as 'always keep the promise in which the beneficiary of the promise has greater need', or, 'always keep the promise in which the beneficiary will have greater expectations of the promise being kept'. Now, it seems, however, that even these two principles may well come into conflict. In the same vein of thought, one would look to further principles that are more general still to work out how one adjudicates between such principles. What we end up with is a set of common sense principles much like Ross's, but each being adjudicated by more and more general principles on pain of arbitrariness. As any two principles may well clash even these more general principles will require further principles which adjudicate between them until eventually we arrive at some first principle which adjudicates between all principles if morality is to be a fully rational process. Here is J.S. Mill, as quoted by Crisp (2014):

...there must be some standard by which to determine the goodness or badness...of ends or objects of desire. And... there can be but one: for if there were several ultimate principles of conduct the same conduct might be approved by one of those principles and condemned by another...

Such a picture of practical reasoning thus uses the fear of arbitrariness to paint a picture of commensurability. There must be, if decisions between actions on evaluative grounds are not to be arbitrary and therefore non-rational, a single principle of moral reasoning which allows for us to decide between the lesser principles when they clash. In this sense, all moral choices must be commensurable by way of the first principle of morality. The looming threat which motivates this thought is that in turning this picture over we are left with a picture of moral reasoning that is based in no more than personal judgement that could so easily be reduced to preferences and biases. Without some final end which governs and adjudicates between other ends it seems we have no rational warrant for choosing between clashing principles or concerns of a lower level. If we wish to claim the

moral sphere is intelligible, we must accept commensurability as a fact of that sphere. And if we are to accept commensurability as a fact we must also recognise that it brings in tow a system by which ends are judged on the basis of a first principle, with all justification coming from one first principle or end. Such a position is central for Utilitarian theorists such as J.S. Mill and Sidgwick (1981). We will come back to the idea of commensurability and the picture of systematisation that it seems to demand from us.

3.1.1 Value Monism: Happiness as the Ultimate End

To start with, however, we will consider a more basic proposal of Utilitarian thought which is distinct from that mentioned above where principles of common sense run up to a self-evident first principle or end which controls them when they clash. The simpler picture is the idea of a moral value monism, where there is only one single value that should be applied to all cases of moral thought. It is not that it overrides or controls the use of other principles. Rather, if we are practising ethics properly, it is the only principle, the only value that needs promoting. I will make three main attacks on this idea to show that the Utilitarian principle, structured in this way, fails as being fit to adjudicate each and every disagreement. Firstly, I will consider the revisionary arguments made for imposing a single end as final as happiness, at the expense of all other concerns, claiming that it is problematic. Secondly I will argue that this idea fails to capture the phenomenology of our moral thinking and the concepts utilised there. Thirdly I will elucidate why such a theory has a distinct epistemological disadvantage.

The Utilitarian, taken as a value monist, seems to have an advantage over deontologists in the simple sense that a single principle has no issues regarding competing principles. Each and every moral conundrum ends with just one thought: 'which of these possible actions will maximise happiness?' Once we have done some evaluative equations, the matter would be simply settled on that score. But what reasons do we have for believing that the maximisation of happiness, not just our own, but other's happiness as well, is the ultimate end for man? Mill's arguments for this position are well known and we will critique them here.

The first step in this tactic is to convince that we do, as a matter of fact, all value our own happiness over everything else and when we act for our own interests we act in the way which is most likely to secure our own future happiness. That is our end. Mill (2001: 35) puts his 'proof' this way:

The only proof capable of being given that an object is visible is that people actually see it. The only proof that a sound is audible, is that people hear it: and so of the other sources of our experience. In like manner...the sole evidence it is possible to produce that anything is desirable, is that people actually do desire it. If the end which the Utilitarian doctrine proposes to itself were not, in theory and in practice, acknowledged to be an end, nothing could ever convince any person that it was so

Mill makes various mistakes in this passage that have been widely discussed. The analogy in the most straightforward interpretation regarding proofs for experience should fail to convince us. That we do see or hear things in no way proves those things are there. People do and see things all the time that aren't there and we might say, 'you *ought* not to be seeing/hearing these things' to such unfortunate persons. Likewise, to claim that something, in this case happiness, is desirable merely because it is desired acts as no proof. In some cases of course we might say that one's own happiness 'ought not to be the thing one desires'. This of course is the naturalistic fallacy of Moore (1922: 11-20) at work: that people do always desire their own happiness is no proof that they should. What is under scrutiny here is whether what we desire is worthy of that desire. Mill in no way allows for us to get a grasp on whether this is the case with happiness. Of course it seems impossible for Mill to *prove* that which is desirable other than claiming it is desired, and surely without starting from an *a priori* position similar to Kant we could never start to believe that something that is desired by no one ought to be desired. But it is Mill's preoccupation with proof that stops him from considering our ability to use everyday reasons to argue that what we desire in each instance should be desired by others too. That would not be enough for Mill, whose commitment to this level of objectivity obscures his willingness to concede to this sort of an analysis for moral reasons.

Mill (2001: 35-360) goes on to say:

No reason can be given why the general happiness is desirable, except that each person, so far as he believes it to be attainable, desires his own happiness. This, however, being a fact, we have not only all the proof which the case admits of, but all which it is possible to require, that happiness is a good: that each person's happiness is a good to that person, and the general happiness, therefore, a good to the aggregate of all persons. Happiness has made out its title as one of the ends of conduct and consequently one of the criteria of morality

The latter part of the argument is also fallacious, this time the 'compositional fallacy'. Believing he has established that to each of us our own happiness is our ultimate end - the only end we desire for its own sake - Mill then makes the claim that, as happiness is such a good for each us, then general happiness must constitute the general good. This is an inferential transition from the first claim 'as my happiness is a good for me, all acts that maximise that happiness are a good thing for me' to the latter claim 'as each person desires their own happiness equally, any act which maximises the general happiness is a good for all persons'. The inference is bogus, however, and it can be easily questioned as to why there is any reason to think, from Mill's position, that a person who cares for their own happiness should care at all for the happiness of others. Claiming that there is this thing, 'a general happiness' - that constitutes the general good and must be made up of the 'happiness of everybody' - is thus equally fallacious. Something is not necessarily true of a whole, here the general good, simply because it is true of its parts; just because atoms are invisible to the naked eye, and you are composed entirely of atoms, it does not follow that you are therefore invisible to the naked eye. The inferential transition should fail to convince us that there is such a thing as the general happiness that each of us should desire, regardless of whether we believe with Mill that our individual happiness is something we do indeed desire and ought always to be seen as desirable.

Even if we attempt to augment Mill's argument, as many have tried to do, it struggles to find solid ground. We might augment it by lessening the appeal to the inferential transference between an individual's happiness and the general happiness, for instance, by claiming that if each person desires happiness *and we have no reason to see her to be mistaken*, and each person sees their own happiness as a good –

then we can claim that happiness is an intrinsic good. From here the maximisation of the good does indeed seem to be a duty for us, assuming we go along with the Utilitarian idea that we define the good first and then the right as the maximisation of that good. All that matters is happiness as:

...human nature is so constituted as to desire nothing which is not either a part of human happiness or a means of happiness [as such] happiness is the sole end of human conduct and the promotion of it is the test by which to judge human conduct...(Mill, J.S.: 39)

Even if we see happiness as an intrinsic good we need not see promoting that good as leading always to the right action. We might also see items such as comfort, justice, wealth, health or a multitude of other things as being things that we desire, and believe that there is no good reason for us to think we are mistaken in so desiring them. Why can't we just think of the above as being things we desire in their own right and not as a means to happiness? To make the leap from 'I desire X and X is a part of Y' to 'therefore I value X as a part of Y', is a fallacy again on the part of Mill. To attempt to subsume the way we value acts of justice under an idea of happiness seems to cheapen our understanding of justice in such a way as to make it unrecognisable. The separation of the good and the right in the manner the Utilitarian creed promotes is forever questionable in the way that Moore judged it to be. It is the Utilitarian's job to make us believe that when we are claiming an act to be moral or just we are merely saying that it is the act that produces the most happiness. But this reference to the happiness of the aggregate of persons seems to undermine our conception of other moral concepts on many occasions, as we will see.

Acts, such as promising, mostly only need to appeal to the people involved within the act itself. To bring into account the entire aggregate of persons and their happiness thus diminishes the bond that should be forged between the person who makes the promise and the person who is the recipient of that promise. Of course it is highly possible that one uses a disaster clause to explain away not keeping a promise – that to keep it would have caused such great harm that it would have been unjust on those who would have felt those harms, and that harm outweighs the injustice of breaking the promise. But this reasoning, while appealing to consequences,

pleasure, harm and an aggregate of persons is not necessarily Utilitarian or consequentialist reasoning. Other, non-consequentialist, theories can appeal to consequences as well; indeed, any sane person will appeal to consequences as a reason for acting in certain ways at certain times. One is only a utilitarian if consequences are the *only* thing one sees as worthy of appealing to in such instances, if one believes that the aggregate pleasure or pain caused by an action is the *only* matter worth considering. And the impersonality that such a move causes makes a mockery of human bonds and intricate moral relations such as those created by promise making.

This is why Williams (1973: 116-117) argues that Utilitarianism strips away the integrity that seems such an essential characteristic of moral agents. If I break a promise as a Utilitarian, based on good Utilitarian reasons, to feel guilt at doing so would seem irrational. I should train myself not to have such feelings, for to do so would only add to the overall store of unhappiness at my breaking of a promise. After all, I stand in no more a special relation to someone I have made a promise to than I do to every other person who makes up the total sum of happiness, and every member of that aggregate should have chosen as I have chosen. If I cannot train myself not to feel this unreasonable guilt, however, the guilt itself must still be taken into account in the calculation of the unhappiness breaking the promise will cause. For there is no judgement on what it is that causes happiness or unhappiness or whether those things are worthy causes of those mental states, only that those mental states do exist.

Under examples such as this we also see how separate Utilitarianism is from any pre-theoretical conceptions of justice, and how subsuming complex moral ideas like justice with the Utilitarian principle appears impossible. Take the example of a community, terrorised by a murderer, who's imprisonment and execution would result in such joy and relief that even mistakenly accusing an innocent man and executing him would have to be seen as bringing more overall happiness than not doing so, despite how much the loss of that man's future happiness and the fear of his impending death and his frustration at being innocent in this situation would count against the act. This example is very telling as it strikes a chord with all of us about how we value justice and how we define just acts. Assuming the real culprit stopped that very day for whatever reason, the innocent man, if he were a Utilitarian himself,

would be forced to admit his death was for the best and therefore not an unjust act at all.

Now suppose that sometime after his execution word gets out that the community has indeed arrested and executed the wrong man and as such the killer is still at large. The execution, that until this point of revelation was a just act according to the Utilitarian society, in the wake of this new fear and panic and remorse of the death of the innocent man, may now be seen suddenly as unjust. The privation of happiness being greater than it was originally in the second wake of fear, caused by remorse for the dead innocent man, the incompetency of the police department and the idea that oneself may also face charges for crimes they never committed. What is crucial here is that it is not that the act was *always unjust on its own terms* and the community have only now realised, but that a just act became unjust at the point of realisation and only *because of that realisation*. One cannot help but look at this example and wonder just how, even if we do take the idea of happiness to be a general good, that it is the maximisation of that good which is the sole and singular indicator of an act being morally right.

Now, I started by saying that the Utilitarian picture has the potential to be seen in a naturalist light. Happiness can be easily reduced to pleasure, conceived as an objective fact of certain various chemicals (such as dopamine and serotonin) being released in the brain, perhaps being created in various amounts by our actions as an indicator of how good the action is. However, Mill seemed to think of happiness as more normative than this simple naturalist position, promoting as he did the 'plurality of goods'. If Utilitarians do not take the naturalist interpretation then they have an extra problem, as now their theory would be reliant upon an idea of types of happiness being promoted which are evaluated as the 'right sorts' of happiness. The anti-Utilitarian literature is flooded with hypothetical examples of people willing away their lives in bliss in opium dens, or flooding the water supply with drugs which would make everyone ecstatic but also leave them living in an ignorant haze. If such critiques are deniable on the basis that base pleasures are not as worthy as the happiness that comes from, say, strong familial relations or the completion of a piece of rational thinking, then there must be some defence for just where this sense of worthiness comes from. And now the Utilitarian principle can have no greater say over which is better, as both are undeniably types of happiness. Such an argument

for 'worthy' happiness being promoted over 'unworthy' happiness would require something like an *a priori* analysis of happiness that would land Utilitarians of the value-monist ilk back onto a rationalist grounding. Otherwise wouldn't what counts as worthy happiness itself need to be delineated? And how one would analyse this would have no further court of appeal in commensurate terms. Wouldn't such an analysis of happiness bring us back to a place where the usual everyday concerns of morality - such as justice, honour, aptitude, success and so forth - be needed to explicate it and control it? If this is so, then the idea that the warrant for such concepts lies, all things told, in how much each is conducive to happiness, and that each only has merit to the level in which it promotes happiness, seems completely misguided. A desire to maximise happiness cannot be the ground for our promotion of what is worthy if a sense of what is worthwhile is a prerequisite for an appropriate sense of happiness to utilise in the controlling principle.

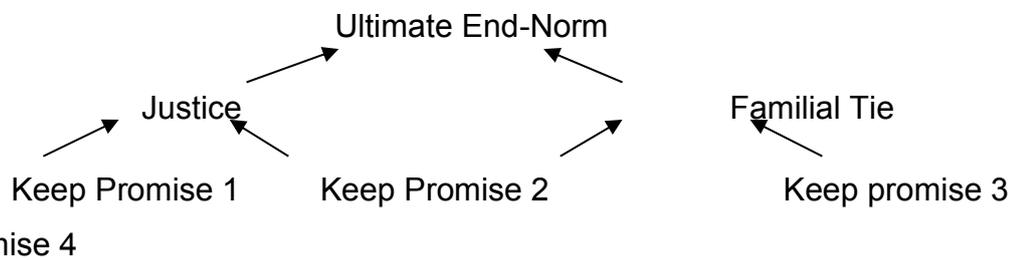
The final point to be made in relation to a purely value monist attempt at Utilitarianism is the epistemological one. When I decide what to do in a case, and I appeal only to the notion of the maximisation of happiness, at what point am I to stop my calculations and actually act? Act A may produce a great deal of happiness here and now but have the potential for great tragedy in the distant future, while Act B may produce little happiness here and now but have no such concerns. At what point am I at fault? At what point does blame for an action, one that was in some small way a catalyst for various actions of others that may rightly be called the consequence of that action to some degree, stop being portioned to me? How would one judge the American and British *coup d'état* against Mosaddegh in Iran, which allowed for the suffrage of women and the modernisation of a nation, when the sharpness of this contrast from a more orthodox Islamic lifestyles and the manner in which foreign secular and Christian nations had a hand in it was also a likely catalyst for the Islamic Revolution of 1979 which has left women in a position with greater restrictions than ever before and the country in political turmoil both domestically and internationally? How does one measure the happiness of such a situation and the blame or praise to be portioned to those who made the decision, when one can only judge on the basis of happiness?

3.2 Moving Away From Value Monism

We have seen then that the value monist notion of commensurability falls flat in various ways, be they methodological, conceptual or epistemological. Given that this is the case, the more sensible form of commensurability which ends in a single principle must pay homage to the way in which other values play their part in moral thought. This is true of the Sidgwickian (1981) approach to Utilitarianism, which makes great use of the principles of 'common sense' and their place in our thinking. These values act as our starting points. One cannot simply tell someone who is going through a moral education that they must in all cases promote happiness. Here, we get to the idea of happiness as the ultimate end via a system which works up from common sense principles such as those promoted by Ross. To theorists such as Sidgwick, what is the downfall of an account such as that posited by Ross is the previously mentioned appeal to intuition or perception to decide between clashing common sense principles. For people who take this route to an ultimate commensuration principle, when common sense principles clash, if we are to decide between them, there must be another principle which outlines what the appropriate concerns are for deciding between them. These concerns will be more general, picking out features which persist in cases where principles come into conflict, so as to help us see what is actually at stake in terms of value. These more general principles, if they are to do their job correctly, must help us prioritise between more concrete cases, allowing for a rational way to decide between disputes. So, while at the level of common sense, one needs principles such as 'keep your promises', when conflict occurs one must search for a principle which would rightly adjudicate which of the promises is of greater concern. This would provide a rationalist system for deciding the practical conflict and giving us a greater insight into what is valuable in each case. I am indebted to Richardson's (1994: 122-144) account of Sidgwick in this part of the section.

Cases can seem easily settled on this score. If there are two promises, equal in all ways, except one served a greater sense of justice in keeping it, or was made to someone closer to us than the other promise, then we would choose that which better served justice or benefitted our closer friend or family member. However, if between the two promises we decided that one seemed more pressing due to the fact that in so keeping it justice was served, where as in the other one a familial tie

was respected, and that it was one of these concerns which was valuable to each case, but not the other, we would then need some further principle, of yet greater scope, to decide between the concepts of justice and familial bonds.



One could easily see a case where an administrator at a hospital was left with the choice of tampering trials to let a sick family member in, when their reason for taking that office was to oversee that no such tampering occurred, as on principle, each applicant had just as pressing need for the treatment. In this case, promises to both office, 'to remain objective and unbiased', and to a loved one, 'to do everything in my power to help them live', violently clash. If such a case is not to be left up to intuition, some other principle seems to be needed to adjudicate between these concerns. Furthermore, it must have the ability to decide between them. But here, the two further principles which govern the ends that made the promises both seem worthwhile likewise clash. Were the administrator not averse to tampering and an acquaintance and a family member both needed the tampering then the more general principle 'respect family bonds' would decide between the promises, assuming in practical terms it was only possible to switch one person into the trial in a covert manner. Likewise, if it was an acquaintance only who asked for the tampering then the promise to act justly would likely stop our administrator from promising any tampering at all for the acquaintance. But in cases such as the first hypothetical in this scenario - where *both* promises require serious attention and neither promise can be separated by reference to the wider principles that guide why they were promised - yet some further and even wider scoped principle seems to be needed for us to appeal to in our decision between our ends of protecting family members and our ends of respecting justice in the role of trials administrator. Given that any two principles can potentially lead to conflict, given the multitude of plausible particularities in any case of practical reasoning, it seems we would require a way to rationally decide between even the highest two principles. This, in turn, would require

a principle that is the ultimate source of warrant; a single principle in terms of which all other ends are outlined, controlled and decided between.

In this manner, the route which provides an ultimate end of happiness as governing action - not directly but via a means of controlling the way in which more common sense clashing principles can be portioned their due - has, *prima facie*, less of the problems of the value-monist interpretation. Unlike the first score brought against the value monist, here the lesser common sense principles are evaluatively sound in and of themselves. They are controlled by higher principles, but one need not only apply the value of maximising happiness. The importance of such things as promises and maintaining integrity are unproblematically appealed to, even if their warrant comes from a higher source than this. In this way, we do not have the issue of the revisionary notions of the value-monist. Ordinary principles and values have their place. The third issue, the epistemological one, also has nothing to cling to here. We know how to apply principles in an ordinary manner. One only searches for answers at a higher level in times of conflict, and when one does so this simply means looking for more powerful and general ends which can act as governing principles. The epistemological quandary of how to apply the maximisation of happiness principle is used not directly to the concrete particulars but instead to other various principles which govern further principles and so on.

The second issue I had with the value monists, however, seems to remain; namely, that this idea of systematisation does not appear to be phenomenologically accurate. Here we seem to forgo how we come to think of moral ends as working in order to keep the strong sense of systematising morality into something that any rationalist will love: being able to deduce from an indisputable first principle down to lesser principles, with all authority flowing from this one point.

Such a picture should already raise our philosophical hackles somewhat by way of the analysis of 1.1 and 1.2. There we decided that whatever ends we have seem to stand outside the deductive chain of reasoning; but here deduction seems to be doing *all* the work of moral reasoning. What lies implicit in such accounts is the same positivist ideal that rules, if they are to be rational, must be underpinned by yet more general rules, the outcome of which can only be a rule (in this case an ultimate end-norm) which governs all other rules it presides over, much like a supreme judge or

arbiter. Much like Kant's single rule for decision making then, this picture is strongly foundationalist. Such foundationalism is inspired by the thought that only deduction can lead to definitive 'oughts' with a primary principle, assumedly an unchallengeable one, as the source of the authority. But we have reason to find the very idea of definitive 'oughts' suspect. To know what we must do, whether that doing is to be understood in terms of belief formation or intentional actions, is dependent upon a normative construct that will always be challengeable and as such never logically compelling.

Eschewing the rationalist assumption that such a fixed ground is necessary for knowledge, systems of ethics like that of Sidgwick's appear to have a motivation that is the reverse of what we would expect in their formation. It is the idea of system that seems to motivate one to look for an unchallengeable authoritative commensuration principle, as opposed to the attempt to understand the phenomenology of moral thought which inspires the search for a system that utilises such a principle. Once we realise that there is little hope of deduction securing definitive conclusions or 'oughts' in practical reasoning, the appeal to it should be lessened. A bootleg imitation of the reasoning that one would ideally like to be available, if one is to accept the assumption that such a thing is required to secure epistemically sound moral and practical conclusions, is no substitute for the real thing. As that ideal is a logical impossibility, given the way that any end stands distinct from the reasoning chain itself, it seems that an exploration of other forms of reasoning is called for. The fear is that reasoning will not be discursive if we drop the idea that each principle we hold is explicable only in terms of another of greater generality. This would cast us back into Ross's territory. Yet we have every reason to hold that just plumping for one choice over another without being able to explain why leaves us in a position that could be appropriately described as emotivist and therefore irrationalist on this basis. What we will need, then, is a way for discursiveness to remain as a key element of moral thought, without appeal to deduction to secure it.

For now I wish to conclude that any system that utilises an ultimate-principle has the major issue of proving why it is that one value which should commensurate all other concerns when other ultimate-principles seem readily available; and how it is that we are to measure all other decisions against it. If the argument is something like the question that Sidgwick poses - 'if we are not to systematise human activities by

taking Universal Happiness as their common end, on what principles are we to systematise them?' (Freeman, 1994: 324) - We might ask why it is that such a systematic view of deliberation is even called for in the first place. Deductive logic cannot hold the sway we wish it to, as any end we choose will always stand outside of the deductive chain and seems to require a judgement on the part of the deliberator. If 'perception' is needed to decide between competing ultimate-principles then we seem to allow an irrationalist element into our reasoning. But perception and deductive reasoning need not be the only choices on the table. It is made to seem so by the appeal of a top-down rationalist system, where the only way of deciding between relevant concerns is by appeal to a more general principle. Following Richardson's analysis, one might well point out here that generality and warrant simply need not be 'run together' (Richardson, 1994: 138) as though each layer of principles presided over those below like a 'court system' of increasingly greater authorities. As he eruditely points out, if this were the case then it would imply that the more general the principle the more true it would have to appear. This, however, seems to be false in our phenomenology of moral and practical reasoning, as:

Among normative judgments that stand most firm are...particular ones...about the evil of the Holocaust...the admirable character of Mother Theresa's work with poor and sick...about the injustice of Idi Armin's rule...Conversely, it is so difficult to frame general principles...that we are constantly indicating our doubts about them by hedging them.. (ibid)

3.3 Evaluation from Bottom to Top:

I would add to this analysis that in coming to frame general principles and properly understand them, we have to use concrete examples to even broach the idea of the more abstract ones. There must be something which is to be *abstracted from* for these general principles to have the force that they do. Take a commensuration value such as happiness. If happiness is truly to commensurate the increasingly more specific principles below then its content must in some way be available in a more specific form than the term 'happiness' allows for. An argument I levelled against the value monist, which a Utilitarian of the Sidgwickian ilk seemed, *prima facie*, to avoid, was the idea that happiness cannot begin to function as a

deciding factor without a specific manifestation of happiness being in place; and then there is no court of appeal left as to which specification is better. Maximising one sort of good to achieve the right action, be it autonomy, happiness, welfare, or whatever else, still runs into issues of incommensurability given that not all manifestations of happiness or autonomy or welfare are of equal footing. It is not that both choices lead to equal measures of the same value that is problematic, but that either choice could be seen as greater than the other on differing conceptions of what that value actually entails when being utilised in a concrete case.

The happiness of basking in the sunshine and finishing a marathon, the autonomy of the right to privacy and the right to roam unhindered, or the welfarist notions of being told that one is dying so they can sort out their affairs immediately, or not being told so they do not have to face the stress of their immediate demise, are all examples of each type of ultimate-principle clashing with itself, where further appeals to the relevant ultimate-principle would be rendered useless for deliberation of what one ought to do. Maximisation seems to fail because any proposed ultimate-principle is more complex than the single principle it pertains to be in guiding everything. As such, clashes seem not just possible but completely inevitable. We posited that the Sidgwickian eludes this particular criticism due to the way the ultimate-principle or end norm can only be understood if it is seen as delimiting the principles of common sense which will be the first appeal in any piece of deliberation, with more general principles coming into play as and when clashes occur. However, when one pictures just what this means we start to see just how the system that Sidgwick envisages is anything but deductive, with reasoning being seen as a trickle-down of warrant from a vague ultimate principle to various more specific principles.

If we are to take seriously the epistemology of this system, we have to see that coming to learn how an ultimate-principle such as happiness can delimit the lower principles would involve, first, taking note of the various differences in the more specific principles themselves. An otherwise poor parent from the first world, who receives an unexpected small fortune might well invoke the justification of greater happiness for either giving the money to Oxfam, to help the starving children of third world countries instead of to one's own relatively needy offspring, or vice versa. For an ultimate-principle to have any deliberative power it must be able to make distinctions between the choices it presides over, and to make distinctions it requires

specific content which can be shown to be present (or present in greater quantity) in one choice over the other. But it seems that varying conceptions of happiness are possible without facetiousness being an excuse. The normativity of terms such as 'happiness' or 'welfare' and so on, lead to a variety of conceptions which have no further, more general, court of warrant to appeal to in deciding between them. The source of these conceptions is in fact our way of understanding the lesser, more particular principles, that the top-down rationalist would have us believe are controlled by the ultimate-principle. When we decide between the relevant content of such ultimate-principles, what is relevant to that decision is still a judgement along the lines of 'what is more important here' – with this more complex phrase being glossed, later, as 'happiness' or 'welfare' as though that term was no longer equivocal. Unless we take happiness to mean an empirical measure of pleasure as reduced to various brain activities, we are always maximising a particular conception of the good that has potentially just as much scope for competing conceptions as we would do if we merely said that, as a principle, one should maximise 'goodness' itself.

Thus, invoking a maximisation of 'happiness' or 'welfare' seems to be merely a place holder, one which requires us to look at the particularities of the situations we are in if we are to give it meaningful content. For in the example of the last paragraph, one's child might complain legitimately about a lack of loyalty or sense of commitment to one's own, of selfishness through self-aggrandising, or lack of immediate compassion, were they not to receive any money. The charity worker might complain of short-sightedness in other's suffering, of a lack of empathy for those far away, of the disparity of the need felt by the third world children in comparison to those of the parent's children. The child might say the satisfaction of expectation would lead to a happiness of knowledge that one's parent truly valued them as they had believed and that a hard life was turned into one of pleasure and possibility. The charity worker might claim that the happiness gained by a well of fresh water in a small community would cause greater happiness than any luxury the would-be-benefactor's child might use the money for. The point is that all these concerns are themselves legitimate concerns regarding who should receive the money. Questions of character and action both have power here. There can be no abstract, ultimate-principle or end which we appeal to in order to decide between

which concerns are most worthy or relevant. Of course we should do as much good as we can. But what that goodness amounts to can only be decided by considering in their own terms which ends are pertinent. Deciding to take happiness as an ultimate end in helping to decide between the two claims here would require deliberating about which happiness has greater claim, which is more worthy. Regardless of the specifics, it is hard to see how it is a commensurable notion of 'happiness' here and not a general sense of what is good, hanging as it does on the various moral concepts of charity, compassion, happiness, welfare, selflessness, gratitude, loyalty to familial bonds (and many more not mentioned) that is doing the deliberative heavy lifting here. If one's conceptions of what concepts make appropriate ends and the content that make those concepts meaningful in deliberation come from living within morality how can it be a simple top-down system based on deduction that explains the whole of ethical thought?

It appears then, that incommensurability is a feature of ethical thought that merely has to be accepted. Even accepting a common ultimate-principle – by our analysis a phrase that merely short-hands an appeal when there is no further deliberative issues to consider- in no way necessitates that the content that such a ultimate-principle will have in guiding actions need be the same as someone else's conception would have it. Once we allow that any linguistic proposal for an ultimate end norm is in need of filling out and specifying via some sort of reflection upon the concepts, those that the rationalist would like to say the ultimate end norm regulates, we have reason to keep open the idea that moral thinking cannot be a matter of mere maximisation of one singular value. Rather, we could only see what one really means by the phrase 'maximising happiness is the benchmark of the moral man' when we see them act upon such a principle and thus the content of what happiness actually amounts to, as they have it, is revealed.

My contention here remains that a principle or end such as 'maximise happiness/ welfare/ or any other proposed ultimate-principle' is just a placeholder similar to 'maximise the good' and each remains as abstract as 'the good' until one sees what concerns are ignored or cherished by the agent purporting to follow this principle. If one does not take seriously the lot of those who will be saddened by an act for selfish reasons - in the same way they might those who would be saddened for benevolent reasons given that 'selfish people would be happier if they learned not to

be so selfish' - then we can start to see that there are more complex concepts at play than just the mere pleasure taken from an act, where pleasure is understood in the more quantifiable way that a naturalist utilitarian would have us take as a ultimate-principle.

Likewise, there is no deductive argument for choosing any one ultimate-principle over the next. The distinction between the sorts of content that we usually see happiness as imposing, which welfare doesn't, when either is utilised as an ultimate end norm of action, is normative to its core. They both act as placeholders to be later specified. It's not merely happiness, but the right sort of happiness that one wishes to promote. This is not to say that such specifications can't happen in such a way that one gets incrementally better at seeing how each has its effect on some ultimate-principle and, from there, use a greater specified version of that ultimate-principle to start to assess these concepts. It is merely to say that the process of specification necessary for an end norm to undergo to start to become useful in any proper sense of deduction to lower norms is itself not deductive. It is, instead, reflective; a matter of reflecting upon what matters the most. Whatever descriptive sortal one uses to denote the ultimate-principle one believes one is referring to, in order to utilise it effectively one must make sure its content is relevant to concrete cases in such a way as to allow for it to effectively have some traction upon those circumstances. Once this is said one could well substitute happiness for welfare or, as previously stated, 'the good' and claim that one should merely maximise *that*; with 'the good' merely being a gloss for the various list of concerns that should matter most to us *qua* rational persons. In this sense the process necessarily needs to be bidirectional (Richardson, 1994: 141 and Rawls, 1971: 20). The content of any thought process moving down from the ultimate-principle to delineate less generalised concerns will be informed by an analysis of those concerns, in their *own* terms.

3.4: Naturalism, System and Motivation

A quick recap of the last two sections, and their main objects of discussion, starts to reveal certain characteristics that are shared by these normative theorists. They each have a desire for moral thinking to have a singular touchstone which can

be utilised to come to a conclusion about an act's moral status. The role such a touchstone plays, be it an ultimate end norm, intuition or a categorical imperative, has two major applications. The first is to give all moral statements a common ground by acting as the epistemological hook we discussed earlier in the thesis. This epistemological hook acts as a defence for one's reasoning. The thought pattern goes: if moral reasoning is meant to be reasoning about a real subject then its content must be categorisable and recognisable as such. The seeming existence of competing claims in particular instances overlap in such a way as to make them untenable as being consistent explanations of the same phenomena. If any such phenomena are to have any reason for our attention then it must be shown that some such claims in some way explain what the essence of moral thinking is about, and actually touch upon a subject matter that is explicable through that essence. A dualism in thought here, that more than one hook exists, would lead to conflicts about what the subject matter under discussion actually amounts to.

Clarification of this first point can be given by looking at how we deal with similar claims regarding other subject matters. If we wish to see if the test results of a particular scientific study were correct then we better clone that study as much as possible, sharing the same sorts of norms about which variables to highlight, which are redundant and so forth. We cannot back up the empirical claim the first study makes by way of rigorous scientific research by asking Mystic Meg. What guides this thought is normative intuition about knowledge – that knowledge requires a definable, consistent and stable standard which applies in all similar cases – which gives us reason to claim that all such events should be decided by appeal to this epistemological hook. Wielding more than one hook would leave us in a Sidgwickian position of having no further court of appeal to decide between them deductively, which explains the lament at the end of his 'The Methods of Ethics' (1981). This would require a judgement, which as far as the orthodox view goes, would leave us without a proper claim to knowledge; neither choice would be ratified by further definitive justifications as being *the* correct explanation. As such, either choice can have no explanation of why reasoning in that way is better than another way.

This last part touches upon the second idea inherent within the draw to having a touchstone: that of a need for a system or method of understanding. In empirical matters the empirical world as shown to us by our senses seems to ground them. As

well as deciding between accounts so as to imbue us with an idea of how we come to know what the essence of a category of thought is, the touchstone also plays the role of showing how that essence instructs us in dealing with particular instances of that category. We have already shown how hard it is to find a way that deduction can play a central role inside moral reasoning **1.1** but the draw to such a feature is almost universal across normative theories. In both top down forms the idea of deducing from certain premises to conclusions is a prominent feature. When confronted with two choices, one utilises the touchstone to clear away doubts. In Kant's moral theory deduction enters via the idea of contradictions. We cannot choose just anything as matching the Categorical Imperative as anything not worthy would contradict with other ideas that would pass the test. Working only with what we see as the essence of morality, we can always come to the correct answer as to what to do by appealing to the principles in place to designate the salient features of that situation and provide an answer; one that could be said to be correct, as long as the theory which promotes that distinct essence of the subject in such a way is likewise correct.

The issue I have opened up in the previous sections was meant to show that it is not in the deduction itself, but in the way by which we arrive at a touchstone from which to deduce, which seem to cause moral theories (and by way of **1.1-2.2**, also theoretical theories) to suffer in terms of making knowledge claims. The same certainty which we strive for in coming to a final conclusion from a normative theory is seemingly absent in coming to decide upon *just* that theory from all the available and possible theoretical positions there are. In light of this, it seems that any claim to knowledge from any particular theory loses its supposed punch, utilising, as any suggested theory must, a claim to certainty on the basis that the moral landscape, as it sits untouched by theory, gives us no definite clues as to what sort of organising and categorising theory might be best. Once this is understood, it becomes somewhat obvious how each theory available is a sort of restricted version of morality as we know it, pre-theoretically. Each promotes one certain type of claim that comes from a definitive type of intuition over all other types, and purposefully rejects the content of certain intuitions in favour of others for the sake of what we might see as the 'form of reasoning'. In short: to keep the norm of system in regards to knowledge in place, we must, due to the conflicting set of intuitions involved in the

moral sphere, choose to group together certain intuitions which can be placed into a system over those that fall outside the system.

When I say the 'form' of reasoning, I am referring to the idea that if one is to say 'X is the right answer here' with any confidence, there should be features which brings one to see X as the right answer; features that are available to all other situations which require that type of reasoning, which could then be applied to it. In other words, these norms regarding knowledge acquisition and what it requires seems to demand we find a consistent method or system of reasoning. This would be a reflection of the sort of analytical reasoning that arithmetic requires; that we utilise the same process for addition regardless of the numbers involved in the equation, or the situation which requires we perform such an equation. The numbers act as *the content*, the situation acts as *the context*, but that which guides the process of actually adding the numbers together is *the form*. Thus, the touchstone, by way of being the element which designates correct behaviour in light of a situation, also acts as the 'rule maker'. It undercuts rivalling intuitions and competing claims by focussing only on a particular sort of claim, ones that can be systematised in such a way as to make a cohesive picture of morality.

Once this is done, we have all we need so as to say that rules can be applied easily, with a system whose output could very easily be seen as knowledge by those who hold the form of reasoning as being so central to calling any output knowledge. But the sacrifice to any such system will always be the content of the competing claims that are ruled out of court as not really having any input to moral issues 'properly understood' by the creation of a cohesive touchstone in the first place. Thus Kantians rule out any notion of promoting goods recognised *a posteriori* and Utilitarians rule out our ability to understand rules that are not guided by the maximisation of a particular good such as happiness. Both manage to create systems which, in terms of form, are workable normative theories, but both do so at the dramatic cost of various instances where to agree with their output requires one to become deafened to one's intuitions. And the literature is filled with such examples, be it dealing with an axe murderer asking for the whereabouts of a would be victim one is sheltering, or understanding how justice would function in a world where maximising happiness should always be the ends of a moral agent.

As far as morality and moral reasoning goes, it is this deafening oneself to intuition which seems a somewhat hypocritical and absurd suggestion at the normative level once one realises just how necessary intuition is in getting to that point. In a pre-ethical world, there is nothing about killing which is inherently wrong, nor anything about looking after one's own needs at the expense of looking after others. Rather, we seem to come to know these acts, or certain examples of these acts, as wrong by way of being inculcated into an already functioning moral world; a world where appeals to consequences in one situation and then appeals to character traits or to acts no one should ever do in the next, does not seem to elicit complaints about consistency.

As pointed out at the end of 1.2 intuitions are a huge part in assembling any moral theory. What norms are taken as read to be good or bad to hold and which ones are left to see how a finished theory assigns them is up for debate, and in a pre-ethical domain none have anything but the beliefs of the person constructing the theory for backing. To find out what matters most in all cases, one requires, at the very least, some norm or norms to be seen as what matters most, in at least some cases, in order to start finding a pattern which is stable and can then be applied to new cases. Again, as much as we seem to need a theory to make sure we are not just working from bias or distorted intuition, we require intuition to start any theory which so designates some claims to be salient and others as not so salient or not salient at all.

3.4.1 Norms of Form:

What I now claim is that the one norm which is present in almost all orthodox theories, theories which end with a rule or rules which are meant to definitively guide behaviour, is not a 'norm of content' but a 'norm of form' (or system, or method). Furthermore, I believe that it is clinging to norms of form as though they are somehow more sacrosanct, more definitive and more worthy of being untouched is both an arbitrary judgement and one which has caused the moral field serious harm in helping to comprehend the norms of content and how they are pervasive in our moral lives. When we have considered theories which use touchstones of morality - touchstones which categorise content norms in such a way as to include certain

content norms that fit a pattern at the expense of others, which often still weigh heavy on our intuitions - the reason for our acceptance that such norms must be mere intuitions (assumedly the product of individual or group bias) is precisely the idea that content norms require to be fitted into a stable and cohesive system to be able to make judgements about the content of morality properly. But alongside the idea that certain intuitions get ushered away by being labelled as 'mere intuitions' by some theory, whereas others become the very essence of morality according to that same theory, we have also seen how the appeal to a strict system which attempts to omit the use of personal judgement and omit error seem to falter.

The Kantian view on system is so strong that it is the form of practical reason and not the content of what that reason should promote that takes the central role in the Kantian account. In theories, such as Utilitarianism, which promote a commensuration system of deduction from ultimate-principles, we likewise see faults in the idea of systematically deliberating. There, the ultimate-principle meant nothing without being given content by the very elements of moral life that they were meant to regulate. This bidirectional process seems at first blush to rule out the idea that such a system could work without the use of proper practical reasoning filling in the content of the ultimate-principle first. As such the appeal to system failed here too - conditional upon how that 'filling out' might work - in creating a theory that included all the aspects of morality that seem to concern us.

This analysis leaves us in a tricky position, of intuitively accepting the idea that morality can be fully systematised only at the cost of wilfully ignoring the call of the very same faculty, intuition, which suggested that there need be such a system. It will be the purpose of this section to show that the idea of a complete system is not a necessary one for us to aspire to, and at the same time consider the ways we do reason about moral issues and the way we can see this reasoning as really interacting with a proper content. In order to do this I will consider a different type of account which runs from a naturalist basis, and which eschews the idea of a stable system for morals. What it manages to do is to leave the idea of the system of reasoning behind without losing the idea that there is such a thing as a hook which could theoretically point us towards an unassailable reason for action in the moral instance. As we will see, such a call will be its downfall. In overcoming a need for a

definitive system however, it fails to take into account the important issue of moral motivation, of what it means for us to be motivated to do the right thing.

My own account of moral reasoning will both disregard the idea that any reason is unassailable in the logical sense and take incommensurable premises as a stated fact; two ideas which I hope to show will no way discredit the idea that one action can be the right action to take when two such possible actions are in play. Likewise, though later in the thesis, the expanded theory should also hold onto a place for moral motivation and explain why and how we are motivated to act for the benefits of others. Furthermore, it will hold open a place, which has already been hinted at, for how we do come to deliberate about moral norms by looking at how we actually come to deliberate about norms in general. This exposition of proper deliberation, about what will then be determined to be an appropriate topic for reasoning, will complete the account.

The idea that we need a system at all comes from the intuition that there need be definitive reasoning for our categorising some acts as right and others as wrong. Without a clear grasp of this, one could move from one case to the next, merely asserting one's own intuitive bias upon the case at hand, and reasoning from there; concluding only relative to one's own biases as to which acts fall into which category. In order to reason to a conclusion that has epistemic credibility, it is thought, there must be some pattern, equivalent to that of mathematics, which allows for certain features of a situation to be seen as more salient than others; with these features being seen as the essence of moral thought. Once we have identified an essence, by way of naturalism or rationalism, which can be said to prevalent in all such cases, we have a proper content for moral reasoning. Theories that seem to go against this idea of systematisation that we have so far seen are those like Ross's, or indeed Moore's, both of whom focus on perception as giving us access to intrinsic goods. The issue here seems to be a sort of explanation as to how this perception of what Moore calls non-natural properties differs from our merely sensing the presence of the natural world. We will come back to this. I believe that such an epistemological case can be made out in favour of such perceptions; not by postulating mysterious faculties, but by showing that the empirically driven conception of perception is not nearly as innocent as we believe it to be. In which case what is being set up is an illicit unfavourable comparison between the sorts of properties that Moore would call

natural and those he would call non-natural. Moore himself seemed to lose confidence in his own theory, essentially allowing a Harean (Hare, 1963: 55-75) conception of descriptivism and following the classic Humean empiricist conception of empirical facts and values that are projected on top of said facts. This in turn reduces evaluations to mere commendations for Hare, something any empiricist would be happy to agree with.

3.4.2: Virtue Naturalism: Eschewing the Norms of System

The theory whose shared feature will make up the bulk of the discussion in this section proposes an option that is quite distinct from other attempts at defining 'the good'. It does so by attacking the very idea that there is such a thing as a distinctly moral category which will spurn distinctively moral propositions; at least in the way that other theorists which we have so far canvassed have proposed. For Thomson (1997), 'good' does not stand for any property or definition yet to be uncovered. Instead, it functions in sentences by merely showing our *approval* of certain states of affair, whatever they may be. Not subjectively, however, like in Hare's account. The way Thomson defies such commendations from reducing to a subjective affair, into something where 'good' can be meaningfully applied, is simply by extending a category of goodness we are all quite familiar with and which bypasses the usual complaints of the fact/value gap (and its cousins, the is/ought and belief/attitude distinctions). It does so, by precisely denying that 'the good' has an extra category of meaning, the specifically moral sort, and dissolving all uses of it from seemingly complex propositional issues into seemingly simple assessments of what features are good (in a non-moral sense) for certain types of things to have.

The paper which I will analyse in this section belongs to Judith Jarvis Thomson in 'The Right and The Good' (1997: 273-298). Her analysis shares features of various other naturalistic based virtue accounts (see (Hursthouse, 1999) in as much as they take a similar route in tying claims which include evaluative statements, such as 'good', with factual statements. A good example of this idea and how it evades the fact value gap is given by MacIntyre. His own account seems to flirt with the naturalist position early on but in fact moves past it. His example in the paper serves as a decent starting point however:

From such factual premises as 'this watch is grossly inaccurate and irregular in time keeping' and this watch is too heavy to carry about comfortably', the evaluative conclusion validly follows that 'this is a bad watch (**MacIntyre, 1981: 57-58**)

There is an obvious draw here to the idea that if man can be made out as a functional concept in the same way that watch is made out, then if we were to know the function then we can know what sorts of principles to use, as they would be the ones which take us from how we are now to the state the proper function decrees; much like we know the steps to fix the watch, they are predicated upon how the watch should work. Rather than take it that there is such a fixed set of ends for moral agents, however, Thomson takes it that we use 'good' contextually in each instance it appears. Whereas naturalists who ground their use of evaluative terms on such an understanding - proposing that a good watch is one that fulfils a distinct essence (accurate and portable) of what it is to be a watch - Thomson would have it that a good watch is one that is contextually fit for the occasion. This means that while the natural teleologist's would claim that a good watch is not made so by fulfilling a purpose outside of its more strictly understood use as a time piece, such as 'throwing at cats', Thomson would say that the sentence 'that's a good watch', is perfectly understandable as also meaning 'this is a good for throwing at cats'. At least she would *if the context* allows that the listener understand it as such. If one is in a watch factory surrounded by angry cats and is asking to be passed things that are being put to use in the context of making sufficient missiles to scare the felines away then, 'hand me that watch, that's a good one' is unlikely to be taken by the person handing out the watches as meaning 'pass me that watch, as it is a perfect example of what we would in most contexts call a good watch'. The context allows that the watch's (or indeed anything's) 'being good' be only understood as meaningful in its immediate context.

This is an important point for Thomson, as she wants to move past the idea that anything need have a specific nature which would allow it to be a 'good' anything, devoid of the context in which it is good. There is no such thing, she says, as *just* a good book, or *just* a good watch, or *just* a good person. Thomson's moving past this position is inspired by her wishing to move away from understanding goodness as a property of any sort, much like 'bigness' is.

Following Geach, she claims that 'X is a good K' is *attributive* as opposed to *predicative* (Thomson, 1997: 277). Saying, 'X is a good K', is not to say of X that 'X is good', and 'X is a K'. It is to heap praise upon X for being 'good, plus adjunct' - that is 'good at...', 'good for...', 'good with...', 'good to...' and so on (ibid)(with ... leaving the adjunct open to describe whatever it is that we are claiming X is good at, with, for and so on). Without the 'plus adjunct' part, we are left with the idea that 'some substantive need be understood' and that substantive would leave the item under scrutiny as requiring that it have some essential purpose or function. Whereas the teleological naturalist thinks there is indeed such a thing as being 'just a good K', Thomson has seemingly shown that there is no such thing as 'just a good K'. 'For what K', she asks, is 'being good for the use of making cheesecakes...being a good K?' (1997: 279) Her point here is that while one might utter the words, 'that's a good cheese', the context would actually make it a perfectly legible statement and use of good, even though one could despise the cheese on every other level - its taste and texture perhaps - *except* when it was used to make cheesecakes and baked with various other ingredients. In which case the context of that statement would mean that 'that's a good cheese' would be perfectly appropriate while in no way determining that 'this cheese' is in fact good all told (whatever that would mean) or good even in its more regular capacity of being just good to eat.

If we compare this to 'big cheese' then we can see the difference. For while bigness is an attributive adjective, some substantive *is* understood here that surpasses the context. That's a big cheese is obviously attributive (that's big *for a cheese*), and not predicative (that's big and *is a cheese*). There is no need for contextual clarification, the statement is simply and plainly factive – it is either correct or not against a substantive standard for what constitutes a big cheese (or a big version of whatever object is under discussion). The question 'big in what way?' is not a valid question here unless the interlocutor misunderstands the comment. Whereas 'that's a good cheese' cannot be plainly factive, as seen against some substantive that remains constant as a property amongst all good things of that type. Goodness is not a singular property then, it is contextual. So, Thomson's account eliminates this issue by not focussing on what makes a good *telos* for a person or object in general. She achieves this by disregarding the idea that there are singular, functional terms that are necessary for our understanding goodness in a particular

object (say, time-keeping for watches), instead allowing goodness to be directed towards whatever function the object is used for at the time, regardless of its 'usual' purpose. If one wants a watch that will show off how rich one is then one coated in diamonds and gold, however heavy, uncomfortable and generally impractical this makes it, will mean that it is, here and now, a good watch. It is a 'good watch plus adjunct' – and whatever that adjunct might be, if its end is fulfilled by the watch (even if it is for throwing at cats) then to say it is a 'good watch' is correct in this manner – but only here and now. On this account of 'first order' ways of being good, there is no such thing as 'just a good anything'.

Thomson uses this analysis as a foundation for 'second order' ways, which while 'resting upon' first order ways, do not follow the formula, 'good plus adjunct', we are told. Justice cannot be understood in this way, for example. That it is 'good for Bert' (1997: 282) is not enough to account for an act being just. In her mafia example (ibid: 274-275) - where Alfred is told by the mafia if he does not kill Bert then they will kill three other people - Alfred not killing Bert is indeed good for Bert, but bad for the three others. But, claims Thomson, there is a reason why we still want to say that killing Bert is unjust, even though, on balance, it is bad for more people. This is because such things as good acts and good people do not fall simply into first order ways of being good. They are second order ways of being good. Being just is good, it seems, because 'it is better for us that the people among whom we live possess this trait [being just] than that they do not' (ibid: 284). And being just is a 'proneness to doing what one owes to others' (ibid: 282). As Bert is innocent of any wrongdoing he is owed no harsh punishment such as the capital punishment the mafia would have Alfred commit. Alfred owes no such punishment to Bert so it is good (or right really, but we'll stick with goodness here) that Alfred refuse, whatever the consequences to the other three. And importantly, Alfred cannot be accused of doing wrong to the other three; it is the mafia who will kill them after all.

So how does this goodness get a foothold in Thomson's account? Well, she seems to ground it upon the somewhat sketchy idea that it is 'better [good] for us' if people have the 'proneness' to committing just acts (as Alfred does in sparing Bert). The 'us' here implies a community. In fact she goes further than this and says

...it is **better** for us that the people among whom we live be just than that they not be just. Indeed, this is not merely better for us, but **essential** to us, since we can form a community at all and thereby obtain benefits which are essential to us and which only community can provide only if a substantial number of those among whom we live are just (Thomson, 1997: 282) (my emphasis)

I think this quote is revealing because it offers two possible readings of Thomson's account, both of which are distinct interpretations of it. On one reading it is just *better* for us that people are prone to the virtuous types of action she will go on to describe and delineate in her usual elegant manner. This would imply that we will be told why it is better than any further option, and as we will see, there are other options available for this position, options she is competing against. As such, we seem to have opened up an evaluative debate about what sort of lives are the type that we ought to live. This moves past what sorts of principles we ought to have in a powerful way. On the second reading, these traits are not merely better for us in some evaluative sense, but they are *essential* to us – necessary features for our very existence as a species (which is taken as an intrinsic good, thus the move from 'better' to 'essential'). I believe that in some sense this is what Thomson is going for. Foot, who she praises early in the paper, had strong naturalist leanings also and she is generally taken to be of the naturalist ilk by various other theorists (McDowell, 1996: 167-197) Likewise, Thomson's analysis of first order ways of being good seems to be grounded naturally, by showing that good can be used non-controversially (at least at times) as a factual matter, once the end is proposed that the suggested good thing is to be the means to, contextually speaking. Thomson's account seems to allow this reading, if not promote it. There is an issue for Thomson if we read the account this overtly naturalist way, indicated by the 'essential' in the above quote; a problem we will discuss later on. I will start with the less naturalist reading of that same quote.

Let us begin again with the idea that, 'it is *better for us* that the people among whom we live be just'. The whole question comes down to the italicised phrase. Given the fact that good is 'always good in a way' - and 'betterness' is taken as an expression of something being 'more good' in a way than another option for that way – then, in *what way*, we should ask, is it better for us that people have these traits such as being just? We are told that 'good plus adjunct' does not sum up any second order

properties; that 'X is Just' is not 'equivalent to 'X is good plus adjunct''. And yet there seems to be only one move for Thomson to make in describing how we come to think of just acts, or a proneness to commit them, as being good, if she does not want to return to the idea of substantives being understood where good becomes once again some type of property. That way is, simply, that it is 'good *for us*', which is indeed, 'good plus adjunct'.

Saying 'that is just' is a way of saying 'that is good' – a specific way. But that specification needs to be grounded upon something in order that various rival views not compete for the content which Justice could plausibly hold. I don't see a change in the sort of form of describing what it is to be good in this way from one that is good in a first order way. 'Good for Alfred' is highly specific. We can say 'X is good for Alfred' quite non-controversially, just in case X benefits Alfred in some way. If the means to that end are sufficient then 'X is good for Alfred' is a simple enough equation. But 'X is good for us' seems vague enough to evade being captured in this formulation in any meaningful sense. The 'gap' (Thomson, 1997: 282) in Thomson's story seems like the most important part at this point; for what it is for 'X to be better for us' in such a way as us to think of 'X being just' seems hard to cash out. To claim that 'X is better for us' seems to be saying little more than 'X is the sort of act that we characterise as being typical of a person who has grasped the best way to live'. But this seems to define nothing but itself; 'how we should live' is precisely the topic under scrutiny in moral philosophy. On Thomson's theory, especially by way of this reading, what makes an act morally good is that it is *better for us* that we perform it. I'm not sure which other theory, even one which would assign a different content to Justice – a theoretical account in which it is thought that for Alfred to be just he must kill Bert, for example - would disagree with this point.

It is clear at this point that 'better for us' is not enough to bring us into alignment with Thomson's idea of what she wants from her account if we are already prone to disagree. A Utilitarian might think here that, as following the principle of maximising happiness is what makes an act good, it would be strange if it turned out that following that principle was not the best way 'for us' to live. What sensible theory would posit an answer to the question at the heart of moral philosophy 'what sort of life should we live?' and then proceed to claim that while 'we should all live that way (or by these principles), it is *better for us* that we don't'? As such, the idea that we

started with, that a proper sense of justice rules out killing Bert seems to be question-begging at best. It depends *whose* sense of justice we are defending here.

The problem seems to come in the standard appeal to end/means instrumental reasoning as a deductive process that can be seen as non-controversial. In the cases of first order, it is exactly that 'good for Alfred' (or, making cheesecakes, reading, lawnmower repairs and so on) makes the equation simple by way of means and ends. But the object at the end of the adjunct, be it a well risen cheesecake or the health of Alfred, needs to be simple enough to already delineate obvious means which will promote their likelihood, or longevity, or whatever else is good in the context given, if the means to achieving this end is to be called good without any sort of serious issue. Once one complicates the ends in such a way as 'better for *us*' - where 'us' implies a whole species (or society) with a plethora of values and a seemingly endless possibility in terms of lifestyle that it would be near impossible to discover a standard for measuring - we are left with more questions than answers. The main question being the one at the heart of Thomson's account: good for us in what way? And those with different senses of justice and so on will employ different definitions which may or may not allow for the killing of Bert. The very definition of justice that Thomson puts forward ('proneness to doing what one owes to others') is normatively vague enough to allow various formulations, any of which may or may not allow for the death of Bert under various conditions. Unless Thomson's quest was a linguistic one, which aimed to show how 'good' (and adjectives such as 'just' which contain 'good' implicitly) function, by acting as a vessel for praise from a specific but undisclosed substantive normative view, rather than make any substantive claims of its own, then on this sort of reading it fails.

Interestingly, it seems like on this reading the account is close to a utilitarian virtue ethic. It is similar to what was discussed when I aimed to show that any ultimate-principle we attempted to put in to regulate further principles of action would need padding out from the evidence and values that we gain from living as functioning members within a community and which really provided the content for that ultimate-principle, which in many ways acted only as a place holder. Thomson's account can easily be seen as an expression of this, which would make the 'better for us' a lot more understandable. It would be saying that, with a wide enough scope to see, it brings a lot more 'happiness' if we act according to a very general understanding of

the virtues (one where Justice did indeed decree that Bert must live – and one which would be guided a lot more by our actual world than an idealised act utilitarian's world). Happiness, here, need not be reducible to any one state but can be informed by our everyday understanding and learning about values at a base level. This would again require a bidirectional process where the knowledge we gain from everyday life and the values that we see exhibited in living such lives act as content and evidence for an ultimate-principle that is constantly changeable by way of this sort of evidence. As such, the ultimate-principle stops fulfilling the function that what we normally think of an ultimate-principle as fulfilling, as it can no longer be seen as the final point of reasoning, or a super reason. Instead it would be a modifiable, malleable and changeable understanding of what's important in life, one that is being added to and altered by an understanding that takes place at levels lower than itself and specified in various ways to fit the situational features it is meant to be ruling upon.

While I agree that the above is something very close to the form that moral reasoning does take I do not believe that Thomson has a right to help herself to this picture. The way in which it is 'better for us' that she is attempting to grasp feels vaguer than this picture, if we are to take her argument as meaning 'better than various other possible ways of setting up a commensuration system' and the fact value gap is to remain open. On this reading, the idea that there is a natural slant to something being good for us – in the same way it is good that a watch be accurate, if the person is wanting a watch that acts in this way rather than merely being a showpiece of one's wealth – fades away. Instead it is a straight up statement of what is important in moral thinking that gives precedence to community formation as a feature for comparison, without much argument to that end. The meaning of terms like Justice would be required to rise up naturally from the living of everyday life, rather than being handed down legislatively from a set understanding that is handed to us by looking at merely the natural, empirical facts of living that life.

The reading I think Thomson is actually aiming for is more naturalist than this; the clue in the above quote being the move from 'better for us' to 'essential [for us]'. I believe she is very much aiming for judgements that any certain thing is 'good' to be read off in such a way as for them to be equivalent to any other natural facts about us; say, 'it is essential for us that we do not all eat poisonous plants'. Morality, as a

practice, in this guise, serves a purpose or has a shared prudential end to which it is put which is definitively decided by our ability to get things done as a social animal – which is something that we just are. McDowell (McDowell, 1996), amongst others, has argued against this sort of naturalist attempt to ground moral statements and I believe those concerns, regarding the idea that a conscious being is one that can purposefully and wilfully subvert any sort of nature that someone attempts to say they must have, in such a way as to make the naturalist argument redundant, are relevant. However, I would like more to focus on an interesting point regarding moral motivation for such naturalist accounts.

3.5 Motivation and Externalism:

This point has more than one foothold in our end account of moral reasoning. It is important in coming to understand how we should view ‘goodness’ as connecting with reasoning, how we should view the ‘make up’ of claims that seem distinctively moral, and the importance of ‘motivational features’ in any account of moral judgements.

Thomson’s naturalism – indeed all naturalist accounts share this weakness - puts her in a position of having to reject the idea that one needs to be motivated to follow a moral belief if they hold that belief. If it turns out to be the case that the best way of grounding the moral ‘good’ is for it to be actually be entangled within a naturalist understanding of praise for a particular thing which is inbuilt to its function (within a specific context or inherently) then when the object of that praise is meant to serve that particular function and it does so well - demonstrably well, in an empirical sense - then it would appear that the acceptance of the claim ‘X is a good act’ need not be intertwined with a motivational aspect in that person’s make up; one that would have them believe that having found it good they ought to perform it. An agent can correctly state that any number of acts are awful and feel no sort of remorse for doing them. It is just another fact that the act is bad, one that would only contingently impact upon the person who makes the statement.

Such a position I will label as ‘externalist’ here. Despite its various uses within philosophy today, I wish to utilise it to signify any position that makes the case that a

sincere moral claim made by an agent need have no necessary link with an agent's motivations. To clarify, this is not a case of 'defeasible motivation', where the agent has some motivational pull towards the moral claim that is outweighed by a stronger claim to act towards, say, their own rational self-interest. Externalists hold that an agent need have no sort of motivational pull at all towards a moral claim to be able to still recognise the legitimacy of that moral claim and make further claims themselves based on the remaining features of the situation at hand.

This is of course not to say that this must be the case. The chances are that even if externalism turns out to be true, people's motivations will indeed line up with the sort of judgements that they deem to be correct. That the objects of approbation and disapprobation (to use Humean language) common to a community in which one is raised will have the requisite attitudinative psychological effects on an agent is hardly going to be a surprising feature of coming to judge them as good or bad (or indeed, just/unjust, honourable/dishonourable,.../...). But for a naturalist, like Thomson, it is essential that this attitudinative aspect is an *unnecessary* extra within the whole enterprise of moral judgements. By the naturalist's lights, we judge from the basis of natural fact only, the psychological component of having a feeling that aligns in the appropriate way with that judgement is more or less a serendipitous outcome of the whole affair.

Many externalist theorists (Svavarsdottir, 1999 – Brink, 1989) have indeed urged that an agent with no such attitudinative extra can still be making correct moral judgements. They refer to this agent as 'the amoralist' (Stocker, 1979) Thus, if externalism is true then it must be possible for the amoralist to be a conceptually probable character, and if externalism is true then naturalism has a strong claim for grounding moral facts. If internalism is true (the antithesis of externalism, whereby one is seen as requiring the appropriate motivational pull –however defeasible – if one is to be making a genuine moral judgement), however, the amoralist cannot be a conceptually coherent character and naturalism must also likewise be false. So there is much riding on the supposed judgements of this agent who can know that an act is wrong while feeling no compunction to act in accordance with that knowledge.

While this formulation of coming to understand something as good or bad certainly has the desired effect of evading the issues of the fact value gap, given that all

judgements would be made on the basis of facts and indeed turn out to be factual statements even if they have evaluative language within them, I will be disagreeing that this is the right way of viewing moral statements. The argument, that is somewhat well rehearsed these days, in favour of internalism, will take the form of striking at a central intuition we have about what it is to be making a moral statement in the first place. This intuition is that if one is not disposed motivationally to follow the judgements one is making then one is not really making that judgement in anything like a sincere manner. It is an interpretation of something that already exists in an authentic and sincere manner that they are mimicking.

To get this across we could well adapt the context of a different thought experiment of Macintyre's (112-3: 1981) and think of ourselves being placed in an unfamiliar territory: Hawaii, many decades before the breaking of the taboo rules by Kamehameha II. If we were to visit such a place and look at all the acts that the term 'taboo' were applied to, we might start seeing a pattern whereby we felt we could begin on our own to avoid those acts that were taboo and only take part in those which were not. In so doing, we may not be motivationally invested in the appropriate sense for these to be our moral compass. There would be a point where, as an outsider, the thought 'this act is not morally wrong, but it is taboo' would cross our minds, whereas to a Hawaiian the thoughts would be one and the same; 'morally wrong' would translate to 'taboo' seamlessly. However good we got at spotting the patterns that were being used by the locals when they were referring to taboo acts, it seems that we are not doing the same as the Hawaiian person is doing in avoiding these acts. The practical significance of avoiding the act is driven, in the tourist's case, by a propensity to avoid offending the locals - however silly she feels about not doing something which she sees as utterly permissible - whereas the locals avoid the acts precisely because they are taboo.

The judgements are, for the tourist, a matter of discovering a pattern by way of which she will avoid punishment or causing offense (the latter of which to her may well be a sincere moral judgement) or simply of anthropological interest to her in documenting the practices of these people. For the local, however, the avoidance of such acts is central to 'living well' or 'acting correctly'. If asked why they didn't commit an act the only explanation that need be given would be 'that it is taboo', whereas the tourist uttering these words would imply a caveat such as, '...and they see this as wrong',

'...and I don't want to undermine their customs', or '...and I do not wish to be punished', which would show how the practical significance of the judgement is separate to and distinct from that of the local in avoiding the act. Gibbard (2003) amongst others, has done much work to show how, without the underlying implication that a judgement necessarily entails that it holds practical significance for an agent who does not stand to the judgement in the way of an anthropologist, words like taboo - or in our community, right, wrong, good, bad, just (and so on) - seem to become empty words, or nonsense words (Gibbard, 2003: 137-159).

Of course, for the non-cognitivist, such as Gibbard, this outcome of the naturalist attempt to show that there are moral facts, by essentially pushing all value judgements into the area of the natural by showing a 'good' something is necessarily tied up with the natural understanding of that something, is somewhat of a coup. By our own analysis it seems we are back to the position whereby moral judgements require an 'attitudinative extra' to be added to a descriptive item (be it a thing, an action, a state of affairs, a relation). And this attitudinative extra, they will claim, does not belong to the thing under evaluation; rather, it belongs to the non-cognitive part of the agent who is making the judgement. The right to appeal to the natural element itself as making the judgement has been repealed, except in cases where the specification of the thing's end allows for certainty in the means being selected as the best means. Instrumental reasoning already allowed for this, however, and few non-cognitivist philosophers - be they neo-Humeans, emotivists or logical positivists - would have denied this position. It seems we must move beyond reasoning about *just* means if we are to give a solid account of moral reasoning which is in any way cognitive. It also seems as though naturalist positions, as I have defined them, cannot allow for this. On our analysis, naturalism dictates the ends an agent takes are either given non-cognitively - in which case the whole affair reduces back into instrumental reasoning, if our use of good is to hang on anything substantial - or they are chosen cognitively; but in such a way as to make that decision seem insincere by taking motivation out of the picture so as to require judgements be parasitical on the judgements of those who just so happen to have their motivations aligned with the judgements of a certain group at a certain time and place. We will now turn to non-cognitivism.

Section 4: Moral Reasoning, Rules and Motivation:

Recap and Preliminary Discussion

In this section I will aim to finally piece together a sketch of just how I believe we should see moral reasoning as occurring. The previous sections might well be seen as the project of beginning to dismantle, by way of casting doubt upon, some strongly held assumptions by theorists who have searched for a ground for morality that stands above any particular set of judgements. The motivating intuition seems to be that the epistemological clarity that a consistent, stable and general ground gives is a prerequisite for intelligibility. There have been various attempts to cash out this assumption: that the ground is self-evident by taking part in human existence such as to reveal either principles or a principle as being the very stuff of morality; that it can be attained by a priori reasoning; from a special sensory perception of non-natural properties; or from the natural world as it is anyway, once we have understood exactly how our evaluative language functions to pick out things that are beneficial for us to see as good, *qua* human, or *qua* societal member. The varied forms of attack on irrationalism and unintelligibility which the assumption represents, under various guises, are all genuine attempts to overcome positions which threaten our ability to say that evaluative claims are more than the expressions of mere biases and prejudices.

The assumption, however, has its flaws. This is made manifest by the much visited idea that each and every theory we have so far examined, while picking out at least one major and important essence of morality, by way of emphasising one touchstone that seems pre-philosophically important, does so at the expense of some further intuition. In some cases that is merely norms of content: to take on board the important touchstone of consequences is to give up on the integrity of persons and vice versa. In other cases it is norms of form: to take a Rossian position that allows both consequence and integrity of persons and emphasises the role of judgement, we have to sacrifice the idea of a coherent system and give over too much heavy lifting to intuition - for which we have no real understanding.

This section will be a dedicated attempt to find a way in which we might satisfy as many of these norms as possible. As such we will need to explain why they are so important to us in the first place, and hopefully give an analysis of the moral terrain that allows it to keep its pretensions for intelligibility, put reason at its centre and keep the idea that moral reasoning has a proper and appropriate subject matter, such that an agent taking part within it can be said to be better or worse at it. We will likewise want for them to be said to be going on in the right manner and in the same way when they are reaching a level of expertise in the subject matter, all without that claim seeming far-fetched or worryingly self-referencing. In so doing, not only must this theory of reasoning capture what is at essence with greater clarity and consistency than the theories which I have disregarded thus far, but to do so while evading the threat of irrationalism and subjectivity which lurks in the background.

It is from the ashes of these assumptions about intelligibility and objectivity that I believe a more defensible and realistic view of morality might well emerge; or at least one which might act as a spring board for a new direction in the application of ethical thought to practical fields.

The theories that we have looked at so far have always been trying to get a grasp of one key question: how might we know that our moral judgements are right? We can separate those that we have seen into rationalist and naturalists camps, whereby, respectively, any claim to knowledge is grounded within either a system of universal rules forced upon us by rationality, conceived in one way or another, or by reference to the way the world is anyway. The problems with positions from either camp have been manifold and strongly reported already within this thesis; however, a more in depth analysis, one which should apply to all examples of each type of theory, will be the first aim of this section. The majority of theories we have analysed so far have been rationalist. By this I simply mean to say that they ground morality, or see it as given some foundation, by way of reason itself. If the evaluative domain is to be more than a mixture of biases and non-cognitive feeling, if the domain is to have an actual chance at intelligibility, then there must be some claims which stand as universally true – that is, that the denial of which would simply be false - from which other claims can flow. We have seen different ways of going about building such theories from theorists such as Kant, Mill and Sidgwick.

In all these rationalist theories the predominant intuition is that intelligibility implies generality in one form or another, and as we have no reason to believe, ahead of any concrete cases, that any given agent is undeserving of being seen as morally equal, these generalities must therefore be seen as universal also. Thus, we arrive at a place where universal rules must really be the very objects of moral thought. What differs between them is the idea each theorist has about what constitutes the touchstone of morality, which acts to generalise the material. Some focus on happiness, others on integrity and autonomy of persons, but each, crucially, sees the proper foundation of moral thought to be made out via some rule or set of rules. Rules are the very essence of generality; as such it is rules which we must follow if we are to be lead in each case by the same ideas, to carry on, as it were, in the same way in each case without bias or prejudice, to aspire to universality and so on.

The very idea of rule following has had much philosophical attention given to it over the preceding century, and each type of theorist we have come across so far has a stake in this argument. The rationalist and the naturalist both interpret rules as the very foundation of moral thought and intelligibility. The non-cognitivist, as we will see, has a connection with the way that naturalist-based moral theorists take rules to be important. The non-cognitivists are against the idea that rules in anyway act as a saviour for rationalist theorists against their claims that one acts according to their non-cognitive desires. Whichever rule one chooses, on their reading, will only show that agent's personal non-rational preferences.

My own position will be what has come to be described as particularist (Garfield: 2000, Little: 2001, Dancy: 2001) It has its roots in Aristotelian philosophy. I will be explaining and defending this position as we move through the upcoming sections. At this very early stage though I need to make sure I distinguish it from a further position which shares its name. This other position states not only that moral claims are not grounded by rules (a position I agree with) but also that rules have no place in moral thinking whatsoever and moral claims are not universalisable. These last two points I stand firmly against. My own idea of how to represent and understand moral reasoning will in fact hold rules to be of great importance in explaining moral thought – what we need is to characterise how those rules function and what form they take. They may not be recognisable from a descriptive level, and it may take judgement to know when to utilise them, but the key point will remain that they are

there. Without some core of generality, intelligibility is under serious threat. Thus the point will largely be that rules are simply not foundational. While the particularist/generalist debate continues, I will maintain that it is a Wittgensteinian epistemology, to be spelled out soon, that best caters to the understanding of how rules are best characterised. This will be made clear by how much easier it is to explain that tricky idea of motivation in moral reasoning that we have already seen the naturalist has problems with. That is a problem that rationalism shares if it is to be made out in such a way that requires rules to be foundational to moral reasoning at all.

4.1.: Non-Cognitivism Re-Visited

Before we get to all that, however, I wish to first take another excursion back into the realm of non-cognitivism. It is through this area that I think we should proceed in order to get across all the points that this section needs to make, as it is via attacking this position that we might best get to grips with two points which themselves each have two facets. The first of these is what is needed for moral disagreement to occur and what is happening when it does so. The second is what it is for motivation to occur and what is happening when it does so.

The latter of our two points makes up the very meat of the non-cognitivist's appeal. After all, it seems that no cognitivist theory we have so far canvassed really comes close to explaining why it is we should be motivated to do some acts and not others. This inability to explain motivation to do good acts is the focal point of any non-cognitivist attack upon normative moral theories. That this is the case should hardly be surprising. Non-cognitivist descriptions of reasoning seem to have an air of simplicity that has an appeal largely in virtue of this simplicity that revolves around motivation. The process seems to be explained simply and thoroughly via the instrumentalist model of desire (end) and rationality (finding a means to that end). Furthermore, in decreeing that such judgements can never be right and wrong, the picture of practical reasoning does not need to be compared with our more everyday understanding of theoretical and descriptive reasoning. The Humean picture of World and Mind remains intact, with the contrast between the two types of reasoning merely being presumed. This is a very important point. What any cognitivist is up

against is an entire Empiricist picture which explains where all of these sorts of judgements ought to sit. Again, the division between mind-to-world and world-to-mind directions of fit is key. If the agent is implicated at all as being part of the source of some mental item they experience, then according to this picture the charge of subjectivity is a likely consequence. Only those items which present themselves to our sense as part of the world 'as it is anyway' are actually objective. And moral claims seem, on the cognitivist rendering, to try to have it both ways. The cognitivist seems to wish to say that moral reasons are objective, but that moral reasons are also born of an agent's will; the right type of will. Thus, it is usurping this Empiricist picture that it seems to me is the main task of cognitivists if we are to overcome this obstacle.

On the other hand, the non-cognitivism, which takes this empiricist rendering as its lead, seems to have a hard time in explaining why it is that normative moral disagreement occurs at all. This seems to be a serious issue for the theory as a whole simply because it is very clear indeed that disagreement over the use of moral terms does happen all the time, and happens in terms of distinctive descriptive content of a term. While we may have some almost jovial arguments over our favourite flavours of ice cream, we do not do so over the application of moral terms to other actual events and features within the world. Unlike chocolate being a good flavour and mint being bad (something which, although true to any sensible person regarding ice creams, has no further court of appeal than one's own sense data and non-cognitive response to flavour and taste) it would appear that disputes about wars, relations, deceitfulness, the moral status of different beings and so on, move beyond this. They 'move beyond' in as much as any position seems to require an explanation or justification as to why one might hold that position over an alternative position. In so explaining, we utilise the practise of reason-giving, moral assessment and argument which we see every day. Thus, if the non-cognitivist is to allow such features she must have a more complex explanation of what is happening than that provided by the Emotivist account - that we simply react to a situation and that is all there is to it. Emotional and evaluative statements, by the emotivist's lights, are not apt for truth at all, as they are in no way a part of the world as it is anyway; explained by science and without any credence given to our contingent responses. Here, there is no explanation for why we utilise reasons to explain our difference of opinions

upon something. Normative disagreement, something which should be able to be explained away by any theory, becomes a mystery in itself. Without some regularity underscoring our application of evaluative concepts we would forever just be speaking past each other.

The bulk of this section will be trying to overturn the key elements of this non-cognitivist picture as the seeming alternative to there being a solid generalist theory which somehow compels acceptance. Having considered and rejected the idea that any generalist theory compels adherence, it seems that the following picture is what is left. It turns upon the components of what we do feel we can know for sure, and then draws an unfavourable comparison in the light of ethics.

The picture of what believe we can just know falls into that which can be explained 'descriptively'; that which has a world to mind fit. The world spreads itself onto our minds. From sensory input we gather empirical data. That which leaves a sense impression upon us is that which we can know is there. For sense impressions we cannot immediately understand – what causes the movements of inanimate objects, say – we can theorise about what causes the impression of what is there anyway, and test them against basic standards for understanding. Anything additional to this, anything that we believe that we project upon the world, ought to be accounted for in terms of this interaction between human mind and the world, seen as two distinct components of that natural world. From this we can talk of subjective experiences.

If facts, truth and objectivity are to fit into this picture, however, it is not to do with what is projected, but the actual circumstances under which it is projected, what 'inner' elements causes the projections, how we learn to vocalise them through social institutions of language and what-have-you. We can give objective and factual accounts of such things as colour experience, a staunch empiricist would say, but for all that colours are not 'there anyway'. They are a by-product of our very particular subjective make up interacting with other objective elements of the world. For such empiricists, our reasons belong to describing and theorising about what is there anyway. Evident here, in the background, is an idea of an 'absolute conception' of a world fully described by empirical data, where every primary and secondary quality was explained or explained away reductively in primary terms. The external world is just the external world, and we, as physical beings, are a part of it, and all our ideas

and actions and thoughts and experiences can thus be explained in these primary terms. There is an appealing sense of simplicity and unity within this picture, and indeed this picture is very much the functioning one that scientific enterprise embraces. Truth, objectivity and reality are determined by what is 'out there anyway' on this rendering. Anything which does not fall into that category can be downgraded to subjective, a mere projection of ours. Thus, in ethical terms, if someone wonders whether you ought to do something, the 'ought', by these lights, is not up for reasoning about, without the ends of the agent involved being known. Whereas for meaningful ethical discussion it seems a pre-requisite that ends and first principles (if there are such a thing) are open to be assessed by rational decree. We need to find a way to overturn this picture, which I will refer to as the Empiricist Theory of Mind. I will discuss the idea of absolute conceptions in Section 5 in relation to moral objectivity

The first part of this section will discuss a more sophisticated non-cognitivism than that of the emotivist. This opponent to moral reasoning will be pitted against the arguments of Neo-Aristotelians who have utilised a Wittgensteinian epistemology; with the focus being on the work of John McDowell. The upcoming section will lay out the key ideas of our antagonist. The following one will give some programmatic and intuitive remarks as to why their analysis seems wrong. Neo-Aristotelians often bring charges from a Wittgensteinian inspired epistemology regarding rule-following in order to make their points. The upshot of this account will be that those who peddle this sophisticated non-cognitivism think that the descriptive elements of what we do can be 'disentangled' from the supposedly projected 'shared evaluative commitments' of communities in which we live. Following McDowell, I aim to show that this description of understanding is indeed counter-intuitive and destructive, not just for moral reasoning, but to our understanding of all reasoning. We will also deal with the charge that not relying upon descriptive generality catapults us into a regress of interpretations. This last charge is serious. If we are not to illicitly take up a rationalist or naturalist ground of our own and yet keep open a space for normativity and objectivity this claim will have to make good on this point.

4.2: Sophisticated Non-cognitivism and Rules

What I think of as a more sophisticated non-cognitivism, and one which has a lot more clout to it by way of explaining normative disagreement, derides its earlier counterpart and rejects the thesis that we simply see a certain state of descriptive affairs and either find it agreeable or disagreeable. Rather, this version runs, when we see a state of affairs we judge it to be good or bad upon the *descriptive features* that are there - and because of the features that are there - coupled with a learned reaction to those descriptive features obtaining in just this way. We do, as a psychological fact, analyse and reason upon whether we should categorise these states of affairs as good or bad. That there is at least some conceptual agreement is evident from the fact that we do not merely talk past each other. This account thus has the benefit of accounting for moral disagreement, and the ways in which using a term in a particular context can be seen as a good or bad use of that term. This section is indebted to the work of McDowell, especially in 'Non-Cognitivism and Rule Following' (McDowell, 1981)

While this, at first glance seems to be a description of a cognitive account, it is not. For what we are actually doing when we make this move is utilising moral concepts that can be said to have two distinct elements: a descriptive element and an evaluative element. Our reasoning is, the sophisticated non-cognitivist will claim, purely based upon the descriptive element; that is, the features that are present in this case and how they ought to be categorised into evaluative language. Whether an act falls under the terminology of courageous or foolhardy, for example, is open to debate and discussion. Whether it falls under the concept courageous or minty simply is not. It is this shared core of description with the addition of different attitudes and sensibilities in individuals that allows for the disagreement that we see day in and day out. Thus, discussions about wars and the such are rarely a matter of whether it is good or bad – thin descriptions - but whether it is just, whether (in a democratic country at least) the population of the country who attacks is aware of why they are attacking and agrees with these reasons, whether the force used is uncalled for, whether the intent of the leaders is not duplicitous, whether the acts of soldiers are honourable... the list goes on.

One could spend many hours looking at all the features of a war and come to a very reasonable conclusion regarding what parts of it were just and which parts unjust, the consequences it had for both countries, the probable intentions of the leaders and so on. But, by the non-cognitivist's lights, at this point one would still not be taking part in proper moral reasoning. All one would be doing is categorising features that are 'actually there' into moral categories that could be utilised via description satisfaction alone. If there were many cases of one's own army risking their lives in the face of danger to help others this would show many cases of courage. If the leaders who called for the war seemed sincerely to care about the ends that they claimed they went to war for this would show good intent and honesty. If the war was conducted with expertise, using technology to reach one's end while causing minimal casualties on both sides, this would point to an attempt to keep the war as humane as possible, and so on. All these thick features (courage, honesty, good intentions, justice, credibility, humanity) seem to show that we have reasoned that the war is, as a matter of *fact*, at least permissible. But to the non-cognitivist this is not so. All it shows is that certain features obtained within the war that *allowed us* to categorise, as we are prone to do, the war as having those features, *descriptively* speaking. There is room for the rightness and wrongness of concept application within this account then – the concept is correctly applied if just certain descriptive features that obtained allow for that concept to be applied. But the addendum that these features speak to an evaluation is wrong. The evaluative element is still non-cognitive for it relies on an attitude that has a definite mind to world direction of fit.

As far as moral evaluation goes, the evaluative element of such thick concepts, as those listed in the sentence before last, still requires an 'attitudinative extra' that shows that no such judgement that this concept applies is true in the full moral sense. While we can say, without worry, that descriptively speaking, the war had all those features, we cannot say that the war was *actually* better or worse for having them. For while, descriptively speaking, the soldiers were honourable to their captives, enemies and allies alike (by way of the relative commitments to concept application within that community) the evaluative addition, that this is a good thing, is one that does not appear in the world alongside those features, but is added on erroneously after the fact and in no way follows from those facts. It is, the non-cognitivist will claim, projected from a very special position of concerns (ours) onto

the world itself. And concepts that are so projected, as opposed to having their basis in the factual domain of the empirical world, belong instead to a particular person or group's psychological state at the time and their customs for applying such concepts and generally being emotionally effected by that which warrants their applications. So while we can say of a war, that it was fought honourably, courageously, fairly and with minimal force by the victor, and we can also say that these features definitively obtained – according to the standards of that group - and we can say that most people who view this war would find those features to be there and to find those features to add up to the idea that the war was won well, we *cannot say the war was actually won well*. The people of this culture applied the concepts correctly upon the basis of an agreed upon definition of those concepts, of what applying 'courage' here amounts to. But their motivation for application and the use of such categories is not held in place by the world as it is anyway. All that can do is account for regularity relative to that community's practices and customs. Those practices and customs do not, however, relate to the world objectively; that would constitute that from a purely descriptive position, there would be evaluative properties in the world as it is anyway.

This account, then, focuses on moral language and the concepts and predicates used in its prescriptions and judgements. There is agreement here between non-cognitivists and my own thesis regarding the way certain types of moral concepts function. Terms such as 'good', 'bad', 'right' and 'wrong' are what we call 'thin' moral concepts. These concepts are 'verdictive' (Dancy 2004: 16) judgements that make up the top layer of moral reasoning. What causes us to make such judgements is forced by the layer of moral concepts which lay underneath these; what have been termed by Williams (1985) and many others as 'thick concepts'. Such concepts are said to have both a *descriptive* and an *evaluative* aspect to them. Concepts such as courageous, delicate, sarcastic, vindictive, petty, egregious, valorous, kind, and so on, all fit this description. The move common to the non-cognitivist tradition is to claim that whilst we can, in some respects, say that these are concepts which have truth conditions for their application embedded in their descriptive element - which accounts for normative disagreement - their evaluative element is still something that cannot be said to be a matter of cognition.

Take the term 'courageous'. The non-cognitivist idea here is that the term can, as McDowell puts it, can be 'disentangled into two components' (McDowell, 1981: 200-

201); on the one hand its descriptive content and on the other its evaluative content. Descriptively, the non-cognitivist would claim, 'courageous' acts are those which can be properly described as something like this: 'to begin or continue a course of action despite having to endure foreseen negative consequences in the face of conflict'. As human beings, who as a matter of fact, apply moral concepts, courage has a descriptive element. Courage becomes a moral concept in that we also 'tack on' on an evaluative element, to this effect: 'and such acts are praiseworthy'. And it is this *evaluative* element of moral concepts which the non-cognitivist sees as being a matter outside of cognition; 'a non-cognitive state that constitutes the special perspective from which the items in the world seem to be endowed with the value in question' (ibid: 221). It is an extraneous evaluative component which cannot be said to be apt for truth outside of the practice in which the concept is applied. We can teach others to apply the concept more or less correctly by tracking its descriptive elements, but we cannot say they have reason to find its being so applicable as good or bad from there.

The importance of separating what is supplied by our practices and what is supplied by the descriptive element is crucial to the account. The picture which is allowing for something to be considered a matter of correct application – that which allows for the normative account of moral concepts that the more sophisticated account of non-cognitivism seeks to keep - is the picture of a purely natural description of the world; an 'absolute conception' or full scientific explanation of what is actually there anyway, regardless of our practices or our own unique perspective of the world. Thus, courage can harmlessly be said to be a concept which is independent of our particular perspective, in that it picks out acts which can be brought together under a basic description, that human beings can, as a matter of fact apply: those which follow something like the rough sketch in the previous paragraph ('to begin or continue a course of action...in the face of conflict'). What confuses the issue - and is a matter which cannot be said to be true, independent of our unique perspective - is that these acts should be in any way commended as being *good* acts. This would require that moral features such as goodness were themselves 'in the world'. And by the non-cognitivist's lights the world is made up only of natural, descriptive features; those features that have a world to mind fit, whereas the evaluative addition has a mind to world fit.

Thus the normativity of moral statements, according to the non-cognitivist, which allows for proper disagreement, is to be found in the descriptive element only. When we meaningfully argue about moral statements we are taking what I will call an 'anthropological' stance towards them, in that we are saying that these people, at these times, would be said to describe act A as courageous because it exhibits the descriptive features X, Y and Z. This disinterested anthropological position is the only one where truth comes into the equation. From the position that appears meaningful to us, the one endowed by the special perspective, is in fact meaningless, because it is irrational. We can happily add, regarding the people who believe they are occupying that special perspective - as a naturalistic description of their mental states - that they would also believe this act to be one which warranted commendation, or would find this application agreeable and not that one; that they would derive happiness from such acts. But we can never say that the act is one which should *as a matter of fact* be commended. And this position holds as true for our own practices as for the practices of any tribe, community or set of persons at any time and in any place. The upshot of this account is that the meaning that we supply to words which carry evaluative content is purely relative to a particular time or culture. It is the evaluative elements that make us establish such concepts, by projecting ourselves and our attitudes onto the world. As any attitude is as meaningless as any other, from the point of view of the disinterested universe - which determines true propositions - correct or incorrect usage cannot be conceived as being how the term is applied here and now in the full sense in which we hear it. Thus, to explain away normativity and moral disagreement, the non-cognitivist places the burden of this with the descriptive content. When we start calling certain things courageous, we are, as it were, contractually obliged to maintain the same uses in the future, with that use being made out at the descriptive level. If we did not do this there would be no locus of proper disagreement, because there would be no means for genuine disagreement left; just a cacophony of variable interpretations with nothing to link them. Such a move thus allows the sophisticated non-cognitivist to overcome the threat of failing to explain the nuance of disagreement, which was the threat which does seem to incapacitate the Emotivist account. We will deal with the threat of interpretability shortly.

Overcoming this particular non-cognitivist threat has been paramount in not just re-establishing moral thinking, but in ushering in a particular way of seeing just how moral reasons work and how they can be said to truly be a matter of cognition. I will follow the reasoning of McDowell (1984) here in taking up a Wittgensteinian epistemology. Doing so correctly threatens the very idea that these non-cognitivists can indeed account for the normativity of moral statements as they wish to. This leaves two options: either fall back into the Emotivist position, which cannot account for normativity, or accept that normativity can be accounted for in more ways than merely pointing to generalities at the purely natural level. Just how much of this attack McDowell gets right, and where his account takes the line of reasoning will be put under scrutiny in Section 5. There we will see that there are some possibilities left for the non-cognitivist even after the following argument has given the cognitivist some breathing room. As the account develops it will also become clear that McDowell utilises certain Wittgensteinian insights that many other cognitivists play down or reject.

4.2.1: Denying the Disentangling:

The non-cognitivist's error is in many ways best brought out by pointing out that thick concepts actually belie a purely descriptive element that is not actually there at all to be understood by someone outside the practices. Mastering a concept such as courage cannot be done by means of description satisfaction alone, which could determine its correct usage by a competent agent in the anthropological sense implied by the non-cognitivist account. Whatever the description given, it appears that purely descriptive content cannot be easily used to capture such a term's correct usage in any particular scenario. In our own description of courage, for example, the terms 'endure', 'negative' and 'conflict', which were used to attempt to reduce courage to its purely descriptive element, are, in themselves, not purely descriptive. Furthermore, scenarios in which the 'more descriptive' features seem to fit are not always ones in which a competent user would want to say that 'courage' is an appropriate evaluation.

Someone who 'endures' the foreseen, 'negative' consequence of 'conflict' when they publicly harm an innocent person are not people whom we would happily describe as

displaying courage. All the same, publicly harming innocent people is an act which will generally cause one to endure negative, foreseen consequences. Surely no one would see this as courageous, even though it fits our description. Thus there is a thought such as possibly adding something like 'for good reason' to the description of courage. Yet what counts as a good reason seems to surpass being reduced to descriptive content. As does something like, 'without being cruel, thoughtless, callous...' as each of these other thick concepts would surely suffer from the vast amount of possible misinterpretations that courage does. (Compare: Croom, 2010). If to apply courage we look for certain descriptive qualities obtaining and then tack on the attitudinative extra, 'and this is a good thing', we cannot appeal to the thought 'is this (overall) a good thing?' to decide whether or not the concept applies. It is whether the concept applies that is supposed to tell us this. As Dancy says, 'verdictive judgements do not contribute to the situations on which they pass judgement'. (Dancy, 2004: 16) Equally as unedifying is the idea that we know all the descriptions in purely descriptive terms of what counts as courageous already. Yet this list might be infinite, and certainly beyond human cognitive reckoning. Understanding thick concepts such that one can apply them correctly seems to require we share a whole host of other evaluative concerns.

This seems to undermine the anthropological position the non-cognitivist wants us to occupy, which they claim explains the normativity of moral concept application. Failing their ability to show that the normativity of thick concepts can be understood from 'outside' the practice, as examples of picking out the right set of descriptive features, it stands to reason that genuine disagreement over such concepts can only be explained by someone who is already within the practice (something they would claim threatens the objectivity of the account as it resists explanation in natural terms), or that the whole thing gives merely an illusion of reasoning (something early emotivists would agree with).

It is perhaps worth pointing out at this point that non-cognitivism, at least the more sophisticated examples, shares this flaw of needing reduction of the evaluative to the descriptive with the naturalist accounts we pondered earlier, with Thomson. The non-cognitivist account seems to require a position from which we can view thick concepts which allows for them to be reduced to a matter of description satisfaction. Naturalists, of course, share a similar position. They wish to claim that someone

being, say, just, is good as a matter of empirical fact, and both the goodness of that deed as well as the deed itself need to be understood in descriptive terms - if they are not to undermine their own account by assuming an evaluative understanding that is grounded upon anything other than the world as it is anyway. Of course, given the different aims of the two positions, and where they start from, utilising this idea has different effects for each account. In the non-cognitivist account, as we are starting to see, there is an issue with the idea that, unless we are understanding morality from within, being able to recognise the requisite descriptive patterns (to identify a thick concept's correct application) as patterns that hold each time without issue, or resorting to further evaluative concepts, seems unlikely. While there are indeed similarities between all acts that can be labelled as courageous or all acts that can be labelled as honourable or just (and so on), there are always examples that either simply evade these precise descriptive features in which a competent user would still find it appropriate to apply the thick concept, or examples which have these precise features, even though a competent user would not wish to apply the concept.

The naturalist account shares this flaw - this idea that we may well be able to show our thick concepts as supervening upon descriptive features in such a way as to remain static and recognisable from outside our practices - whilst also having the misfortune of not providing any notion of how one becomes motivated to be courageous just from coming to know that it is 'good for us' to have these traits. Even if one were to grant to the naturalist that there are definitive patterns underlying each thick concept that allowed for them to be cashed out as a matter of description satisfaction alone, or a naturalist was to attempt to ground the sort of epistemology I am utilising here to explain moral reasoning in a naturalist ontology, she would not be able to explain why this should *motivate* her in her actions. That is, were it obvious that doing action A at the descriptive level would mean her act could be labelled as courageous, whereas action B would make her act, by way of supervenience, as being labelled cowardly, this in no way gives her any motivating-reason for doing act A over act B.

What is a flaw in this naturalist account is the very area the sophisticated non-cognitivist exploits to strengthen their own claims. While we can recognise acts at the descriptive level – allowing as they need to for moral disagreement, moral

education and moral assessment – the additional aspect, that these acts are to be avoided, or promoted, or done at all costs, and so on, is not entailed by that recognition. It is my contention, following McDowell (1981) (as well as others), that this ‘disentangling’ of thick concepts – into one of recognition and one of motivation to action – is not available to us. To be available to us in this manner would require us to be able to step outside our own modes of evaluative thinking and be explicable to anyone who shares none of our interest in our modes of evaluation.

This point has, of course, been used by McDowell, amongst others, in defence of a virtue based account. I will also take this route, but I wish to give a clearer account of what is happening when we are reasoning, building upon the idea of bi-directionality available in holism to show that argument can be made out. While my account is broadly Aristotelian it makes less use of virtues as such, focussing on how values, empirical facts and norms interact. If we do this I believe we can actually give a lucid account of reasoning in general - moral reasoning included - which will in turn allow us to start assessing reasoning about ethics in the medical field well-armed, with an idea of morality that has a strong epistemology, explains away ontology, explains motivation and gives us an idea about how moral claims can be seen as an exercise in truth and objectivity.

4.2.2: The Place of Rules in Reasoning:

The non-cognitivist concern here, regarding objectivity, easily transfers over to the conception of how rules function in our thinking, a topic which Wittgenstein was famous for focussing on and elucidating in such a way that allowed for the McDowellian ‘anti-non-cognitivism’ account to emerge. Everything, Wittgenstein showed us, depends upon the position from which we account for the generality that is manifested by rules. The non-cognitivist thought is that if the rules that fix terms such as ‘courageous’ are not understandable from the anthropological viewpoint, then the rules don’t fix in place anything at all; there would be no point of reference by which to say a concept has been applied correctly or incorrectly. But this is an adherence to the scientific conception of objectivity, where the only things that can be said to ‘really’ be there are those things which can be understood from outside of specific human practices. And this adherence is, as Wittgenstein has pointed out,

suspect and most likely misplaced. Our perspective, he demonstrated throughout his later works, makes far more appearances in all of our areas of reasoning than the Empiricist rendering of the relationship of mind to world allows for. I wish to show that morality has only a few features that are peculiar when compared to science or mathematics. Each domain hangs on adhering to a practice that cannot be defended from a position outside itself or that is defensible *a priori*. The features morality exhibits which are not found in other domains of thought, I will endeavour to show, are not enough to disqualify it as a domain which can be understood as objective; instead these differences are precisely the sort of *suis generis* features which make it an independent domain of thought and which provide it with a unique and distinct content. Without such differences one would struggle to distinguish its content from purely descriptive content. Thus, to take such differences to act in favour of believing less contestable domains as granting access to truths, in a way that morality does not, is to already have accepted, implicitly or explicitly, that the contestable Empiricist position to be correct.

To do so I will follow Neo-Aristotelians like McDowell, Garfield (2000) and Little (2001) and utilise these important remarks from Wittgenstein's later philosophy regarding rule following; ones which challenge that very position by way of exposing the idea that certain data comes to us uncoloured by our assessment of them from a specifically human position. Having used this position against the non-cognitivist arguments - that facts and values should be conceived of in such a way that interprets the latter as including a move that precludes them as being fit for cognition - I will move on to show how it likewise debunks the idea that intelligibility requires generality at the level of storable general principles. This should undercut the intuitive adherence to rationalism, for those who wish to show that the moral domain is indeed still intelligible, without concession to descriptive generality or system based rationality which runs from a touchstone definition of the right or the good. The implication, here, will be that the same theory of mind has been accepted by those who see a certain conception of universal rules as being a necessary feature of defending moral reasoning as being cognitive as those who argue the whole matter is non-cognitive. The next section will address a way of conceiving of motivation which is amicable to this same position.

We will, then, begin with our assessment of what it really means to follow a rule. I will follow Garfield (2000) and use rules outside of morality to make my point. Take these two rules from the world of soccer:

- 1) Fouling a player in your own penalty box by tripping or pushing an opponent prior to taking the football from them will result in a penalty kick for the opposing team
- 2) If, whilst committing such a foul, the player also prevents a *clear goal-scoring opportunity* then that player should be shown a red card.

Whilst the two rules both direct game officials in decisions as to what should happen in the event of either situation occurring, there is a significant difference in how each rule is learned and applied which is also the crux of the matter regarding moral rules.

Regarding Rule 1: The rule can be learned by reference to description satisfaction alone, and its application in any given instance be deemed as correct or incorrect on this basis. No examples need be given to show how such a rule works. Instead, one can simply learn the descriptive formula, 'a penalty kick should be given when a defending player trips an opposing player in the penalty area prior to kicking the ball' and then go on to apply it to future cases without much issue. One can get better at applying the rule to a point. But it is the rule, not the judgement that its application is relevant here, which appears to be *foundational* in determining it has been used correctly.

Regarding Rule 2: The rule cannot be learned in such a way that it might be applied correctly *without* examples of its correct use, via inculcation into the world of football officiating, being given to the person learning it. What counts as a 'clear goal-scoring opportunity' calls for judgement beyond the scope of description satisfaction alone. Instead, one can only learn from seeing various cases of goals scored from various angles and positions, all of which depend on a myriad of factors, including but not exhausted by such things as: ball control at the time of the foul, the supposed shot back-lift needed to score, the skill of the individual in question, the angle and

distance from the goal they are positioned at when fouled, whether the severity of the contact constitutes a trip, and so on.

How we come to grasp and understand rules so we can apply them to future cases then is paramount to the non-cognitivist account. Being inculcated into a practice such as soccer officiating is, for the non-cognitivist, insufficient as an explanation, as it leaves open the threat of subjectivity and arbitrary application of the rule. That is, the non-cognitivist thought goes, if we refer only to how we *do* apply a concept or rule to *justify* future application, then we lose any account of how we can be *wrong* in that application. If we are to apply a moral term such as courage correctly in each case, its application must be fixed from the anthropological position and conform to something more like the example of Rule 1; where a series of counterfactuals are insufficient as the only explanatory device for saying someone is going on in the same way when applying the rule or concept. But we have seen that, at first blush, evaluative language evades precisely this sort of description satisfaction. Moral rules, if there are any at all, as we have seen, resemble Rule 2. They resemble it, that is, in that attempting to give a purely descriptive explanation of what acts, states of affair or agents a concept applies to seems impossible. We seem to require further thick concepts in our explanations of the correct application of any particular thick concepts. But if this is the case the non-cognitivist position is under severe threat. The epistemological hook – the generality of the natural world as revealed to us by our senses – cannot be thought to account for the normativity of concept application.

The non-cognitivist response to this is that, rather than learning mere responses to certain factors through a lengthy enculturation which requires counterfactuals and prompts (think of Rule 2), we can postulate a ‘psychological mechanism’ (McDowell, 1978: 206) which allows us to grasp patterns which are, in the descriptive sense, actually there. These patterns, however complex, once divined from the examples we are given, then lead us to make more or less or correct applications of the rule or the concepts the rule determines (depending how well we have grasped the pattern) whenever that pattern (or if we have not grasped it thoroughly, similar patterns) appear. This postulation of the psychological mechanism, then, alleviates the threat of apparent arbitrary application that the counterfactual learning needed for grasping rules like Rule 2 seems to throw up. For example: we cannot describe fully the

pattern which accounts for the phrasing 'clear goal scoring opportunity' – this is just an anomaly of our linguistic abilities – but this complexity does not necessitate that such a descriptive pattern does not exist and ground correct applications. The psychological mechanism can do the work that our linguistic minds cannot and tacitly find the descriptive patterns underlying correct applications. We could, theoretically, with precise language, describe every example of a 'clear goal scoring opportunity', but we use this phrase as short hand instead. Thus, with enough care and time, we could indeed explain the term to someone who takes up an anthropological position to soccer officiating. If we cannot do so, we are using the term arbitrarily. As such the concept is not open to normativity, as what is a right or wrong application of it has no grounds.

But this appeal to the 'psychological mechanism' has been shown by Wittgenstein as being misguided. And without such a mechanism to help us grasp rules from the anthropological position we also have reason to doubt that our use of rules can be grounded as such. Here, the mathematical examples readily available in Wittgenstein (1953: §198), McDowell (1978: 205) and Kripke (1982: 9) are edifying in helping us come to see how even rules which we believe to be completely separated from our being involved in the practice which employs them have the same dependence upon our responses as those in practices such as soccer. We like to think of mathematics as existing anyway; as its rules are understandable platonically, from a position outside that of being an agent who practices arithmetic. As such, the consoling thought goes, the rules which govern this practice, at least, are not open for interpretation in the way that Rule 2, or a given thick concept seems to be, and from our preliminary discussion, moral rules seem to be. As such, surely this is a prime example of the mechanism at work.

Kripke (ibid) is quick to point out that even the rules of arithmetic seem to be open to vast interpretation, if we hold onto this purely objective position from which similarities need to be grounded. For example, Kripke offers the example of someone who interprets the plus rule as the '*quus*' rule. *Quus* states that: for any numbers <57 when you see the + sign then add those numbers together. For any

number >57 , when you see the + sign the answer will be 5. Whilst a child learns the rule, we believe initially that they have indeed grasped what we believe to be the plus rule. But once numbers larger than 57 start appearing, we become confused when he starts reeling off the answer 5 to each problem, and even more confused when he states with certainty, 'but I know that I have the right rule in mind'.

Furthermore, the + sign could be interpreted in infinite further ways, meaning that the plus rule could really be identifiable with an infinite range of behaviours. Thus, until an agent has shown us he can add every conceivable set of numbers, we can never be certain that he has in any sense 'grasped' the rule. But a certainty in the grasping of such rules is precisely the thing the non-cognitivist account assumes. By their lights, rules cannot be grounded by the responses of agents who are following the rules of the practice in question. This would allow for the arbitrariness of application they need to avoid so as to say that we are really going on in the same way at the descriptive level – that at least *this* is a matter of cognition. The same worry ought to be shared about anyone who leans towards an empiricist theory of mind.

The non-cognitivist may well respond again, at this point, that Kripke's argument - that a whole myriad of potential rules of the 'quus' variety could well be what is interpreted by someone learning the rule - is somewhat beside the point. One such rule will in fact be *the* plus rule as we know it, and that teaching rules such as this just consists in continuing the teaching until that person has shown that they really have grasped the rule by answering a variety of examples and always in line with the actual rule they wish to say that he has grasped. When this has happened, then the non-cognitivist can still say with some certainty that the 'psychic mechanism' they postulate will interact with *the correct rule* and guarantee that the student will now exhibit the correct behaviour whenever faced with this rule. The psychic mechanism essentially divines the rule, which is there anyway, independent of our practices, and places its dictates clearly at the forefront of one's mind whenever one is asked to follow it. But just what form this dictate takes becomes the crux of the issue, and in showing that no such form seems sufficient or plausible so as to explain the capture of the rule, the idea that the postulation of the psychic mechanism is 'an idle intervening step' (McDowell, 1981: 206), as McDowell argues it is, becomes all the more appealing.

The idea, following Kripke (1982: 8-27) is that what 'appears before one's mind' has to take a form that thus explains the kind of predictable behaviour that understanding the rule affords us if we are to see the rule as existing independently of our practices. But the form itself seems debatable in any format we wish to give it. The idea that every infinite possible application of the rule appears there, all at once, seems simply implausible; after all, we have finite mental powers and cognitive abilities. A second option is that just some applications appear before our mind, with a prompt ('and so on') to carry on the algorithm. But the idea that the prompts can guarantee future behaviour is just to assume that one will carry on the rule correctly, which is exactly the idea we are questioning. All we have done is move the doubt from the practical sphere to some mental state. What was under question was 'how is it, that in practice, when given a rule to follow and after only seeing examples of it in practice, a student comes to apply the rule correctly, rather than follow one of the other infinite range of behaviours the rule could dictate?'. This second suggestion just moves the same problem into the mind of the student, away from the practice, but the same scepticism still recurs here as to how a few examples and a prompt can lead to correct behaviour. Without a way of being able to describe precisely what it is that the psychological mechanism does in order to delimit certain options such that the student can grasp one, it seems to be an extraneous asset in the explanation of our coming to know and follow rules. If we instead take the idea that we are just disposed to follow the rules thus, then we lose our claim to normativity, which is what this sophisticated non-cognitivist picture wishes to keep. Our being communally disposed to act like this and not that does not seem to amount to objectivity. Other dispositions might have prevailed and they can indeed change. We will have a fuller discussion of this point in **4.3**. It is this set of arguments that Kripke poses which McDowell turns against the non-cognitivist.

This is precisely why an idle intervening step seems to be a suitable way of summarising the 'psychological mechanism' that Gibbard (Gibbard 1992: 278) and other theorists of this ilk postulate. The importance of the Wittgensteinian account is that it shows that the fear - that this means concepts cannot be applied correctly or incorrectly - is misplaced. Leaving the learning of rules as being embedded within the practices that those rules come to govern does not, as the non-cognitivist seems to think, determine that the authority and objectivity of such rules are a matter of us

merely assigning them authority and objectivity in some arbitrary sense – as though one might simply stop seeing them as valid and use some alternate set at will, and claim that these are her rules on the matter.

...obeying a rule' is a practice. And to think one is obeying a rule is not to obey a rule. Hence it is not possible to obey a rule 'privately': otherwise thinking one was obeying a rule would be the same thing as obeying it (Wittgenstein, 1953: §202)

Normativity, then, is not something we can obtain from the anthropological position, as it leaves open the ability for vast interpretations of the application of rules. Understanding just why a concept or a rule applies here requires that one is immersed within the practice that the rule is said to govern. Just like with Rule 2 from the soccer example, one is required to understand the game of soccer *before* one can come to apply the rule. This 'coming to understand the game' is a sort of education by enculturation; by means of which the rules in question come to have the meaning and force that they do to participants who understand the practice or 'form of life' as active participants within it. From a position that is not within the practice, certain relevant features may be spotted as endowing an act as falling under a particular evaluation to those people and thus warranting the use of a concept that such an evaluation determines (think of our tourist attempting to determine through descriptive features alone what acts count as Taboo in Hawaii). But there are no guarantees here; and to make the mistake that a concept can always be categorized into requisite descriptive features, which make up the rule upon which the non-cognitivist supposes the correct application of the concept is to supervene, simply misses the point. As McDowell puts it:

Supervenience requires only that one be able to find differences expressible in terms of the level supervened upon whenever one wants to make different judgements in terms of the supervening level. It does not follow...that the set of items to which a supervening term is...applied need constitute a kind recognizable as such at the level supervened upon (McDowell, 1981: 202)

McDowell, following Wittgenstein, thus shows that not only does it not follow, but that such a position will leave one unable to comprehend vast swathes of rules, meaning and concept applications that one would understand if they were to go through the

process of inculcation and learn the ‘techniques’ that this form of education bestows upon one when one is asked to ‘follow a rule’ within a particular practice common to that ‘form of life’. Whether the rules under question are moral, or arithmetical, or otherwise in nature, they still have the same dependence upon those who understand just why they apply here being immersed in the practices which make up our distinctive ‘form of life’. And it is *this* which holds together and gives our practices the meaning and sense that they do.

That the rules are not grounded from a position external to our practices does not render them merely an ‘agreement of opinion’ (Wittgenstein, 1953: §241). This dependence is not in some way lesser than that proffered by grounding them externally. The lack of grounding them externally seems suspect in itself, but this is more so due to the fact that we cannot start to explain just how any psychological mechanism we possess could come to really interact with what we know to be the right application of the rule, when, from this anthropological position, any interpretation of it seems as good as the next. Without the ‘apparatus’ for discerning which interpretation is the right one, the idea that a rule can be assessed as correctly applied by someone standing outside our actual practices, and the agreement in judgements such an inside position works from, falters fatally. The Wittgensteinian reflections upon rules, thus drives a wedge between intelligibility and generality that is made out at the purely descriptive level. And this pushes up against the image of our collecting in purely sensory data which we split up into portions for the application of meaningful concepts and can be taught from the former level. I will again follow McDowell (1981: 214) and refer to these combined arguments against the non-cognitivist idea that we can disentangle thick concepts and find patterns for correct application at purely the descriptive level ‘the uncodifiability thesis’

The psychic mechanism, I suppose, is meant to be something like that of pattern recognition; but that we have the mental faculty of pattern recognition in no way means that the patterns we recognise could be meaningful or even comprehensible without inculcation into the practices in which those patterns have meaning.

Extending a series, or applying a concept, requires that we first grasp the practice in which such patterns are meaningful - and we cannot appeal back to those patterns to explain how we come to understand the practice we are taking part in without circularity. As Wittgenstein says,

If language is to be a means of communication there must be agreement not only in definitions but also (queer as this may sound) in judgments (Wittgenstein, 1953: §242)

While all these ideas give us reason to doubt the idea that we can cash out rules and extensions of concepts in purely descriptive terms, it has problems of its own. The next section will consider some of these in details. In short, the worry is simply that even though we know by way of communal inculcation into practices that certain things are right and wrong, there seems to be no way that we can say that this is definitively right and wrong. Were the practices different, we would have learned differently, and there seems to be no further recourse for appealing that the practice we do have is better than another we run into, or some hypothetical one. As such, while we can chide the non-cognitivist for reducing normativity to the purely general, we now have to show that we can hold onto it all, or that it is not merely a culturally relevant matter. We surely want more than this for the account to qualify as realist. These worries will be overcome or quelled by the analysis of the next section.

4.3: Getting Things Right without Foundational Rules:

At this point there seems to arise some serious issues regarding what it is to make claims rightly or wrongly. The terminology of the anthropological position does not help matters. It seems, using such language, that we push up against an intuitive idea that what is happening when we get things right is more than merely coming into step with a community, or a culture who are separated from other cultures and communities by way of the sorts of judgements that they make. This leaves open the distinct possibility, it would seem, for more than one conception of truth. In fact, if the anthropological position is unavailable, one from which we can know what is *really* happening by taking a God's Eye view, the community claim seems to be the best we have. And yet there is an intuitive idea that it is not the judgements of a particular community or a particular culture as to what truth consists in.

Any sort of realism, then, seems to require both, that, as a matter of necessity, there is no way to grasp a rule that can be made out as consistent from a God's Eye View, and yet that we have to do so, because truth cannot be fixed by way of our

judgements. We come to learn language by way our being integrated into a particular society, even at a particular time. The uncodifiability theory seems to pay homage to this. The differences that pertain between the correct uses of particular words at a particular time cannot be fixed from a position that is somehow outside our own special perspective, in a way that could be mastered by merely tracking the presence of certain descriptive features. As we saw, language users seem to require the evaluative aspect in order to be able to recognise and utilise the thick concept; yet we know that different societies in different times and places have different evaluative components to match the descriptive components we have. Without a position to tell between these – which appears to be something equivalent to the God's-Eye View – we seem committed to a community based truth.

But then from where do we obtain our sense of correctness that seems to be implicit within our saying we are carrying on a concept the *right way*? The whole business, under this Wittgensteinian analysis, seems to break down into something like a sort of anti-realism, whereby what is seen as correct application is done so by way of ratification from the community. The issue being that this opens up the question of whether the community can ratify their own applications in a way that seems independent enough of their own judgements to claim that they might be correct in a more fixed way. There seems to be two positions about the status of statements being true, both of which give up on a specific intuition while holding another firm:

- 1) The Independent Reality Position: This holds onto the intuition that those things that we say that are true are so independently of what we say and do. The intuition this runs up against is ever having a vantage point to reality itself. We seem to need a form of life to start explanations, but all explanations can only be couched in that communal language, and the community is always once removed from the independent reality by way of this language. Objectivity and normativity are maintained conceptually but seemingly inaccessible.
- 2) The Communal Truth Position: This keeps the intuitions that we learn our way around the concepts in such a way as to explain how we come to understand normativity and have varying views, as well as giving us something we can have access to in easy to understand epistemological terms. The intuition it rubs against is the idea that things can ever be 'thus and so anyway,

independently of our ratifying the judgement that that is how they are' (McDowell, 1984: 222). Objectivity and normativity are kept but diminished to a communal artefact.

This is the precise problem that McDowell tries to grapple with in his paper, 'Wittgenstein on Following a Rule' (McDowell, 1984: 221-262) the article from which I took the quote at the end of 2). There McDowell wrestles with how one might reconcile these two ideas of the normativity of concept application – the same getting applications right and wrong, the absence of which I argued was a flaw in the emotivist account – and the intuition that when we say that X is indeed an apt extension of a concept we are doing more than just appealing to the learned behaviours of a particular community's consensus upon the issue.

It was indeed with the community that this thesis left the issue. The sophisticated non-cognitivist turned our objection against the emotivist against us. They argued that if we could not show how the same concept applied to different objects from a position external to the practice in which concepts were applied (i.e. the anthropological position) then there seemed to be no room for right and wrong. Every application would be instantly deemed to be correct because we deemed it so. As such, there seemed to be nothing to say that we are going on in the same way, as independently-ratified. We settled with the claim that as there was no way for us to understand just how it might be that we could ever grasp rules from outside our own practices - using this to cast doubt upon the claim that independent-ratification was available – and, as such, took a position that could be read as making the claim that internal ratification or inter-community-ratification would have to do.

But the inter-community-ratification of 1) seems hard to stomach as well, precisely because it seems to make the claim that all there really is to our agreement in judgement is that we do agree. As McDowell was the theorist I followed in disagreeing with the non-cognitivist - and his use of Wittgenstein in emphasising the falsity of a psychic mechanism, thus, seemingly placing the normative constraints for going on in the same way as being inculcated into a community - I will again follow his lead and see if he can ward off this assumption that 1) or 2) are the only options. I will map some arguments from 'Wittgenstein on Following a Rule', here, with some elaborations. The hope of this section is to free us from worries, further down the

line, that giving up on 'explanatory theories' in ethics will mean following the fashions of our communities. The aim of this section is to see if there is a way to find a position between 1) and 2) above which keeps hold of both of the intuitions in a satisfactory way. McDowell does have an answer that I will try and bring out and expand upon in small part. But this answer has been seen as issuing in a sort of Quietism, which is the highly disputable position that philosophy ought only be descriptive of our practices and should never put forward explanatory theories. The problem here is that the answers available to a challenge such as Wright's seem light on the ground. If, as we absolutely must, we are to keep the possibility of certainty within our practices, then it may well be that an explanatory theory of such certainty will always be challengeable - which once again puts that certainty into doubt. I will look into this issue in Section 5, along with a few other issues that the McDowellian account brings up. For now I will track and elucidate the arguments between McDowell and Wright, which the former puts forward in 'Wittgenstein on Following a Rule'.

The main competing view in Wittgenstein on Following a Rule is indeed that of Crispin Wright, and his position as explained there is essentially 2). As such he prescribes to the view that the best position we can aim for is 'anti-realist'. Both he and McDowell express normativity of meaning in the form of a commitment to using concepts in a certain way – the right way. How they envisage this commitment being fulfilled is quite different, however, as we will see. It is a commitment that McDowell describes as normatively 'contractual', which can be put like this: If we use a concept once, then, if normativity is to be possible, there must be some pattern of ways of usage of that concept that one is committed to. Or in Wright's (1980) explanation of the pattern idea (one he eventually disavows) "committed to certain patterns of linguistic usage by the meanings we attach to expressions".

Given this sort of analysis, we might add that applying a concept in the right way is simply always using it in accord with that pattern, while making a mistake is coming out of step with the pattern (normativity of human judgement seems to loosely presume fallibility), and getting things wrong is not having grasped the right pattern (despite, one would assume, at least some overlap with the correct pattern's extensions).

Wright's position, following 2) above, is that "there is...no rigid, advance determination of what is to count as [a concept's] correct application" (Wright, 1980: 21). From this we get the idea that objectivity of the intuitively pleasing sort which McDowell seems concerned to try and salvage, is discarded here by Wright. It is done so largely due to his reading of what Wittgenstein says regarding explanations. When we try and learn a concept and are trying to find our way with the supposed pattern, any explanation that we receive as to how to go on is couched likewise in this same language that we are trying to grasp. Thus a gap opens up, between what is meant to be gotten across to us, and what we interpret that to be. The same applies for samples of examples, as well as definitions.

Wright thus takes this cue, regarding the ever present interpretability of language that Wittgenstein mentions, as pre-determining the idea that our understanding requires a leap, "...we move towards the idea that understanding an expression is a kind of 'cottoning on'" (Wright, 1980: 216). As such, the best understanding we can ever have is 'idiolectic' (McDowell, 1984: 224); an understanding we believe we have privileged access to, given the representational quality it has, as my best guess of the understanding the pattern that you were trying to get across. The issue with this is that it seems impossible to find a vantage point from which to distinguish between cases in which I am getting things right, as such, and cases in which I merely believe I am. This closes off the pattern idea as existing at an objective level, given the lack of vantage point here for knowing I am following it or not. The representational quality of the pattern, as I happen to see it, is the best I can ever have.

This relates, quite obviously, back to Kripke's idea of a 'quus' rule from our attack on the non-cognitivist. Unless my instructor were to push me for every conceivable application of the plus sign, it would be quite possible that my interpretation of the rule were not correct in as of yet unuttered applications, even though I assumed it to be so. Eventually, however, I might go out of step with my community – and this would show that indeed I did not have the pattern right.

Wright points out, however, that while appealing to the community's application is a way of getting back in step with them, it is not a way of trying to show that there is some access to what the pattern in and of itself ought to be. For the problem of every interpretation standing behind the next that motivates the quus rule just reappears at

the community level. Each member was taught at some point in the same way we are being taught and there is simply no way of bridging the gap between their joint application (assuming it is joint and at some point each won't come out of step) and a pattern that exists anyway as a right extension of the concept – the same pattern we wish to say we follow contractually when we start to use the concept. The idiolectic commitments, it seems, become a sort of colloquialist commitment to certain uses. But from here there is no court of appeal left. Person A can come out of step with the community and be corrected back into that pattern as best they can, but there is nothing the community seems to be able to appeal to in order to find their feet if, say, they are met with new situations where a concept needs applying. As such, there is nothing for them to come out of step with. And as such what the pattern ought to be rests *with the community*. A driving Wittgensteinian theme behind this notion might well be that abbreviated slogan, meaning is use. The 'pattern' is only meaningful because it is the one we use. Our changes to it therefore change the meaning of the concept, as opposed to the driving intuition behind the objectivity McDowell seems to want to salvage, where it seems we would change our use to accord with the pattern that best captures the actual meaning of the term – at least the one we are normatively committed to. In order to distinguish his position and make it credible then, McDowell will be required to show that meaning and cultural use can conceptually come apart in such a way for the community to use to be deemed in certain instances as incorrect by an individual who holds their ground on the meaning.

So from Wright's picture, the idea of objectivity in applications which stands above the community's attitudes to what is right disappears completely, and he shares this part of the anti-realist picture with Kripke. The key premise seems to be the idea that the pattern always requires our reaching out to grasp it, and yet this seems to always require an interpretation on the part of the person learning. In terms of objectivity, both thus condemn to failure the idea of a pattern that determines correct applications, because both claim the best we have is a representation of what is being put to us. The best we can get is a community-ratified-understanding. The community has nowhere to turn because ratification-independent patterns cannot be plausible. The looming intuition is that McDowell's own position is sending him back towards the Platonism we have just tried to leave behind, where rules track features

of a fully independent reality that exist beyond any perspective. As this is one part of the non-cognitivist's ploy of disentangling thick concepts, and also the illusory measure of certainty they claim we get from empirical claims, against which, post-Wittgensteinian analysis, we can now safely say they unfairly compare evaluative claims, McDowell has every reason to refute the independent reality position. However, without thinking of a way that makes it the case that we can understand concepts in way that defies the interpretation picture, then, we seem stuck with this conception: Interpretation stands in the way of our understanding anything more than what the community decrees, and even this normativity arrives *after* the individual's judgement – as their own understanding of the community's pattern is itself an interpretation. If the community changes its collective mind there remains no position for the individual to claim they are wrong. The position is a strong one and it is crucial that we overcome it if we want to hold onto realism at any level.

4.4: Rejecting the Interpretative Regress Picture:

McDowell courts two pictures of breaking through the interpretative picture, both of which Wittgenstein puts forward as implausible. The first is the Platonic idea of a 'super rigid machine' (McDowell, 1984: 231). On this picture we interpret what the person puts across to us, trying to find an interpretation which holds the application in place as our instructor seems to require of us. In each case, however, the Kripke-steinian sceptic is readily poised to claim that the mental rule we put in place as the real interpretation of what our instructor means is, likewise, open to interpretation. But if an interpretation could be found which could not be argued with, then from hence we could work back towards correct applications. We can perhaps think of this like an epistemological touchstone, mirroring the idea in the moral theories that we have canvassed, that there must be some sort of grounds from which we can run deductions that stands apart from, above or in special relation to any intuitive position or interpretation of application for the concept. From this view, however, such a stable rock seems worse than implausible to find and we are always cast back into the idiolectic conception of understanding. No such super-rigid machine can exist; the arguments against foundational rules preclude this.

The second idea – a ‘concomitant’ (ibid) of the first - is the idea of ‘rules as rails’, where the pattern exists outside all our practices and goes on to extend itself to all possible cases. This is another picture that McDowell, Wright and Kripke’s arguments are set to reject, with all three reading Wittgenstein as rejecting also. The picture after rejecting these two most obvious ways of finding an understanding that is not an interpretation therefore resembles a ‘dilemma’ (McDowell, 1984: 230) I will paraphrase: Horn one is the idea of infinite interpretability, while horn two is the idea of doing the impossible and finding an understanding or fact which is not an interpretation by way of one of the two ways above.

Given the talk of practices and communities that is ever present in Wittgenstein, Wright takes it that what Wittgenstein is claiming is that we cannot accept either horn of this dilemma, and so we ought to see truth conditions of a concept’s correct pattern of application as being those simply which the community ratifies. If the community changes its mind then so be it. The community cannot go right or wrong, rather ‘...the community just goes’ (Wright, 1980: 220). Interpretation is allowed because it is simply a fact that we cannot step outside our own position. This would be Wright’s version of bedrock: my spade is turned when the community decrees. One might want to say: This is just what we do. Wright says that community cannot go right or wrong. It sets the standard for individuals but it sets its own standards without a standard to be measured against.

The pattern cannot extend itself, nor can interpretations ever come to an end. As such, as McDowell makes clear (McDowell, 1984: 236-238), Wright’s explanation of Wittgenstein shares quite a lot with Kripke’s. The upshot being that both claim that the notion of truth conditions needs to be relinquished in favour of something more fitting of this analysis: for Wright, justification conditions, and for Kripke, assertability conditions. I will largely focus on Wright’s account, however, as I believe firstly that there is already good evidence for rejecting Kripke’s reading on the basis that I cannot buy into the sceptical solution idea, and secondly, because Wright’s account is less sceptical it moves us past the idiolectic conception and into community ratification as it does, and therefore more attractive because of it readily explaining the normativity we experience in concept applications, whereas Kripke almost loses all sense of meaning from the get go.

Thus, we have an issue as to whether we have, with the two positions - 1) and 2) - used up all the logical room in which normativity can still get a foothold. McDowell, in searching for objectivity which moves beyond judgements that have already been made and can be made out at the level of patterns, seems to require that the facts are just 'out there' (a phrase he courts but doesn't take up (1984:255)), whereas Wright wishes to claim that there can be nothing out there that wouldn't require interpretation, and thus all normativity belongs between individual applications and community ratification. There is then a challenge for both theorists. Wright will need to show his position can in fact keep normativity, and I will put the challenge that McDowell makes to this case forward next. But McDowell also needs to make sure that when one says the pattern is 'out there' he is not illicitly holding onto the second horn, the one Wright quite rightly rejects, and making the case for a pattern that extends by itself.

What McDowell questions is Wright's ability to be able to 'purge' the picture of the idea of ratification independent patterns, and yet keep the normativity of meaning – the idea that we are contractually obliged to extend the pattern of concept application in this way and not that, which allows us to be right or wrong in our applications. The issue for Wright is of course that all that this normativity consists in is determined by the community ratification. Our being right or wrong is decided post-judgement for Wright, as nothing can determine our being right given the possibility for attitudes to change and the impossibility of anything grounding those judgements that does not itself require interpretation.

McDowell makes the case, however, that the unsatisfactory nature of answering as basic an order as, 'bring me a yellow flower' (1984: 234), in the idiolectic case, like this: 'I brought you the one which gave me a feeling of satisfaction' – an explanation Wright would agree is unsatisfactory due to its inability to register normativity against the community standard - can be easily transposed to Wright's communal case with this answer: 'I brought the one that received approval from the bystanders'. That neither answer seems satisfying raises the question whether Wright's answer can actually be properly reduced to the idea of identifying a yellow flower, thus completing the order in a way that respects the normativity implicit in it, such that only certain actions will satisfy the request. And McDowell doubts that it can so satisfy it. There is no such thing left as just 'calling things yellow' (McDowell,

1984:235) on Wright's account, or indeed a thing's being yellow. The cognition of what is happening by Wright's picture is vastly different than merely applying a concept because some fact about the world (the flower is yellow) obtains.

Everything then hangs on this idea of how we might say a fact obtains. If it means that yellowness extends itself as a pattern logically independent of us which guides our actions on receiving the order then we fall into the trap of believing we have a way of recognising this pattern which would end all interpretations. As this judgement-transcendent pattern seems to be open to varying interpretations – on the basis that no super-rigid machine can exist which cannot also be pointed to as the right interpretation the real 'meaning' of yellow - we have no point of reference for our actions being labelled good or bad in accord to our actions. If my idiolectic understanding of yellow flower matches that of the person who issued the order then this is mere good fortune. Regardless of how sincere my personal belief that I am going on the same as before, there is no fact I can cite outside that sincerity which can confirm this.

The same lack of a fact outside our judgements is true of the community as well. They can keep me in line with their own applications of yellow by showing contentment or discontentment with my bringing this or that object, or extending a pattern or a series just so. But the same problem arises, of their having nothing determinate above these judgements to appeal to as holding them in place. So there is not independent fact of the matter as to why just this command sets just this pattern of future actions.

So there is no fact outside of our own judgements which allows for me to extend this pattern and not that one. Accordingly: objectivity, along with facts and truth, drop out of the picture entirely. There are no truth conditions involved in my making a judgement. My coming to feel normatively constrained in the application of my concepts thus cannot be made out at the normative level, the level of contractual obligations. The explanation of my actions thus requires description at what McDowell calls 'the basic level' (ibid). Something's being right is no more than something's attaining communal agreement of its being so. And its being so depends, it seems, purely on the judgements of the community which are colloquialistic. From this position, judgement precedes truth and as such *normativity*.

With no norms to appeal to, it seems that judgement has to be described at a non-normative level: and this is the basic level. My wanting to get something right or wrong is not explainable by my valuing right actions to be better. Rather, we are just trained, at the basic level, to commit certain actions when confronted with certain sorts of vocalisations and gestures and respond in this way and not that. If then the training changes, due to the fact that we stop receiving affirmative vocalisations when we bring a certain object, we are re-conditioned to start to do different actions. McDowell calls this a 'brute sounding off'.

What we have is a full ontological reduction of the normative and meaningful actions and thoughts such as following a rule or signpost or command, to the basic level and non-meaningful behaviours laid out in empirical terms. At least, this is how I read McDowell's analysis. On page 235 while explaining the basic level, he italicises certain phrases:

...does Wright's reading of Wittgenstein contain the means to make it intelligible that there should so much as *be* such an action as calling an object "yellow"?

It is problematic, however, whether the picture of the basic level, once entertained as such, can be prevented from purporting to contain the *real truth* about linguistic behaviour

The problem for Wright is to distinguish the position he attributes to Wittgenstein from one according to which the possibility of going out of step with our fellows gives us the *illusion* of being subject to norms

Due to this emphasis it seems that McDowell's reading suggests that what Wright is saying commits him to an ontological reduction of the normative to the purely descriptive. The normativity of our actions via our relation as sentient beings to the world would thus be an illusory emergent property which itself ought to be dispensed with. 'Real truth' exists outside normativity on this reading. Yet, clearly, Wright wishes to preserve a role for normativity. It was only the appeal to ratification independent patterns and determinate concept usage he wished to purge the picture of. He still believes that we have contractual obligations, but that those obligations cannot be made out by appeal to an antecedent truth. He might, then, think that the reducibility of the normative to the non-normative is simply explanatory. There are

simply two distinct levels of analysis: the level at which we consciously operate and interact, and the level at which those operations and interactions are best explained. So Wright might simply maintain that norms are not illusory mental phenomena. What is an illusory phenomenon is the idea of objectively set patterns of concept application that we somehow grasp without interpretation. All his arguments show is that 'Y is Yellow' is explanatorily reducible to 'Y is an item which receives community assent when I am asked to bring or point out something yellow'.

This latter position, however, runs into its own problems. If Wright insists on holding onto the idea that there are no truth conditions for statements and that normativity as experienced reduces to a normativity of community acceptance - community acceptance being that which allows for justificatory conditions to hold - then anti-realism, as a theoretical position postulated within a common language, has nothing to appeal to in order to prove its own truth. Its only appeal would be made within the parameters that it sets itself. This conclusion has the double defect of being both self-defeating and absurd. By denying truth conditions are in some way antecedent to certain judgements, anti-realism of this variety has, to borrow a phrase from Cora Diamond, sawn off the branch on which it is sat.

The absurdity comes in the idea that our theories, that which give us a sense of the relation in which we stand to everything else – be it other subjects, objects, reasons, actions and so on – can, on Wright's account, be made good from a position that itself only utilises the norm of 'assent to community'. This can be made out by trying to consider precisely how one might change the mind of another without appeal to any further shared commitments than this. There would simply be no explanation that was better than any other. All 'betterness' would consist in the communal assent. Disagreement itself thus becomes a major problem from here. Holding one theoretical position and not another needs justification, and yet all justification would need to be made in terms of community assent. That communities that share a common language and form of life are often vastly divided on opinions, and it is quite possible to hold the thought 'my explanation is better than their explanation' (with 'their' encompassing almost everyone else there is) it seems mysterious where the normativity for such a mental condition can come from. Assumedly, it is explained in terms of an idiolectic conception of what is meant from just these terms (the ones that pay homage to theory creation) that pays no heed to the community standards.

But as community standards are the start and the end of all normativity in conceptual pattern application, then stepping outside the confirmation to use these terms this way would make one wrong – or at least make one’s position lose all justifiability conditions. If one convinced their community to accept *this* and not *that* reading, it likewise seems mysterious as to how the reasons one employed to do so would manifest themselves in order to do so; they would only become pleasing reasons *once* the community had determined them as such as, “there is...no rigid, advance determination of what is to count as [a concept’s] correct application”. That is, there is no understanding of a rule which is *not* an interpretation according to this reading. There is no interpretation which can move beyond community assent for its ratification.

In diagnosing what has gone wrong here, it seems sensible to turn to Wright’s conception of what he is arguing against. And from here I believe we can see precisely why it is that he lays all the weight of normativity in the remit of the community. Wright’s argument is Anti-Realist, as he firmly believes there is no position from which we can say that X is just right or wrong independently of our judgements. But just like the obtaining of facts, everything hangs on from where we picture this independence as arising. Not the super-rigid machine or the self-extending pattern, certainly. Without these in place, Wright seems to think that we have lost all notion of the independence of judgements. He equates Realism with having to defend one of these two positions which Wittgenstein rules out. This is made clear by his insistence that “...the root idea of objectivity is that truth is not constituted by but is somehow independent of human judgment. Realism gives this independence the obvious interpretation: logical independence - the idea that for particular true statements it is either unnecessary or insufficient, or both, to meet our most refined criteria of acceptability in order to be true” (Wright, 1980: 199)

One might well take Wright to be correct here – and yet this seems a worrying position. Kripke’s position seems to usher in the ‘abyss’ where a sceptical solution is all we can have. Interpretation follows interpretation with nothing holding anything in place. As such, in our daily lives, we constantly oscillate between this idea - when we analyse the situation – and the comforting idea of the super-rigid machine. The other option, Wright’s, which avoids the oscillation, is to place all the weight onto the idea of conformity to community patterns – and yet this puts everything into either a

reduction that defies normativity or one which creates a self-defeating thesis of meaning. Kripke's and Wright's positions are seen by McDowell as equivalent to the twin monsters of Scylla and Charybdis (McDowell, 1984: 242), which we must steer a course between. If this course is an option we can show that all the available logical room was not used. Were it used, we would almost be forced to fall back into the idea that meaning is an illusion – which is such a counter-intuitive conclusion as to cause major issues for every theory there is. As McDowell says – there 'must be a middle position' (McDowell, 1984: 256)

4.5: The Middle Position:

It seems utterly imperative for own purposes, then, that we find this way of explaining just how we come to gain normativity in a way that respects objectivity – in the sense that we do not create right and wrong decisions by making those decisions. Other than showing these other two options as falling into chaos, we must find positive arguments that do not do the same.

From what we have said so far, however, it would be quite understandable to see the project as one that is doomed to failure. If we accept a picture of these two levels, one where normativity at least appears prevalent, and one where there is no more than a mere sounding off of brute behaviours, it would appear at first as though any analysis either starts with the former – which would lead us to an ungrounded and ever idiolectically interpretable account of meaning – or the latter – from which it would appear meaning simply cannot arise. Yet surely these are the only two positions available to start to ground any arguments towards a solution.

McDowell seems heavily to endorse the Wittgensteinian suggestion that there is a way of grasping a rule which is not an interpretation, but which is exhibited in "what we call "obeying the rule" and "going against it in actual cases". The endorsement gains some small amount of detail is this quote:

I think the thesis that **obeying a rule is a practice** is meant to constitute the answer to this question. (McDowell, 1984: 238) (My emphasis)

This quote is inspired by his reading of Wittgenstein (1953: §§201-202). It is here that the paradox of infinite interpretation is stated explicitly. Kripke takes it that Wittgenstein buys the paradox and offers a sceptical solution. Wright, like McDowell, does not accept this reading. They both agree that the second half of §201 and §202 both work against this reading ('there is a mistake here' (ibid)). Wright seemingly takes the suggestion of a practice to relate to the consensus. But McDowell is concerned, for the reasons already stated, to find more within this idea of a practice; more than what Wright makes of it, which might be thought of as such: one conceptual scheme amongst many possible ones, contingently ordering an otherwise unknowable reality that we are brought into line with conceptualising in this particular way by way of conditioning at the basic level.

I do believe that there is enough in McDowell's writing to make a decent case for the idea that we can do better than that which Wright is offering; holding onto an idea of both normativity in meaning and objectivity at the intuitive level, although he doesn't lay out a detailed or systematic case. This itself is barely of any surprise; Wittgenstein did not seem concerned to give a reading which could be laid out in a theory of meaning, given his reservations in his later work for such theorising, but was instead concerned with highlighting certain errors, assumptions and oversights in the way that we thought about these topics which could be employed against certain other theories. What I will try and do here is to give a decent account of the picture I think that McDowell is trying to get across from certain elements of the paper we have been discussing. I will add a few small remarks of Wittgenstein which I think will help his case. It will be brief, and programmatic. But as there needs to be a third way – as to not do so would be to endorse the nonsensical idea that anything is meaningful – then we need to at least have something.

The elements that I believe are the key to McDowell's reading take the form of key rejections of Wright and positive suggestions of his own. One of the key comments he makes about Wright's suggested picture is to be found in his rejection of the idea in Wright's writing of reality being completely detached from our actually judging things, and thus realism needing to find a position that was 'somehow independent of human judgements'. Of this McDowell says that:

A genuine fact must be a matter of the way things are in themselves, utterly independently of us. So a genuinely true judgement must be, at least potentially, an exercise of pure thought; if human nature is necessarily implicated in the very formation of the judgement, that precludes our thinking of the corresponding fact as properly independent of us, and hence as a proper fact at all. (McDowell, 1984: 254)

And:

We can find this picture of genuine truth compelling only if we either forget that truth-bearers are such only *because they are meaningful*, or suppose that meanings take care of themselves, needing, as it were, no help from us. This latter supposition is the one that is relevant to our concerns. If we make it, we can let the judging subject, in our picture of true judgement, shrink to a locus of pure thought, while the fact that judging is a human activity fades into insignificance (McDowell, 1984: 255) (my emphasis)

That this is a rejection of Wright's idea of realism and logical independence thus means that we need to start to think of realism as being formulated in a way in which our judgements do not immediately disqualify the judgement as being true – due to it being (at least) once removed from the reality that is 'out there' anyway. And the way in which we go about this ought to give us a stronger ground for thinking of judgements as true. If we can make this picture appealing and defend it well, then we can reinstate, with less worry, the idea of the truth conditions that Wright's picture rules out as a default. The only way of so doing requires itself that we find a way of understanding that is not an interpretation. But the very idea of interpretation itself is ushered in primarily by way of the fact that this is the picture we paint for ourselves. There stands outside us, and independent of us, all of reality, all of the truth there is, and then we are left to interpret it like so many shadows dancing on the walls, what casts them being forever out of view.

This platonic image is not an unwarranted one, I believe, as an analogy for the picture that Wright is painting, and yet it seems to beg the question about realism from the start. It is also problematic for Wright. The problem becomes clear when we think of it thus: If truth necessarily stands outside us, then whenever we have said before that X is true, then X has not been true. For Wright, this is made out as 'while

X has all the justifiability conditions to make talking of it true correct – it is *relatively* true, given the scheme we have taken’. But it is not true as such. Thus we have, implicit in Wright’s account, two standards of truth. The first is the way the world really is. The latter is how we best characterise it given the tools we have. This is pre-determined by his conception of realism as demanding truth that is logically independent of our judgements: “...the idea that for particular true statements it is either unnecessary or insufficient, or both, to meet our most refined criteria of acceptability in order to be true”.

What McDowell seems to want to say, is that there is no flaw in thinking that our very own ‘refined criteria of acceptability’ are what determine truth conditions. Wright thinks that *all* truth really takes place from the position of the basic level. And yet as there is no normativity there, there is no truth either. To be true is to be meaningful, and to be meaningful is to be normatively related to other judgements. Truth as a concept does not exist beyond our forms of life. It is used within it to classify certain states of affairs in a manner that is meaningful and practical and is intrinsic to that form of life. Something is only meaningful if it can be right or wrong; if it is true it is right. But truth requires a perspective to be true from. This is why I started the section with claim that the picture of the anthropologist’s position is problematic. It gives the confused idea of a ‘perspective-free’ perspective – of some creature anthropologically studying us who judges us right or wrong against how things really are. But all the arguments of interpretability that have gone before, for us, would just reappear here for our supernatural anthropologist. Her truth could only be her own ‘most refined criteria of acceptability’ established in the practices that she has been trained into. Either truth is equivalent to ‘our most refined criteria of acceptability’ which allows us to keep the concept, or those criteria do not seem to demand that what we say needs to be true. This whole thing then opens up the question, which of these two ideas of truth is the true one? Is it true that truth is a concept we cannot properly use? The confused nature of such a question seems to show there has been a mistake here. For anti-realism to be true, it requires that we consider it from the position that it states, as part of its formation, that we cannot have access to, which means we cannot know that it is true. If it instead claims that it is true according to our most refined criteria then one of the criteria it has to be using is that it is justified. So take any true claim you want – the earth is not flat – and using all

our criteria we can ask whether or not we are justified in believing it true. If we come out with the answer 'yes' – then anti-realism is false; we have found a statement we are justified in thinking to be true. To then keep the thought 'but anti-realism is also true' would contradict your previous claim that 'the Earth is not flat' is a true statement. The only defence against such a statement as 'the Earth is not flat' being true is to once again appeal to the 'perspective-free perspective' and then we are again left with the idea of not being able to say that anti-realism itself is true according to its own criteria.

So on neither the perspective free perspective of truth, nor the most refined criteria for truth can anti-realism be true. Realism, however, can be justified and *is* justified on the most refined criteria picture; the picture at which norms and justifications are readily available and admits of the necessity of some perspective for truth. The picture which purports to reject it is one that is naturally self-defeating.

Wright's entire picture seems to be infected with this image of a perspective-free perspective. It is also the inspiration for the idea that our world of meaning is reducible to a picture of the basic level, where all we have are brute movements and sounding offs. It is hard not to take this picture seriously, whether or not we are thinking about where or not truth applies. It cannot help but seem that if we could step outside ourselves we could easily be analysed in this way; from some position, to some other intelligent species, who were yet unaware of our customs, manners and language, this would be how we would appear. And this, I believe, is the biggest challenge for McDowell. Given that such an analysis seems possible – given, that is, that it seems necessarily true that we are organic beings who train each other in ways that could be explained as brute movements, inner states and brain chemistry – how do you imagine meaning ever getting a significant foothold? Given what we have said about the perspective of truth there seems there must be something to say here to quell the worry that meaning is an illusion. Such a position would require the idea again that this form of analysis is the 'real truth', 'out there' in the 'world as it is anyway' independent of all human interaction. From this I believe we can rule out the ontological reduction – something we have already cast doubt upon. What McDowell would seem to want to say is that this picture of the basic level is best conceived as an explanatory reduction.

Given that explanatory reductions didn't sit kindly for Anti-Realists, we may at first think we have worries here. But the reason it didn't sit well was Wright's insistence that there are no such things as truth conditions. For McDowell's form of realism this is not such a worrying problem. What opens up as a possibility for McDowell is that two statements regarding the exact same phenomena can give true accounts of what is happening, regarding that phenomena, both of which are true, but where each of which take the form of a distinct type of explanation to the other.

...the denial of ratification-independence...yields a picture of the relation between the communal language and the world in which norms are obliterated. And once we have this picture, it seems impossible simply to retain alongside it a different picture, in which the openness of an individual to correction by his fellows means that he is subject to norms. The first picture irresistibly claims **primacy**... (McDowell, 1984: 248) (my emphasis)

This idea of primacy, I believe, will go a long way to explaining just how McDowell could hold both pictures in place as reasonably co-existing on separate planes of analysis, both of which have a claim to being true. In terms of Wright's account, where ratification independence is denied, norms are obliterated because they are secondary; they are an illusion we have of our understanding that is properly conceived of as the interaction of a certain species of organic being, with just this sort of brain chemistry, situated in just this sort of environment, interacting with the primary objects around him within the confines of certain natural laws of motion.

But given that from this position there seems to be no idea of either anti-realism or realism being a justified position to take we have reason to doubt this idea of primacy of one picture creating the other and thus constituting the *real truth*. The intuition that is driving this whole idea may be one of temporality. The universe, we rightly think, existed anyway as a set of natural laws into which we were thrown. Then via our interactions with it, we began to analyse it and thus we came to place upon it a perspective. Thus everything we know should be reducible to that which was already here prior to our perspective. Thus that picture of brute forces at the basic level demands primacy.

What I think McDowell wants to suggest is the opposite. And it is indeed more fitting to what we are properly discussing to do so. The argument we are making is about

meaning and understanding and the ideas - such as fact, illusion, objectivity, relativity, truth - which go with that. In terms of understanding things, what came first is, in fact, this understanding at the most normative level: a full blown 'naïve realism', if you will. It was then from the concepts that emerged from within this language of normative interactions that the idea of other layers of analysis became possible. And once we see this there is this possibility: that this other level of analysis we imagine – the supposed perspective-free perspective – *is still our perspective*, just another one of them, for a different type of analysis that is as much a part of our practices as the more normative ones are. We study other phenomena in this way, claiming it is a disinterested position. But it is not and cannot be fully removed from the language in which it arose or the interests or practices of the members of a community or form of life which utilise it. Our actions can be explained in a way that does its best to strip them of our everyday concerns and interests – but it is still of interest to us. It is still of interest because it makes up a huge part of *our* practices and explanations.

Section 5: Painting a Non-Reductive Moral Picture

5.1 From Brute Movement to Norms:

What we have said so far can now lead us, at least to some degree, to answer the question of how norms and meaning can exist in an external world that can rightly be thought of as norm free, without contradiction: the two are both perfectly acceptable explanations of the same actions, objects, relations or events. As long as both qualify according to their respective 'most refined criteria for acceptability' that we have, then two explanations can be true at the same time, as long as they exist on these distinct levels of analysis.

And this two level conception thus not only explains how we can have norms while also appealing to brute movements and the like, but also a way in which understanding need not always be seen as an interpretation. I will aim to bring this out in due course.

McDowell hints at the two distinct levels on a few occasions. The most obvious being the way he appeals to Wittgenstein's concept of bedrock as functioning. He quotes Wittgenstein to help his analysis.

what has the expression of a rule—say a sign-post—got to do with my actions? What sort of connexion is there here?—Well, perhaps this one: I have been trained to react to this sign in a particular way, and now I do so react to it.
(Wittgenstein, 1953:§198 in McDowell, 1984: 238-239)

When Wittgenstein's interlocutor reacts in the way that we would expect of someone who was gripped by the anti-Realist picture,

“But that is *only to give a causal connexion*: to tell how it has come about that we go by the sign-post; not what this going-by-the-sign really consists in” (ibid)

he responds thus:

—On the contrary; I have *further* indicated that a person goes by a sign-post only in so far as there exists a regular use of sign-posts, a custom (ibid) (my emphasis)

Thus, the training at the basic level is what explains the lack of interpretation. We are trained in a way which means that *sometimes at least*, there is a way of understanding a rule which is not a mere interpretation. The explanation for the certainty and the lack of interpretability is thus accounted for on these certain occasions precisely because the training in question is what we start from and as such defies interpretation. We are conditioned to accept that certain things just are what the people who train us say they are. We are, that is, trained to accept a form of life, of finding certain similarities and differences salient, and as holding firm certain judgements at certain times. If we are to question this we will become unstuck from the glue which holds us together. We will have misconstrued a foundational pillar of our communal understanding - from the normative level - as something built normatively on top of it.

From our actual perspective, a norm-laden perspective, from which blossomed this idea of an alternative conception – which while detached from our cultural norms is still meaningful enough to strike us as meaningful (i.e. for it to be a right or wrong analysis here) – what is being taught is just another normative statement. At the basic level what is happening is a sort of conditioning which explains away the lack of interpretation available and thus grants us certainty at the very beginning of the language game. As some evidence that Wittgenstein wouldn't withdraw from this view we can quote these sorts of passages with McDowell:

Someone asks me: What is the colour of this flower? I answer: "red".— Are you absolutely sure? Yes, absolutely sure! But may I not have been deceived and called the wrong colour "red"? No. The certainty with which I call the colour "red" is the rigidity of my measuring-rod, it is the rigidity from which I start. When I give descriptions, that is not to be brought into doubt. (Wittgenstein, 1978: VI - 28 in McDowell, 1984: 240)

And:

How do I know that the colour that I am now seeing is called "green"? Well, to confirm it I might ask other people, but if they did not agree with me, I should become totally confused and should perhaps take them or myself for crazy. That is to say: I should either no longer trust myself to judge, or no longer react to what they say as to a judgement. If I am drowning and I shout "Help!", how do I

know what the word Help means? Well, that's how I react in this situation.—Now that is how I know what "green" means as well and also know how I have to follow the rule in the particular case (Wittgenstein, 1978: VI-35 McDowell, 1984: 240)

So we are *trained* at the basic level and *taught* at the normative one. These two processes are one in the same, but seen from two distinct levels of analysis. In our day to day lives we inhabit the normative level. As such, we think of things which in themselves are normative as explanations of what we do. We follow the sign post or the rule, we obey orders, and we justify things according to these ideas: we play all the games that we are taught to play. On the other level of analysis this is not what we are doing. From this perspective we *act* as we are trained to act when someone makes this sound or these gestures. Between these two sorts of analysis, says Wittgenstein, lies bedrock. Bedrock thus plays the role of demarcating two separate and distinct modes of analysis and explanation. Above it, at the normative level, lays justification; below it lays causal explanation. At PI §217 Wittgenstein says this:

"How am I able to obey a rule?"—*if* this is not a question about causes, *then* it is about the justification for my following the rule in the way I do. If I have exhausted the justifications I have reached bedrock, and my spade is turned. Then I am inclined to say: "This is simply what I do." (Wittgenstein, 1953: §217) (my emphasis)

Here Wittgenstein clearly demarcates justification as presiding above bedrock, while causal explanations lay below it, thus the 'if – then'. Clearly the justifications we give to each other are effected by the sorts of training that we have received. Gaps in someone's training would mean that certain justifications fail to come across as seeming salient to them. Thus they might become confused as to what seems to the rest of us as the obvious justification. If we reach bedrock and no justification we have given convinces our interlocutor we should be happy resting on the idea that he has not been taught properly and put the mistake down to being a causal one. Such a diagnosis can be useful and we may have him undertake the training again. Wittgenstein would be quick to warn us not to conflate these two sorts of explanation.

That was perhaps the point of the quote, above, regarding the colour red as the 'rigidity of our measuring rod'. As a normative comment we think it can be justified. But at bedrock the justifications run out and any questions need to be made in causal language. Wittgenstein seems to think our conflating the two sorts of explanation (which is hardly surprising when the language of the basic level arose from the language of the normative level) is the root cause of major issues in philosophy and thought in general.

I'll leave this section with the same quote with which I started it. It seems to me to be the quote which best sums up the reading of Wittgenstein that I have made, while following McDowell's line of thought. McDowell is less explicit about the two levels of analysis. Although the idea of there being a 'basic level' indicates there must be at least one more level. More likely is the idea that McDowell was letting Wittgenstein speak for himself through the aphorisms he quoted. Likewise, McDowell was more concerned with pointing out the flaws of accommodating normativity in Wright's anti-realist picture than in promoting another. All he needed to do was show a route from the final dilemma that did not force us to live with Scylla or Charybdis and which accommodates normativity whilst staying true to idea that certainty is not underscored by our minds interacting with a logically independent pattern that extends by itself. That picture does not and cannot accommodate normativity, leaving us without truth conditions and thus without certainty. He certainly does enough to intimate this route.

Before the quote though, it is worth saying one more thing. Our answer to the problem of understanding always being an interpretation (or being best described as involving interpretation) seems to occupy a similar area as that of the super-rigid machine. But it was always going to have to if it was to avoid the idea of interpretation all the way down, as it were. If we are going to incorporate this idea of bedrock, which explains where we get these intuitions of justifying normatively, while holding onto the intuitive notion of objectivity and diagnosing the cause of so much conceptual confusion in philosophy, then what we have to say, with Wittgenstein is this:

If language is to be a means of communication there must be agreement not only in definitions but also (queer as this may sound) in judgments (Wittgenstein, 1953: §242)

5.2: Objectivity, Analogy and a Space for Morality:

McDowell's argument against Wright, then, is one which wrestles with the underlying assumptions which underscore Wright's own anti-realism - which is, as previously noted, an attempt to show that objectivity must be reducible to how groups happen to think and thus is always one step removed from what is there anyway. The 'is there anyway' is itself given force from the prevailing notion of objectivity which is the mainstay of the empiricist: that which discerns the label true is that which could be thought as existing prior to and independent of any sort of perspective whatever. The falsity in such a picture is that it rids us of any hope of ever discovering the truth and thus makes truth (and its related concepts of real and objective) a redundant concept for us, when in fact these concepts do play major roles within our thinking. It is only if one buys into the idea of a perspective-free perspective as necessarily underscoring the correct use of these terms that one need to buy the sort of anti-realism that Wright is pedalling in the argument he has against McDowell. Wright's picture does change in his later writings (Wright: 1995) but it is key that we see this earlier attempt at least as flawed and to see why the idea of some perspective is necessary.

However, the argument between Wright and McDowell, and the fact that McDowell manages to show Wright's thinking to be flawed, is only a partial victory for the moral realist. For the argument is a generic one which covers the use of concepts in general and does not show that moral concepts are necessarily of the type that can be salvaged. That is, McDowell may show that in general one can say that X is objectively correct, if X stands for the bringing of a yellow flower, or if X is that diamonds are hard. Or perhaps more cautiously, he at least shows that the idea that it can never be correct - due to the fact we can never escape our perspective and

see whether what we think of as being yellow can be saliently grouped together regardless of any perspective - is a flawed worry. But the generic shared concerns that reinforce the idea of 'shared forms of life' could only be said to stretch out to encompass moral concerns if one were willing to completely overlook the sorts of differences there are between what one would usually mean by 'shared concern' when applied to such moral concepts and what one would mean by 'shared concern' when considering how we seem to know whether we are, for example, playing a game with strict rules, or merely throwing a ball around for fun.

In the latter option we see just how practical our concepts in general are and how necessary it can be to be integrated into a form of life to comprehend the ends and purpose of what we are doing, and why two distinct things like throwing a ball around aimlessly and baseball can both be categorised under the concept 'game'; both share so many similarities and yet both be seen to imbue similar actions with such different meanings when we are undertaking them. It is explaining this to an outsider that we believe that we will have very little chance of doing without her first having grasped at least something of what it is to *be* like us. In the moral sense so much is meant when we think of shared concerns, namely, do we believe this thing (whatever it is) to be bad, good, laudable, reprehensive or whatever. Concern here does not imply merely having the right grasp of the situation to know how one is to go on using a concept, but to take up a motivational stance which informs actions and the use of concepts; to feel how we feel towards certain things. This is what is meant by Mackie (1977) when he says that the same duality we find in thick concepts – concepts which both describe and give reasons for action – seems to imply a queer metaphysical entity.

So McDowell rescues the idea of truth and objectivity for our general concepts from Wright, allowing that the invocation of Wittgenstein need not throw us into a communal subjectivity which leaves truth forever outside our grasp. In so doing, however, he cannot be seen as yet rescuing morality. If one stops here one loses sight of just what makes the moral a distinctive category in the first place. Objectivity may require a level of perspectivalness, but whether that level stretches to the sorts of perspective which makes moral concepts appear to provide us with reasons for action (the very thing which makes them distinctively moral in the first place) remains to be seen. Clearly, sophisticated non-cognitivists could grant that description might

be such as to require the former type of perspective without having to relinquish that the extra level of perspective - the type which sees certain concepts (thick moral concepts) as providing as part of their very meaning reasons about how we ought to behave - is added on as a motivational extra by the individual agent. With this being the case, the challenge would revert back to the idea that if thick moral concepts are really (from some perspective) descriptive and evaluative then the two components could be disentangled such that the concept can be separated into its two components. The idea of an outsider being challenged to grasp moral concepts only based on pure descriptions no longer makes sense, as we have been shown by the McDowellian analysis of Wittgenstein that they would not just struggle with moral concepts, but concepts like game as well. Thus the challenge must be altered to fit someone who shares enough of our form of life concerns to grasp descriptions of complicated cluster concepts that are underpinned by shared concerns, without having to share our evaluative concerns.

Thus, the challenge might be put thus:

that an agent who doesn't share the usual reaction to such concepts - that is, someone who felt no compunction to be honest or courageous - would nonetheless be able point out honesty and courageousness wherever they were to be found by an agent who did feel such a compunction.

But straight away it should be clear that such a challenge admits of a flaw. It is the same flaw we found in the externalist idea from earlier when we were discussing Judith Jarvis Thomson. That any sort of guess by the supposed amoralist, whom the reformulated challenge assumes exists, would be necessarily parasitical upon the normal understanding of the concept which would carry with it the evaluation within the meaning. As Wiggins notes in his argument against Gilbert Harman (Wiggins, 2000: 150–3), what keeps such concepts grouped together is the shared perspectivalness of those concepts with their full meaning – that is as answering to our moral concerns.

It is important we make this point clearly as it is a crucial piece of the puzzle as a whole. The position we are trying to defend is one of a full blown non-reductive moral realism, where the meanings of moral concepts include the evaluative portion. This certainly makes the whole category quite odd, but the oddness is a double edged

sword. Here is what I mean by that. Mackie notes that the very essence of moral concepts is quite unlike anything else postulated because they claim to describe the world while also giving us reasons for action. Unlike the non-cognitivist, he takes morality at face value as a describing game. But, he claims, those concepts describe nothing at all. They are like concepts such as magic or witches, which may appear to have properties which fit their descriptions but which in fact refer to nothing at all. If they did refer, the things they referred to would be awfully queer. The classic method for deciding whether something qualifies as objective is to define what makes something apt to qualify and then to determine whether the candidate has the right materials. We have noted that such a tactic seems to beg the question, for where one draws the line is rather arbitrary. The empiricist assumptions that Mackie clearly brings to the table are the demarcating point for him, but for all that the oddness which we see in moral concepts does indeed give us pause for thought. These are no normal concepts.

However, as the sword swings the other way, it seems that this may well be a blessing for morality, for of course moral concepts are unlike other concepts which are not moral, this is why we have a distinctive category for them in the first place. The only defence we are likely to have for moral concepts will then be a *suis generis* defence which matches their *suis generis* character. We are not arguing that all moral claims can be seen as objective claims in the precise same way which perfectly descriptive claims can. For them to do so they would need to *be* perfectly descriptive claims; and clearly they are not so. Thus, while *prima facie* their distinctive character works against them, in as much as they do not simply fall into any preconceived notion of the empirically objective; the same distinctiveness is precisely what would be expected for any sort of concept or property which is not an empirically objective concept. Otherwise the distinctions between the categories of concept would be mysterious at best and both pointless and unfathomable at worst. It is only if one has pre-set assumptions to the effect that only purely descriptive properties can be objective that this would call into question the possibility that moral concepts can be objective. And in the current argument such assumptions would be question-begging in a most obvious manner. As such, an argument must be made that the category of the objective is not nicely demarcated along such lines.

Wiggins and McDowell have both claimed that the sort of empiricist assumptions that Mackie carries can only be seen as marking out the borders of what counts as objective at quite a cost. The cost at its most destructive is the idea that we cannot explain how we come to know what another is thinking, that we will be left interpreting the thoughts of others at all times. The Wittgensteinian argument regarding concepts in general which was detailed above is the focus of this. The pragmatic points there went to show that meaning cannot be simply thought of as arising from the mind's interaction with a fully mind-independent world. There needs to be some level of communal agreement at the start for communication between participants to occur in such a way as to provide a meeting of the minds. The implication here is that objectivity ought to be conceived as pertaining to the sorts of judgements which hold firm once the perspective which allows for the meeting of minds is admitted of. Once we have brought the idea of human perspective into the fold – something we are forced to do to allow for a coherent picture of certainty to be manifest - then appealing to that which we see as fully independent of us as also being that which defines what is true is problematic, as full independence is seen as a fallacy. The way we think of subjectivity is thus now understandable as not simply the input of some perspective, as even those domains of thinking we see as characteristically aimed towards objectivity (the natural sciences) admit of some perspective, it is just that it admits of no particular perspective. Thus, such an input cannot be akin to relativity.

What must be made out is just where the perspective enters into the equation and what influence it has. And Wiggins and McDowell both believe that there is truth to be found in certain domains which a strict empiricist would likely see as subjective in the pejorative sense of equating to relativity. By explaining how a shared perspective enters into these domains they lay out an alternative account of what amounts to objectivity which, they believe, can be used to bolster, by way of analogy, the hopes of the moral cognitivist. The aim here is simple: to show that ontological concerns (the idea of primary qualities) needs to be replaced as the demarcation point for objectivity by epistemological concerns - standards for what we can know about, with an emphasis on the 'we'. From hence we can start to understand which domains of thinking are subjective in the pejorative sense and which speak of a shared perspective but still allow that the content of the mental states we are in when we

have them are proper belief states. The two analogies (neither of which are seamless analogies for morality, but which open up the idea that 'things which can be said to exist outside our perspective' cannot be the right metaphysical rendering of all that can be objective) are those of colour categories and of mathematics.

Before we begin with the two analogies then it is crucial we consider just what is at stake here. What we are trying to uncover is what might demarcate the objective from the subjective, how these categories stand in relation to each other and whether or not, at least sometimes, moral statements can be counted as belonging to the former category. The reason I mention this is that McDowell has decried any such attempts for us to try to formulate an explanatory thesis for what it is for something to be objective or what objectivity consists of. He opts, on the basis of his reading of Wittgenstein that I posited above, that the best answer here is a sort of quietism to such affairs. That is, he believes that we cannot and should not be explaining the metaphysical status of variant domains of thinking. To do so would be to stop describing what is in front of us and to instead start explaining what, according to his Wittgensteinian groundings, we cannot explain – namely, the metaphysical status we ought to say those things have. I have sympathy for this position and am inclined to agree that there is an exegetical case to be made that Wittgenstein took it at the very least as serious contender for how to address certain questions that arise in philosophy. As I read Wittgenstein, he rallies against the position of the perspective-free perspective more than any other, that is, against the Platonic ideal of something like a transcendental realism. This initially can make him seem like someone who would be better labelled as a transcendental idealist - allowing for empirical realism only from this immanent position. However, the question of transcendental realism only arises once one has asked some very serious questions about what various items of knowledge consist in. And for me, Wittgenstein rejected those very questions as leading us into philosophical fictions. This means that ultimately the transcendental idealism position is itself rejected as the answer to a dubious question. The ideas which carry it – those of meaning being carried along only within linguistic communities with shared forms of life – are not themselves a positive thesis, but best thought of as an attack on the Platonic picture, a picture we needed to escape from. If, once the rejection of Platonism is complete, we then reject the sort of questions which place transcendental idealism in an

explanatory position, we are left with merely the empirical realism – that which is, from a human perspective (the only perspective we can actually have) truth. This, as far as I can tell, differs somewhat from McDowell's reading which, especially in later writings, keeps the idealist pretensions, but all the same ends up with the closing off of questioning about what various things, objectivity included, really consist in. For a quite recent defence of that quietist position see his 2009.

I will openly discuss possible conceptions of objectivity here, especially in relation to morality, and I believe that this can still be done and respect what is good about McDowell's work. Firstly, what it leaves open are arguments from descriptions of how we do act and how we do think which, as Wittgenstein maintained, were the hallmark of good philosophy. As we will see, those thinkers who do demarcate the objective and subjective along strict theoretical divides are best countered by way of descriptions of areas of thinking which, however intuitively, do not follow these divides (the previously mentioned analogies will play a major role in this). A further point is the doubt that explanations and descriptions come apart so easily as to themselves demarcate distinct and contradictory categories. Rather they seem to be interdependent categories held together by the very act of reasoning. As humans, we do things for reasons, we do things which are justifiable and this includes describing. Describing is more often than not an act of clarity or demarcation which allows us to reason more clearly. As such, one only describes towards an explanatory end, and yet one can only explain that which is describable. Thus they are both things that we 'just do'. What we mustn't do is add metaphysical fictions that take us further away from describing what is actually there, but conversely, we should surely not refuse to give an explanation of what is describable, so long as the explanation is not seen as diverting our attention from what is actually in front of us. The key point of quietism is to not try to go so far with philosophy as to try to describe the inexplicable or explain the indescribable. Here we will have run into fictions of the metaphysical sort. Surely it is not beyond us to say something sensible about objectivity. This may not be a global thesis of everything that falls under the term, but this does not stop the project of wishing to say whether certain domains can be ascertained as being more or less objective in their own terms, and what those terms might be, even if this is not from some transcendental standpoint. McDowell's own position indeed goes a long way to showing this for morality, by

way of descriptions of how we think of other domains of thought. From such descriptions certain things follow, if only the rejection of various metaphysically fantastical pictures. By thus narrowing down the options of understanding the domain we can get a better grasp on explaining what that domain is all about. Thus, while I sympathise with McDowell's position, I am led to believe that simply ignoring various challenges regarding the objectivity of the moral could be counterproductive to anyone who believes themselves a cognitivist and wishes to recruit others to that cause. If one prefers to overlook the points at which I may be seen as delving below bedrock as unsatisfactory, that's fine. But one might also rest easy with the idea that such comparisons of language games needn't be thought of as merely explaining one to be more or less objective in comparison to others, but as elucidating the descriptive differences between them, even if such differences are seen as not casting any connotations on their metaphysical status. In this way it can be seen as an exercise of a deeper description of each language game, of what gives each its distinctive character. In my own mind, given the vague and contentious nature of Wittgensteinian comments regarding this matter, I will envisage the idea that we should not offer explanatory theories as being a check against offering theses which go beyond our right to make them, rather than a total ban on offering any explanations whatsoever.

As noted, McDowell does make a certain case for moral objectivity based on an argument from analogy. Alongside David Wiggins he postulated that there are examples of statements which we see as being simply true which cannot be upheld by the empiricist theory of mind which underpins the accounts of detractors such as Mackie. Mackie, as previously noted, held certain staunch empiricist views which he barely hid at all in his arguments against morality, from queerness and from relativity. The arguments from analogy do not only work against those such as Mackie, but they are particularly effective against such an opponent, one who effectively seemed to believe in a Quinean naturalised epistemology. Mackie's argument was that morality is a describing game which falls into error due to the fact that the properties it aimed to describe simply did not exist. No property could be both descriptive in the full sense and also move one to action. Any property with such a dual characteristic would, according to Mackie, seem awfully queer. Furthermore, these properties did not seem to imbue a central level of communal agreement. Different cultures utilised

the same ethical concepts to describe different properties and even within cultures disagreement can be rife. This is not so with clearly descriptive claims, Mackie noted, and thus these two arguments went on the record for the cognitivist to debunk. As previously mentioned, the queerness need not be seen as an issue; one can simply assume that this is the laying out of a feature of a domain (the descriptive) in which we are happy to assign objectivity and then asking the cognitivist how it is this can be so and yet moral concepts and properties can be seen as objective. This is not an attack, more the laying out of a task any good cognitivist would believe themselves able to rise to. (Grice: 1991)

To rise to it will necessarily include a departure from the conception of objectivity which Mackie brings in tow. And it is precisely this that the cognitivist project that McDowell and Wiggins are suggesting uses a Wittgenstein analysis to achieve. In doing so they set up an alternative to the strong and unitary empiricist picture which gives the grounds for buying into error theory and also non-cognitivism of the sort one finds in Blackburn and Gibbard. Once this is achieved we will have good reason to think that the sorts of differences that these theorists love to posit between moral claims and descriptive claims set up a telling unfavourable comparison only if one does not relieve themselves fully of the overturned picture. The argument that we cannot understand any sort of knowledge as occurring without a perspective has been well rehearsed now, so these two analogies must offer something more than this, and they do. They offer a picture of how it is that a shared perspective can be a foundational pillar of a domain of thought without showing up in the specific judgements that are constitutive of that domain of thought so as to threaten the judgements at this level. The domain is beholden to our perspective, but once we adopt the perspective the judgments need not be thought of as such. The concepts and properties of such domains transcend the pragmatic explanation for their existence such that beliefs generated about them can be apt to be true or false in a proper sense. As we will see, such a picture is particularly complimentary for the moral domain.

Let's take colour first. Colour properties are properly anthropocentric. Take away the first order perspective colour is perceived from and colour stops functioning as a proper category. That is, there is a fundamental explanation about how we come to perceive colour which can be made out in purely physicalist terminology and

includes instead a description of light bouncing off variously textured surfaces to create variant wavelengths and enter our optic nerve in such a way that will produce in us a perception of something which we as a creature of a certain sort finds interesting. Without our physical make-up, such surfaces may produce no such differences, and without our interests we may not find such differences of interest enough to assign them different colour concepts. It is perfectly contingent on us being us which means that we have the category of colour. Colour does not 'exist anyway' in the empiricist sense. This is what makes it a *secondary* property in the empiricist literature. The texture, the light, the optic nerve and the brain exist anyway. But the colour property and the corresponding concept do not.

Pushing up against this point are the various Wittgensteinian considerations we have previously rehearsed about there being perspective-independent worlds, but also another point – that explanations about colour are clearly thought of as being true or false. It is simply true that blue and yellow will make green. And there is nothing one adds to this by adding 'from our perspective', other than to demarcate that which we categorise as fundamental building blocks of the universe and that which we do not. But to be reductivist here is to discount the point and purpose there is in colour explanations and the sort of consistency and convergence that is necessary for such a point and purpose to be manifest within that explanatory space. To discount the truth of the whole category is to discount also the truth of the explanations therein and yet, as McDowell succinctly puts it,

A [colour] is a property the ascription of which to an object is not adequately understood except as true...in virtue of the object's disposition to present a certain sort of perceptual appearance (McDowell, 1998: 133)

Thus, categories of thought needn't aspire to being objective in the empiricist sense to also be thought of as apt for cognition. As long as one shares the perspective then one can start to conceptualise the consistencies that there are within the category and make inferences from them. The contingent nature of our own subjective psychological make up is fully on display here. But it doesn't itself enter into the explanations of one who has access to the category. As Wittgenstein showed, there is no transcendently realist 'true anyway', to be distinguished from a perspective-free perspective, so true and true for us do not come apart meaningfully. Deep

enough down everything is 'true for us' in that one must admit there is no truth outside of a perspective. The argument thus pertains to the idea that perspective can be omitted from the equation when we are considering whether concepts can be seen as fit for cognition - assuming we want anything at all to be fit for that purpose. Thus, explanation takes over from ontological considerations of the fundamental sort in demarcating whether we can cognise a property. And we have an explanatory space for colour which has strong internal relations, allows for convergence between agents in judgements and has a point and purpose which is wide ranged and holds essential functions in our way of life. As such, the explanatory base transcends its physicalist base. The physical elements which allow for the explanations to be meaningful are a necessary element of colour explanations but they are not sufficient for a full reduction of the explanatory base to the physical components.

Key to this point is the idea that, without the manifest image that the empiricist wishes to explain away with reduction, there can be no explanation of colour as there would be no category recognisable from the basic level to explain; as Wiggins puts it, we could not 'saliently group together' the items constitutive of that category without having an understanding of the category from the position of those who have access to the explanatory space of colour. The properties which constitute the manifest image are conceptually and explanatorily prior to, and are a necessary pillar of, the empiricist *explanation* of them, whilst the primary properties the empiricist points to are a necessary element of accessing the manifest category. This is not enough for reduction. The physicalist rendering will not be able to be swapped out non-controversially for the manifest image rendering enough to show the former is an identical but 'more objective' equivalent of the latter. In the case of water reducing to H₂O, reduction is possible, precisely because this swapping out can explanatorily occur. H₂O is 'more objective' in that it allows us to differentiate between water and other liquids which might share all its supervening properties with water other than its molecular make up.

If we grant the above, which I firmly think we should, we also come to a further realisation. What counts as the standard for what constitutes good and bad colour categorisation and conceptualisation can only be ascertained from the point of view of the ideally placed observer within the first order perspective. We cannot even consider the idea that we 'all have colours wrong' without trying to understand colour

from outside of our perspective – but from there, as mentioned before, the colour concepts and their corresponding properties cannot even be sensibly grouped together. We thus judge others' ability with colour concepts against an agreed upon standard set by the practice of assigning colour properties to objects that cannot be understood from outside the perspective. Thus, our analogies are giving us a new conception for conceiving of domains of thought which admit of a perspective but where that perspective is considered as a key element in conceptualising an independent explanatory basis, made out as a space of reasons, which is not reducible to the physical components which allow for that perspective to take shape.

The key disanalogy between colour properties and moral properties is that moral properties merit the application of the concept to them, whereas colour properties causally invoke the response of colour perception. This is indeed the point at which the analogy breaks down. But here the cognitivist can utilise the mathematical analogy. Maths leaves us with no other way to go on but *this*, once we have become embedded within the practise of doing maths. It has, from a pragmatic analysis, as part of its make-up, a necessity embedded within it which is intertwined with the practice of maths itself. Whereas we are struck by colour as a category we simply cannot deny the existence of, despite the fact it is shared from some perspective, the internal relations in mathematics gives us a grasp on the idea that we are involved in working out the internal relations and the idea that those relations are pre-set by the necessary nature of the domain by our involvement. Alan Thomas (2006: 38) has usefully described the two analogies as reflecting, on the part of colour, an analogy with the category of value as a whole – a set of concerns that cannot help but make an impression upon us as being very real – whereas on the part of mathematics we get an analogy of a domain of thought which, once we have grasped it, has embedded within it deliberative standards which elicit a deliberative must or ought from us. Wiggins in particular draws heavily on the mathematical analogy to draw a distinction between a conception of objectivity that is heavily enforced by perception in empiricist terms and the type of perception that comes from a person occupying a particular perspective. The opponent he has in mind is Mackie.

Wiggins' analysis of Mackie is illuminating as it reads him as a moral constructionist, which is indeed the best place to situate Mackie, and explains the alignment between himself and Korsgaard in thinking that realism must be explained as a sort

of Platonism, implying that moral properties are real and would require a Moorean type non-natural sense for us to detect. Mackie believe that ethics was an attempt at describing real properties, but it described properties which are not real, but merely postulated as a response to the interests that we happen to have, such as living civilly together and co-operating on projects. The sorts of moral properties we 'find' are best explained as constructions which aid us as groups of persons. But they are not there in the way those objects which are apt for our senses to grasp are there. Thus belief that these properties are real is identical to the idea of believing concepts like witch and magic have corresponding properties. They are an invented set of prescriptive bonds which allow us to get along. Wiggins takes umbrage with such a conception of morality when seen through the filter of the error theorist or projectivist.

His first complaint is against the constructivist conception. If morality is to act as means to some end, he points out, Mackie should be able to delineate just what that end is. And yet, as Mackie goes through his book, showing us how variant concepts can be used in such various ways dependent upon context, mostly in order to bolster his own claim of relativity, what he reveals is the fact that morality as a whole has no singular end which is understandable from the outside (or the inside for that matter). Its concerns are various and manifold and with such an analysis the constructivist picture seems far less likely. For something to be invented as a means to certain ends those ends must be clear. But in fact, in order to understand the various places we discover ethical thought as being necessary, and how those change in various circumstances, requires that one already be a functioning member of a moral community.

The second point that Wiggins makes is against the idea of queerness. The 'queer nature' of morality, in postulating properties that are not available to the senses, is equally there in mathematics. In that domain of thought we start to understand of a structure of mathematics which is almost unarguably cognitively graspable but which does not exist as a fundamental component of the universe. Rather, it understandable only as a means of explaining how we come to categorise things in a way which is indispensable to our form of life – the counting of objects. No sense organ can detect 'twelve-ness' without one being armed with the concepts that come from the learning of mathematics. This is because, as Putnam maintains and Wiggins endorses, perception properly understood is an 'exercise of our concepts'.

The 'our' here is key, and points back to the idea that there are various examples of non-natural properties that are only painted out to be fantastical and in need of specially postulated sense organ if one buys the simplistic empiricist picture. Again, to use Putnam's example, one does not have a sense organ for 'detecting elation', but once one 'acquires the concept' one cannot help but find it in the world. The objection to Moore's work comes from the idea that all the properties which there are all the properties which 'are there anyway' – that is, all the ones which can be identified by independent means. But this opens the question of just why this is the case.

The argument that concepts which pick out properties which defy independent physical ratification is furthered when we consider how many properties unlike this are required to reason about those properties and come to any conclusions at all other than conclusions such as 'X property is there' or 'X property has these further physical features'. Reasoning takes agreement about fundamental background assumptions about levels of analysis, good theorising, what normative standards are required for good theorising and which indicate bad theorising based around a grasp of consistent assertions, inferences, circularity of thought and so on, which pay no homage at all to physical properties. And yet we can see the property of consistency and inconsistency in arguments and theories and we criticise these in those we think lead to poor levels of analysis. Anyone who argues against this point will be manifestly doing precisely that in their counter argument. The kicker with this argument is that an argument to the contrary is itself an exhibition that non-primary properties can constitute knowledge.

So the analogy with mathematics further gives us license to believe that morality has some chance of being made out as a domain in which truths and knowledge can be ascertained, despite the fact that its properties are non-natural. All we need to show here is that non-natural properties do not require some special sensory organ to detect, but that perception correctly conceived is not best thought of as explained in empirical terms. Mental content requires that we are utilising our concepts and the justification for those concepts are best thought as being postulated in the public domain. The argument from Mackie which stays on the record and which is the hardest for the cognitivist or objectivist to argue away is the argument from relativity. Here maths gives no respite. For convergence and the sorts of necessity we find in

maths is constitutive of its domain, whereas in morality this is clearly not the case, at least in the same manner. Disagreement is clearly rife, not just between various societies but within them as well. Overcoming this problem will be the one which will be telling for morality and ought to shape the way we think about morality.

Furthermore, it is a problem which is easily explained away by any non-cognitivist who is willing to take notes from the Wittgensteinian lesson.

The analogies with colour and maths thus only take us so far. What is not shared by either of these domains and yet is the mainstay of morality is that there is a motivational response, a psychological response, to the sorts of concepts and properties which are in play. It is this mix of description and prescription which is what makes the moral domain so distinctive in its constitution. As mentioned at the beginning of this chapter, the non-cognitivist who argues for a distinction between these two features of moral concepts in which there is a description made out in empiricist terms will struggle, as Mackie does, to explain how this idea of description can be successfully made out. It is not, I have argued, a hallmark of an objective property that it need to be instantiated by an independently verifiable structure made out in primary quality terms. Therefore properties can exist which defy this rendering. But the non-cognitivist needn't take this turn in particular. Description and prescription can still come apart if description is meant as applying to properties which are conceived as finding their verification not via their exhibiting independent empirical features, but instead as being ascertained from a perspective which implies a shared interest or set of judgements, where an appeal to an observer who is embedded within that form of life is necessary for understanding when that concept is being used well. The classic example will be games. Noticing that something is a game is noticing distinctive features which defy articulation in terms of reducing the concept to its physical properties. Similarities at this level evade us for certain concepts. Rather, one would have to be inside the form of life to be able to understand the connections, made out as something like Wittgensteinian family resemblances, which allow for something to be categorised as a game. But those various concerns do not have any active force on the will of the person who notices that something is game-like. This is not so with moral concepts. If morality is to be a non-reductive domain then, the evaluative or motivational aspect that is carried within thick concepts like honour, will be required to be a part of their meaning and

not tacked on non-cognitively. And this task requires that the cognitivist can indeed tell a story in which the desiderative nature of moral concepts can escape being categorised as (non-empirical) description plus non-cognitive like or dislike, which becomes psychologically active once the property is recognised in some agent or action, and is merely mistaken as belonging to the concept, such that it can be analysed away. Instead, the story must show that an agent who is able to reliably pick out the concept must have just these evaluative interests as a part of recognition of the concept. Minus the proper interest the concept will fail to have any semantic bite.

5.3: Moral Motivation and the Idea of a Standard

Thus, it seems we need a picture in which the appetitive element does not enter the picture at a late stage, as this will always invite a non-cognitivist rendering of the situation. And it is this picture which the cognitivist/realist David Wiggins paints for us, in both his 'Needs, Values, Truth' (2000) and his 'Ethics: Twelve Lectures on the Philosophy of Morality' (2011).

A huge player in Wiggins' account is the idea of how our needs and interests can manifest themselves into concepts of certain sorts. Drawing inspiration from Wittgenstein but focussing mainly on the work of Hume, Wiggins makes the case for conceptual practices which both depend upon and yet transcend their humble origins in our actual needs - which are made out as elements of our first nature. This picture allows those who share those needs and have access to the conceptual armament of some community, however broadly defined, to adduce conceptual 'ideologies' and 'ontologies', which are themselves seen as standards of justification and modes of reasoning which are both structurally dependent upon but explanatorily independent of, those needs and interests – standards that members of this community will 'steadfastly adhere to as if by second nature'. (Wiggins, 2006: 329)

The second nature is the level of analysis at which semantic concerns come into play. But highlighting the structural relationship here between the pragmatic and semantic features helps us to better understand just why reduction of what the non-

cognitivist will wish to say ought be reduced, namely the part of a thick concept which claims 'and this is a good thing', would be to misunderstand the nature of that relationship. The motivational and appetitive concern does not come after the recognition of the concept, as in the non-cognitivist picture, as an extra, but in fact underscores the distinctive nature of the category of the moral in the first place, as being one which necessarily invokes emotional responses from us and provides those who are within a fully functioning first order ethic to respond to certain properties, moral properties, with certain actions guided by certain feelings.

The non-cognitivist can be seen as essentially laying down a challenge. It cannot be true that honour is good because you desire it (the non-cognitivist claim) and that you desire honour because it is good (the cognitivist claim). But here the mistake is made that these two claims need to be made out at the same level of analysis. The additional claim that is that once this is done, the non-cognitivist rendering seems by far the more understandable. You come to see the property as you come to find any other property in the world, and then you additionally add onto this an evaluative extra which is not included within that recognition. The alternative seems to be that honour would have to have goodness built into it as a part of it. And yet it seems unfathomable as to how that would be so. Here we truly would have a non-natural property that seemed mysterious and that we may well need a special faculty to be attuned to.

But Wiggins denies the challenge made out as such really represents circular logic because he sees both claims as existing at different levels of analysis, with the levels being separated as a structural claim and an explanatory claim:

- 1) Honour is good because we desire it
- 2) We desire honour because it is good

Wiggins here makes the claim that the 'because' in each claim works differently, and this stops the idea of a circle in the illegitimate sense, and instead highlights the fact that there are two different sorts of analysis here. In 1) the 'because' should be seen as a demarcating a structural relationship between desire and the category of things we find good (one of which is honour). We necessarily need to have the sort of psychological elements of approbation and disapprobation in order for there to be at some point a manifest category of the good. These act as the marks of the interest,

which is then filled in by content to make a specific category. But just as our physiological and psychological make up need not enter into the category of colour, so too should the sort of motivational interests that demarcate the moral not enter into our category of morality.

In contrast, 2) is an explanatory relationship which is made out once 1) is already set up. Desire thus enters twice into the category of morality. Firstly in an unspecified form that, via our existence within a specific world, transcends that basis and become a standard against which we judge the actions of others. Desire here informs, as it were, the sense of the entire category of the moral. But this subjectivity does not enter into the specific judgements made by an agent who is in the position to grasp that category. Asked why we (as functioning agents whose second nature is so attuned to the standard) desire specifically honour, we will explain the desire in terms of the properties and concepts which are constitutive of the category.

Alternatively, asked why 'we' (as a group of human beings) come to find honour good, the explanation will make reference to the contingent nature of our interests as social animals, the fact we have appetitive natures, the fact that bonds of trust are important and so on. Again, *pace* those theorists such as Foot and Thomson, these considerations should not enter into the normative explanation of why we find a particular thick concept good or bad.

Thus, not just the 'because', but also the 'we' and 'desire', exhibit equivocations between their referents in the above supposed two point circle. And a circle only applies as a point of criticism if the two propositions are mutually self-supporting in a manner which ends up being fallacious. This is not the case here. Thomas (2006: 42) makes the claim that there is an irreducible virtuous circle of philosophical explication. But in some ways he is wrong. There is no circle once we have seen that matching terms which provide the content of each proposition ought to be seen as equivocations and thus two distinctive propositions.

This way of reading Wiggins' position should also not be seen as a sneaky way to manipulate an honest question asked by the non-cognitivist. It is supported strongly by an analogy with colour. There too we saw how the reduction of colour to its physical components would have caused confusion. Anyone who wishes to grant some level of perspectivalism should grant that properties need not be independently

characterisable in primary quality terms. Instead, within categories which allow for assertions to be made and convergence, and which do not seem to allow for their content to be easily reduced or explained in other terms – that is, for those categories which seem to develop an independently understandable ontology and conceptual network – we ought to consider the status of the property in terms of how we come to encounter them, and not judge our ability to cognise them properly by way of an analysis of the property in some independent manner.

This puts the non-cognitivist into an uncomfortable position. If she is to hold that perspectivalism does enter into the equation then she must allow that we can have descriptions that find their generality exhibited in the way it used by those who occupy the first-order position. This, it will be maintained by the cognitivist, will require that as part of the pragmatic explanation of this generality, she allows the interests of those persons to play a structurally explanatory role in their grasp of the language game. Now, moral concepts assume the sharing of a wider and more motivationally efficacious set of interests than more descriptive concepts such as game. The interest presupposes, once the sort of genealogical account that Wiggins has in mind is complete, that one will feel motivational compunction to do or withhold from doing certain things when one sees the property in question. However, to use this as a stick to beat the cognitivist with would require the perspectival non-cognitivist to disallow that these sorts of interests are relevant to the concept. But to do so would be to deny the distinctive nature of moral concepts as action guiding. Yet, as perspectivalism has already been agreed upon as a necessary feature of descriptions of all sorts, this blocking move would appear to be motivated by and pay homage to the sort of Mackie-esque claim that meaning is decided by purely descriptive qualities.

This move would push up against the idea that the contingent interests we as a species have can license descriptions of the game sort. Thus, the non-cognitivist would either be pushed back towards a position where supervening terms ought to find their generality in independently determinable primary features, or to draw an arbitrary line as to which sorts of interests can contribute structurally to semantic content. In the case of colours, no notion of motivation is carried with the concept. Thus the analogy breaks down here. But while this implies a distinction is to be drawn between how these categories characterise their respective contents, to see it

as a position from which to attack the moral cognitivist is to move away from the idea that we ought to consider the legitimacy of the property in light of how we come to grasp it. The non-cognitivist would thus be left with three options. Firstly, to conceive the property independently which would require generalising at the empirical level, which I claim with Wittgenstein is impossible. Trying to distinguish in independent terms why we use rude in variously different ways in various contexts would leave one either desperately postulating physical similarities which are impossible to make out (the same goes for terms like 'game' which would thus stop being classified as descriptive when no similarity could be made out) or to claim that there is massive error in various uses of rude and game, which is a disguised version of Platonism – that rudeness can be made out in a rules as rails format. The second option and the other pole of the debate, seems to be to accept the analysis of the cognitivist. The third option, and middle ground then, is to claim that certain sorts of descriptions such as game and colour do speak to categorising properties in terms of shared interests creating an independent explanatory ontology, but that specifically moral properties and concepts are unlicensed to utilise this model of justification because in so categorising its properties the interests which foster the categorising follow through to play a role as an active psychological feature in the decisions of agents.

As a quietist sympathiser, I can here see the point of the McDowellian warning that we are now in danger of creating a dispute that might be thought of as extrapolating meaning from metaphysical distinctions which have no merit. If a category of properties ought to be conceived in terms of our coming to grasp them, and the numerable concerns they answer as an agent from that first order position, then surely the answer shuts down there; there will be no further merit in postulating that distinctions between those categories should then imply a difference of status. There is something stopping me making this move, however, and it is this. The difference of the psychology of an individual agent as pertinent in morality – while it is not in the status of colour - reflects the fact that the moral reasons we draw are practical reasons, whilst both mathematical and colour reasons are theoretical. Someone might not have the best ability at picking up differences of colour, but the standard is set, quite neutrally in the theoretical disciplines such as colour recognition, by the best occupied observer. An agent may not see the difference between types of green, say, but there is no harm in saying he is just bad at colour recognition. The

standard is set somewhat holistically around the idea of an idealised observer in an idealised setting but can be formulated in terms independent of any particular agent.

Moral reasons, however, are practical, and this is reflected within the way we already do conceive of the standard. It is not set perfectly independently in the way that colours are. Colour charts exist and mathematical formula exist and they are standards that easily transcend, theoretically, the position of any one agent. An agent who disagrees is simply wrong against the independent standard. If one were asked to justify why they were in a belief state as to why they saw X as red, or why they took the answer 1000 as coming after 998 in a number sequence of +2, they would point to the object or the +2 as constituting their reason. This is the hallmark of theoretical reasoning. In the practical case, however, the motivational aspect of the reason which is required disallows this simplistic justification method without adding an extra; namely, something about the agent themselves. This is the distinction between internal and external reasons. Internal reasons invoke as a part of the explanatory efforts the motivational set the agent happens to have in a way that external reasons do not. Thus, the middle ground for the non-cognitivist who holds firm to the idea of non-cognition and also denies the empiricist rendering of description will be this: that the reasons which there are ought be external even if this externality need pay reference at some point to the fact that the agent occupies a form of life. The intuition this would grasp is one which pays homage to an intuitive notion of objectivity that was on display in McDowell's 'diamonds are hard' equation. In contrast, non-objective reasons would apply to someone who has the ability to understand all of the meaning that we ascertain from our form of life but who does not share, for whatever the reasons, a motivational set from which good practical deliberation will deliver the correct reason.

Personally, I am not impressed by this attack. It once again takes the form of pointing to a metaphysical difference which one can simply point to as a hallmark of the moral and tries to use it as a club to beat morality with. But, the differences can only be taken to be of note from outside the first personal stance. It is however a problem which Alan Thomas takes very seriously. He believes that the attack, which is mostly associated with the work of Bernard Williams (1985: 146 – 155), ought to be met if were to salvage morality. I believe that the argument stands and falls with the proper understanding of the idea of a transcendent standard. All the argument

does is to ask us once again to retrace our steps and to reaffirm once we have done so that, no matter the differences, the similarities there are in building a standard which arises from contingent human interests (whatever those interests may be) allows one to appeal to the standard - and not the abilities of any particular agent to conceptualise the standard via their own neurological make up, psychology or what have you – as independent. Once this is done then if we are impressed we will have given ourselves licence to say that the properties the categories pick out are real and that we can have knowledge of them. As far as I'm concerned this amounts to us having moral facts. There can be factual knowledge about moral properties. There will be further elucidation of distinctions between realms of facts and how these ought to be conceived. However, once we have the basic picture, a picture which Thomas shares (or to be specific, shares with the addition that we should take up a negative reconstruction of the same idea, a standard which could not be rationally rejected rather than one whose content is positively assented to) I struggle to see how the Williams inspired argument gets off the ground.

Williams charges that theoretical reasons are external while practical reasons, of which moral reasons are a type, are internal. This is because internal reasons 'sever the link' between explanatory and normative reasons. This is put forward with force and clarity by Thomas:

Consider theoretical reasoning directed towards the acquisition of true belief. A judger comes to accept a belief P, because this is seeing the matter aright, and one way to gloss this remark is that the judger comes to believe P *because* P. A philosophical understanding of this 'because' would take it to be both *normative and explanatory*: that given the norms of reasoning governing the acquisition of beliefs of this class, and the judger's proper application of these norms, there is nothing else for him or her to think, but P. In addition, from either a meta-level perspective or an explanatory one, an account of the same judger's cognitive economy would see the fact that P as being causally implicated in the aetiology of the judger's belief. In this case we have a compatibility between a justificatory and a causal account and a fully satisfactory philosophical understanding of the 'because'. (Thomas, 2006: 71 - emphasis is mine)

While this is of importance, whether it downgrades the claims of morality or merely situates their position once properly understood is yet to be determined. And I would claim that the whole affair comes down to where we consider the restructuring of what can count as an explanatory reason that we have been trying to foster throughout the last section as acceptable. I will explain what I mean.

In claiming that internal reasons sever the link between explanatory reasons and normative reasons, what Williams means is that there is an explanation of our coming to have the normative reasons that morality seems to bestow upon us, and which we believe we are reaching to when we justify a moral action, which does not end in a '...because P' format. Rather, in explanatory reasons terms, we believe that 'P because we are a type of creature who has just this set of motivational propensities raised in just this world at just this time with just these sorts of social influences...'. This explanation explicitly rules out the idea that the sorts of properties which moral reasoning must presuppose can exist within the world. However, this is now a picture which clashes with precisely the sort of claim that we have been making: that these properties are there in the world.

On the rendering I have been giving, moral properties fit the analysis of 'P because P', except that the belief of the property comes with the caveat that one has the eyes to see just those types of property. Granting this, the link between the explanatory and the normative need not be severed at all. That there is a separate explanation that can be made in terms where the property does not *really* exist is just, I would claim, recognition of the fact that all non-fundamental properties have a physical basis upon which they supervene. But this basis is not enough to enter into any explanations without there being explanatory interest from some perspective for grasping just those types of property. Williams needs a strong analysis of '...because P' here. Are we considering this to include descriptions of the purely empirical sort? If so he falls to the basic Wittgensteinian critique. This, I believe, is how Alice Crary (2009) reads him in her analysis of his absolute conception. Alternatively, given the focus on motivation and its link to prescriptions, is Williams allowing that descriptions which rely on some shared perspective might still be objective? For instance, would colour fit the format 'P because P', or would it fall foul by taking the following formulation in explanatory terms 'P because I am a creature who has just this sort of physical make up for detecting differences in light waves which strike me as both

phenomenologically different and explanatorily important once I am raised in a community which takes explanatory interest in these phenomenological differences conceptually in various ways'? If we allow that colour does allow the 'P because P' formulation despite the perspectival nature which is implied in colour descriptions, then Williams will have to contend with the arguments detailed above about drawing arbitrary lines between domains of thought which utilise different anthropocentric features to demarcate acceptable properties. If he does not, he will fall into the non-perspectival empiricist camp which has even larger concerns to contend with, as drawn out by Wittgenstein. If he is making a point about motivations then we can clearly leave him on the record as rightly showing that there are differences in the deepest fundamental explanatory terms about which parts of the human psychology and physiology are required for various properties to be conceptualised well. After all, as Wittgenstein showed, all judgments pay homage to an agreement of other judgments. To believe something is to stand in epistemic relation to it. And thus this sort of agreement plays a structural role within the explanation of coming to believe anything at all. Thus, even regarding the simplest formulation of 'P because P', where P could stand for as simple a proposition as you wish, the 'because' does not come for free, explanatorily speaking; rather, it disguises a set of norms and interests. The 'because' thus implies both norms and fundamental properties.

We can thus agree with Williams that 'P because P' in the moral case involves the inclusion of a motivational set as a part of the explanatory basis of the 'because', while disagreeing that this in anyway disbars moral properties as being there in the world. That is one analysis – the fundamental analysis - but this needn't be the correct level of analysis for moral concepts. This is because, and this is key, the entirety of moral sense, or rather the sense of the entirety of morality, is what is presupposed by ones having such a motivational set. Certain motivations are a necessary feature of one who grasps the point and purpose of morality from a first order perspective, just as certain physiological and neurological features are necessary for one to grasp the point and purpose of colour categorisation. Both categories rely on shared judgements and these judgments supervene upon physical features, but unless those features (characterised in precisely the way we characterise them) are to be made out as fully independent of our capacities and explanatory interests, then there is nothing against which the former properties can

be cast in an unfavourable light. The better way, I would suggest, is to conceptualise the world as made out of an unknowable amount of possible properties and that those that we do conceptualise will be those which are beholden to our explanatory interests. Those properties must always be thought of as supervening upon an independent physical world, but as soon as we cast our explanatory net over this independent world it loses its independence and becomes conceptualised by a group who do have explanatory interests and a limited way of conceiving and conceptualising those properties. The only way of conceiving 'the world' as a comprehensible whole is thus from a perspective. As it is the perspectival world not the independent world which grounds the concepts of truth, objectivity, rationality and so on which structure it; complete independence becomes a nonsensical position

This still allows, however, for us to differentiate between those explanations we see as more or less fundamental and those which are not. Here objectivity stops picking out what's there anyway, contrasted with the subjective which will include some blend of mind and world. Thus, that distinction must be given a separate form. And that seems to be, following our earlier discussion, that objective properties can be thought of as those which include reference to some standard impartially considered apart from any specific agent's viewpoint and which pay homage instead to being categories we cannot see as anything but indispensable and which have an understandable ontology and conceptual network which can be taken as explanatorily independent of the make-up of any specific agent. This point is crucial as it lays down a distinction between those theorists from the cognitivist side who make a virtue of the independence this brings and those who do not, those, that is, who remain open to the threat of the externalist/internalist distinction. The very fact of motivation implies that morality is internalist but also explanatorily independent.

The latter camp, I believe, contains McDowell. His attempt at creating a transcendent framework does not make use of the idea of a social standard rising up from but transcending a set of agents with certain concerns. Rather, he utilises the idea of the idealised observer or *Phronimos*. This is an extension of the idea that we saw when we were developing a standard. The standard for what counts as green is what would count as green to an ideal observer in ideal lighting. But this acts as only one structural move in developing the agreed upon standard and cannot replace the

standard once it is developed. It is the tacit agreement amongst the numerous people who take part in the standard which gives that standard its clout. Without the second pillar the idealised observer seems structurally unsound, for different persons with differing motivations/physiological make-up could herald a different ideal observer. To someone with quite weak moral grounding – and we must admit these people exist – are more likely to postulate a much less morally strict *Phronimos* than we think they ought. We thus would not be appealing to the same thing when were to justify our actions. There is no transcending here, and given that we have, with Thomas, settled on the idea of an internalist account which aims at objectivity by way of the development of an impartial standard, we should claim that McDowell's idea of an externalist account based on the *Phronimos* seems unlikely to get off the ground.

Morality, we have claimed, has the possibility for providing an explanatorily independent ontology. Yet there remains one sort of major issue left for morality here that can be made out via two formulations of relativity. Those standards seem to change vastly between different communities – say, Sparta in the 5th century BC and today's Britain – and also internal to certain communities regarding specific judgements. If the standard is impartial, why are there so many disagreements that seem impossible to solve?

5.4: Relativity and Pragmatics:

We have now reached a point at which the unfavourable comparison can do little work to make us worry about the status of moral claims. Subjectivity enters into the very core of morality at the level of a shared sense of what structures the conceptual terrain, and in so doing need not be thought of as infecting the picture at the level of reference. This expansion of where and how subjectivity can show up in accounts of domains of thought can then be placed alongside a corrective to the view of objectivity as truly being independent of us as thinkers occupying a particular perspective, even when we are considering the most fundamental properties we believe there are. With these two corrections in place we have reconceived the objective/subjective distinction in such a way that domains of thought can contain

propositions that can be both subjective and objective. Objectivity thus pays homage to the idea of an impartial standard, which can be thought of as a complex of ontology and ideology, and which while not strictly theoretical and as such contestable, can still inspire convergence between analyses. It is crucial to remember that this is the nature of the domain. Disagreement at the highly specific level not only, for good Wittgensteinian reasons, implies agreement at a background level, but also in many instances can be explained via normative reasons in such a way that each of the disagreeing parties can understand the position the other takes. Subjectivity can still be considered as the necessitation that at some point a conceptual scheme which pays homage to individual make up or community standards will, from an explanatory reasons point of view, be implicated within the explanation of how an agent reaches a reason. The remainder of this section will try and situate what form morality takes between these two poles. I have, I believe argued that there is enough for morality to be an objective category. However, I hasten to add, this will not mean that all moral decisions are objective. There is then still the philosophically elucidatory task of considering just how morality fits the bill in the case of specific judgments and how these different types of judgment ought to be classified.

What I will be trying to bring out here is how a final challenge to objectivity, borne from a final intuition about the same, which is embedded within this vision of moral objectivity which has been borrowed from the McDowellian/Wiggins school of cognitivism, ought to be deflected so as to give the moral domain its most realistic characterisation. The challenge is one of the most obvious and numerous mentioned in the literature: if morality really is objective (in our terms because it comes with an impartial standard that transcends the bonds of the directly subjective) then why are there so many variant answers to moral questions? How can we square the idea of an impartial standard with the undeniable fact that in so many cases people actually do disagree? Shouldn't the standard aid people to delineate between right and wrong?

I believe the above formulation works well as a description of the challenge at the inter-community level. There is however a second challenge, that of differing moralities between communities. And this might be laid out in a thought experiment.

Objectivity as made out via an impartial standard seems to imply the standard could be recreated even if (via some tragedy or another) human civilisation disappeared and then came back with only the most basic understanding of previous social institutions and our human nature. If this were to happen can we really believe that there would be anything that might resemble the moral standard we adopt? And if not, is the whole affair not too local?

I think both these challenges need to be taken seriously. There is a common idea that can be found amongst Wittgensteinian inspired philosophers that the latter question asks too much. Perspectival or local objectivity seems to offer a way out of answering it while, by way of leaning on the same perspective as a pillar of support, allows us to answer the former question by placing the objectivity of the domain as the background ontology which allows us to see the positions of those we disagree with as complete, defensible and reasoned conceptions of how the terrain ought to be organised, despite that disagreement. I am concerned that is offering too much to the relativist; that one would be stretching the terms of impartiality far too thinly if one were to claim that something is impartial only between, say, those of a particular small town or other group. This looks, *prima facie*, like bringing in a post-modernist rendering of truth and objectivity, where one group's truth needn't meet another's in anyway and thus the idea of truth and objectivity can be abandoned. To take this route, I would claim, makes a mockery of what we fought for in allowing perspective in the first place. To answer the above questions then, I want to make use of two ideas we have already seen. The former is a Wittgensteinian epistemological point; the second is a complimentary point about ontology of the moral domain.

The Wittgensteinian point might be put thus: for there to be such a thing as disagreement there needs to be a backdrop of agreement. The ontological point can be put thus: not all instances of moral thinking need to be categorised as wrong or right in any necessary sense. There are varying level of goodness and badness that one can reach to and which define the difference of a specific individual outlook and which we may disagree with more or less strongly but which we needn't see ourselves as morally obliged to put a stop to. We can call this domain the permissible. The deontic 'must' sort of moral claim might inform us contextually about how we can delineate other moral norms without running up against it and remaining a consistent outlook, but it does not do this alone. It does this alongside

various amounts of other norms that provide us with meaning; namely, those which we find in cultural institutions which moral norms are meant to partly inform. These are largely culturally enforced and may well take the perspective of a particular form of life in this narrow sense to understand, but they are informed by a development of what are necessary concerns of human nature, however contingent that nature might be from some perspective-free perspective.

One theorist with whom I share a lot of ideas, but who for me, transgresses regarding a proper rendering of objectivity by giving too many intuitions about objectivity away, is Alice Crary. In 'Beyond Moral Judgment' (2009) she makes use of the same sorts of Wittgensteinian insights which I do regarding changing the subjective and objective distinction to allow the idea of perspectivalism to carry the weight of objectivity. She does this via what she describes to be a pragmatic account of language. She believes that once one has understood Wittgenstein properly one ought to see that language is properly thought of as 'a moral acquisition'. Pragmatic features stop the idea that moral judgments are the full hallmark of morality. Whether a (moral) concept applies must take into account the pragmatic considerations of how we do utilise expressions, not just the meaning of the concept, as if it could be extracted and its meaning abstracted from the ways and contexts in which it is used. The learning of language, of how expressions ought to be used, thus impinges upon the propensities and sensitivities one utilises to discern beyond meanings understood as atomistic semantic definitions. A philosophical focus on semantics stops us from seeing just how far into our everyday lives our theory of value goes, and how it is best understood. To use Machiavelli as a guide here, one cannot semantically define the difference between the attribution to oneself of either 'untrustworthy' or 'cunning' by way of definition. It depends upon the context of ones actions, one's own position (and other agents' positions) within that context, the consequences the act begets, one's actual intentions, one's perceived intentions, one's manner, the success or the failure of those intentions, unexpected consequences of the action... the list goes on. These are pragmatic features which impinge upon how we think that a concept ought to be used, which do not speak to some conditions of that concept which can be separated from the rest of the world and brought to bear upon it when circumstances semantically laid out allow it to fit into place like a jigsaw piece to a whole picture. Rather, the picture cannot be seen

without one having developed (or been ‘taught’ – non-theoretically) the sorts of sensitivities and propensities needed to know one’s way around.

There is little to be disagreed with here. I buy Crary’s picture of pragmatic concerns as guiding complicated concept acquisition without argument. I have made many of these points myself. However, I do not buy completely what she claims follows from it. Crary believes two things which might be thought as startling claims follow from this. The first is that thoughts and utterances which make no sorts of mention of overtly moral concepts, and which are not judgments of any type, can still be moral thoughts and utterances. The second claim is this: once we have the pragmatic picture we cannot ask for such a thing as an ‘abstraction requirement’. I see the benefit in both of these ideas, however, without caveats I cannot find either fully acceptable, and taking them as Crary takes them, I believe, is damaging to an intuitive notion of objectivity which is necessary not to fall into post-modernism – something I believe that Crary is in danger of falling into at the end of her work when she applies her moral epistemological theory to feminist theory (2009: 164 - 191).

The difference between Crary’s analysis of Wittgenstein’s attack on the Platonic notion of objectivity and my own is her purposeful omission of an aim to a standard which is impartial but internal. When Crary talks of ‘abstraction requirements’ and ‘absolute conceptions’, towards the beginning of her work, what she has in mind is the idea which should be discarded by shedding Platonism: that the world ought to be thought as fixing the meanings of terms (‘maximally’) independent of any perspective. Thus her argument against the ideas of semantics as conceived minus the pragmatic features which contribute to conceptual application. Upon that we can agree. However, this means that what is determining objectivity is a perspective as contributed to by the interests of the people who occupy a first order position. And there is now no telling between the sorts of interests people happen to hold and the sorts they ought to be able to justify. Let us call what the standard gives us an ‘impartiality requirement’. If we are to understand each other properly, there must be something upon which we agree, which we hold firm to in order to be able to make points, namely that which Thomas calls a standard of reasonableness, however tacit this might be. Crary aims to provide the very same with only her idea of shared pragmatic sensibilities and propensities. These are what allow for the agreement which can be the base of disagreement. But I would argue this is one pillar which

holds up the standard and cannot be equated with the idea of a standard itself. A reminder regarding interest might be apposite here. On the picture upon which we are working a broad conception of interests underscores that which we share, and that which we share allows for the agreement which gives us access to the conceptual apparatus and ontology of the domain. Individual interests can vary, thus the disagreement (that which is constitutive of the character of the domain); but if these interests are such that a person who doesn't share them, despite otherwise sharing the background, can never deliberately reach the insights they provide, then interest is not just settling the domain at the broad level. It is also doing too much work at the level of specific, individual interest. I say 'too much' because this invites with open arms William's critique.

Crary puts to work her idea of the full blown rejection of an abstraction requirement in defence of a conception of feminist theory which walks the line between postmodernist views on objectivity (there is no such thing as objectivity) and traditionalist views on objectivity (objectivity is to be conceived as picking out features of the world which are independent of us). Between these pillars is a fully comprehensible version of objectivity which fits her (and so far my) conception of 'wider objectivity' – namely one within which we can find general levels of agreement by way of occupying a position which is grounded by pragmatic features – the concerns which we come to ascertain by way of being inculcated into a form of life which conceives the world by way of a conceptual network which we learn to utilise not merely by application of semantic rules, but by way of learning when these rules apply according to various contexts.

All of the above I agree with, but her wider conception itself is, in my humble opinion, too broad. The end at which it meets the traditionalist conception of objectivity is clearly demarcated by the idea of a rejection of abstraction requirements, platonically conceived. However, the side which touches on post-modernism blurs into the post-modernist critiques she criticises. It does so, I believe, because she places too much weight on distinctions that can be made between perspectives of inter-community groups specifically defined - namely, given the nature of her analysis, feminists. The same would apply if we were discussing racial theory or class theory. This critique is not, in my view, anti-feminist. Instead it correctly attacks a conception of objectivity that does not do enough to demarcate it from the properly subjective. I will soon

suggest a corrective to Crary's account which, while born from a structural ideal which works against it, will in fact strengthen the methodological chances of feminism.

What demarcates my own account is the standard which aims at impartiality and aims to find, from first order thinkers, those properties which we can appeal to which cannot be reasonably be ignored. None of these properties are thought of as fully independent from us, but nor are they thought as only ascertainable from some particular perspective which is imbedded within a wider perspective of understanding. That is, if one can grasp the general gist of another's complaint, and both those persons share enough background assumptions for them to also find agreement in various cases about the use of moral terms, by way of my account, there must also be enough shared between those persons that one of them could, *if correct*, show the other where they have *deliberatively gone awry* by way of the utilisation of the concepts which they share. I will leave open the means in which one might make this argument for now, although I certainly do not believe this deliberation need be a strict theoretical exercise. I believe that this much is implicit in any account which wishes to combat the critique of non-objectivists such as Williams. If this is forgone, then what we will be left with is internal reasons which do not even allow for an appeal to impartiality. And such an account is a fair description of Crary's account as it is laid out in *Beyond Moral Judgment*.

Everything here hangs on the idea that if one is to have a form of perspectivalism which still allows for objectivity, then the perspectivalness must be reined in, at some point, so as to not allow any individual perspective or perspectives to enter into the ways the owner/owners of that perspective is/are deliberately cast adrift from others who utilise the shared conceptual networks, such that those distinct perspectives cannot come to a meeting of the minds on the basis of good deliberation. If the concepts characteristic of feminist debates, say, such as 'oppression' and 'fairness', are not to start to take on additional instances of correct use, such that a member of an anti-feminist camp and a member of the feminist camp would be merely talking past each other when they utilise them, then there must be a way to deliberate about whether those terms are being correctly used in these instances. If we forgo this idea, we forgo a level of analysis at which we can say 'P because P' about moral terms, such as 'oppressive'. The correct analysis will

not be that we share enough of a motivational set to fix the sense of the terms, which then allow discussions, however contestable, about whether a certain set of actions, practices, institutions or what have you count as oppressive. Rather, the analysis would be Williams' analysis – that if someone does not share another's motivational set there may well be no way for her to ever deliberate towards a motivational set which includes disapprobation of those purportedly oppressive actions, practice and institutions. And this latter formulation looks a lot like non-objectivism. Moral truths would become very much ways of conceiving the world from certain identities like so: 'it is *true for* ... that oppression ought to be descriptive of ...' (where '...' stands for any agent or group of agents you like, and any action, practice or institution you like, respectively). It is the emphasised '*true for*' that I claim reveals the doorway to subjectivity inherent in Crary's account, and does so in a way which threatens the good parts of her Wittgensteinian analysis.

That Crary takes this view is made apparent in her analysis of the perspective one requires to see the truths made clear in the feminist literature.

If we are to fully understand considerations for identifying a particular abuse that women are made to suffer, we need to survey matters from a perspective informed by an appreciation of the injustice of sexism (2009:176)

and

Appreciation of the insidiousness of particular forms of gender inequality is a necessary precondition of a full understanding of the [wrongs] at issue (2009:175)

There is a clear way in which these are quite innocuous. One does require the ability to grasp the perspective in question for a full understanding. What is less clear, but implicit in the work of Crary, is that those who do not have the perspective may very well not have the means to grasp the perspective. Rather, the perspective is a pre-requisite for the grasping. This is what worries me. At no point does Crary challenge this assumption. If I am correct in her intention I fear we have entered a relativity rabbit hole. If I am not then there is no problem to be had. For in order for one to be able to grasp the extension would be precisely a deliberative process that moves one from one perspective to another by way of any mode of reasoning which is

available. And in so doing reaches out to grapple with terms, concepts, ideas, principles and whatever other deliberative tools are available to both converter and converted in order to inform the imagination and awareness of the latter so as to open their eyes to the possibility that the more nuanced and credible perspective allows for. Minus this, there are simply two perspectives which differ and nothing by which to say that either is better. It is this impartiality to perspectives that Crary leaves out and in so doing allows in relativity and methodologically hinders the work of feminist theorists by fostering the belief that those who do not share their perspective cannot be so converted. Conversions are not guaranteed, but they must be theoretically possible via deliberation if we are to keep some vestige of objectivity by disallowing the subjectivity necessary to underpin the ontology of the moral into that very ontology itself.

This correction regarding a distinction between abstraction requirements and impartiality requirements is one I believe accounts such as that of Wiggins and my own can offer to Crary. Frustratingly she does consider the idea of impartiality (2009: 198 – 204) and very nearly agrees with my own position here, noting that she is for the idea of impartiality as long as the approach does not “...insist on conceiving impartiality in narrowly rational terms” (2009: 204) However, being generally for impartiality is one thing, having it inbuilt into the structure of moral discourse is quite another, and I believe that this is what the standard offers us the chance to do in explicitly wide conception terms. What cannot be overturned is a general standard of concerns which give morality its very sense and meaning, a standard of the ‘unforsakeable’ in Wiggins’ terms and a standard which ‘is not to be reasonably rejected’ in Thomas’s. From hence, morality can take many forms down many different pathways of different lives inspired by different experiences and cultures. But this standard locks down the idea that if morality is to be anything, it better be the same thing, in at least some small way, wherever it appears.

Moving on to Crary’s second startling claim, I believe another correction can rightly be made, or at least a reconceptualization of what I believe she is postulating is called for. The second startling claim mentioned above was that thoughts which were not judgments and which made no mention of ‘moral topics’ could be thought as being moral. This is grounded upon her claim, from her Wittgensteinian analysis, that all our concepts can be thought as practical because different concepts belong to

various language games and between different games, and different contexts within the same game, both the meaning of the term and the criteria for its suitability can change. These differences cannot be easily theoretically tracked and included within the definitions of the concepts, and so are best seen as underpinned by pragmatic features of our language acquisition which are grasped by a practical sensitivity once one is inculcated. As part of our interests and sensitivities are properly thought of as moral, and as interests cannot be semantically demarcated, this means that language is rightly thought of as a moral acquisition.

My complaint with this claim is that it places too much emphasis on the moral. Crary does not merely claim that an agent utilises practical sensitivities and propensities in all areas of language and also in morality, but that we are 'justified...in representing these same sensitivities and propensities as *composing* her moral outlook' (2009: 2-3) Whilst I agree with the former claim, that pragmatic features of language acquisition, in some broad conception, are necessary for one to grasp all areas of language – this is just the rejection of the Platonic semantic claim - I cannot agree that the very same propensities and sensitivities which ground understanding the point and purpose of, say, describing the formation of electrons and protons around an atom, are the same sorts of sensitivities and propensities which allow one to delineate courageous and foolhardy actions on a battlefield. Getting the former right and the latter right require distinct sorts of analysis and distinct sorts of interests and this must be the case if we are to distinguish between moral concepts and other sorts of concepts. But Crary explicitly rules out she is looking at moral judgments (such as which concept applies to which action) and instead says that thoughts that are distinctly non-moral in conceptual terms can be thought of as moral. The question becomes what she means by this.

One interpretation would be that all practical discernments are moral; that the moral and the practical are indiscernible from each other. This position has some merit but I am not sure that it survives reflection. A set of psychopaths who had no shared feelings of interest which could be traced back to benevolence may be very different to us, but we could surely still see that they had practical reasoning. They would have ends and means of attaining those ends, and upon knowing those ends we could rightly say they had reasoned better or worse than others. Morality, contrariwise, includes the invocation of certain ends which act as bonds between

each other and prohibit certain behaviours. If I am dehydrated and in the desert and I see water, discerning the quickest route to that water does not seem to have any moral inflection to, as I am benefitting an unconsidered desire and no one else is affected.

However, in most cases I am not so isolated and in variant ways my actions will interact with things that, upon reflection, I can see as having a moral dimension. If this is the point that Crary is making - that once we have ideas of benevolence thrust into our set of rational concerns, then even decisions which themselves contain no specific reference to moral concepts can still be reflectively or from a third-personal perspective morally judged – then I have only no quibble with her claim. Morality is not just about what we do, but also how we go about doing it. Helping is selfless, but helping begrudgingly or with constant reference to how helpful one is can be classed as self-aggrandising and prideful.

Backing up this second conception are Crary's various examples of literature.

Regarding Austen's marvellous *Pride and Prejudice*, Crary gives examples of various characters with various character traits, showing how their everyday and seemingly non-moral actions and preoccupations which include, "Mrs Bennett's culinary arrangements, the prospects of her daughters marrying, Mr. Collins' social habit, etc." can "bring us to a good understanding of these things by getting us to reflect on them in a manner informed by the shift in our modes of responsiveness...the shift that...is internal to an appreciation of what pride is like...and further, that this shift in the novel's central moral teaching" (2009: 143)

Pace Crary, however, the shift does not indicate there is such a thing as a stretch of moral thought which pays no homage to moral topics and includes no moral concepts or judgments. Rather it pays direct homage to what must be seen as a pre-existing idea of pride and its various extensions. It may add to these extensions, or alter one's perceptions of them, but the 'mode of responsiveness' which allows for this analysis is precisely a moral stretch of thought that has us bring pride under the microscope. We are expanding or specifying how and when we ought to judge within the mode of responsiveness – as such, the mode is presupposed by that very expansion or specification. Not all moral judgments really focus on concepts themselves. Rather, in judging, we bring moral concepts to bear on otherwise non-morally conceived items. It is every day actions which are thing under analysis, but

the analysis is dependent upon a shared mode of representation which includes that one has concepts, such as pride, at their disposal. Rushing into battle with a sword in your hand is just a description, if you like, of a non-moral action. Sure, it is possible that one might think, 'this is my chance to be courageous'. But one might simply think, 'this is the best way out of this mess', or 'this is my order', or 'that hill needs to be captured', or what have you. Cray has not given us access to moral thoughts which are beyond moral judgment. She is describing that which is common to almost all Wittgensteinian commentators regarding what moral judgment consists of – namely, that perception involves the exercise of concepts, and that seeing moral concepts in the world requires being imbedded within a shared moral outlook.

This *is* moral judgment. We can cast it upon *almost* any situation we come across. We look back at our actions and wince at our cowardice or brashness and feel pride at our courage our humility. We sometimes first-personally try to decide what the best action would be in a way which takes into account what it really means to be a good person here, what would be the most courageous route. But we often just act and assess why we acted later on. Reasoning so often comes post-action and our motivating reasons are often disguised to us until we do act and reflect. This is a psychological fact of people. But I cannot buy into the idea that the analysis of a man who runs from danger and thinks 'that's the safest route' or to go back to Austen, Mr Collins' overtly practicality-focussed marriage proposal to Elizabeth Bennett, should themselves be seen as moral stretches of thought. They are practical stretches of thought which – like almost all practical thoughts and actions - can be assessed morally. The thoughts or utterances in question *exhibit* the moral character of the actor that can be assessed as moral, but I can't say they are themselves moral, they are merely indicative of a way of being. I use 'way of being', and not 'moral outlook', because the second would assume that the way that one treats people is the way in which one believes that they ought to treat someone. And these are so often distinct. Akrasia, anxiety, self-destructive tendencies, sickness, intoxication and numerous other issues can intervene between a person's instinct and their reflective considerations of what we owe to each other. And it is the second which, I claim, are moral stretches of thought. Moral thought is the bringing to bear of moral concepts upon a situation, either prior, during or after it. It is not the characters doing everyday

tasks which have moral thoughts in *Pride & Prejudice*; rather, it is the characters judging them - either third personally or reflectively - and the reader.

I started the section by claiming that two ideas would help us to deflect the worry about the indisputable fact that we disagree with each other all the time about specific moral judgments. The first was the Wittgensteinian point that for there to be disagreement there must be some agreement. The second was the understanding of the structure of moral thought between the deontic and the permissible. I think what Crary gives us with her analysis of practical reasons being viewed as open to moral assessment – the reminder, as it were, that morality is not merely about knowing what to do but the manner in which we ought to go about doing it - is a fantastic rendering of how we ought to conceive the permissible and the sorts of cues, examples, deliberation and reflection we use to conceive this moral domain. I have quibbled, rather than disagreed, with the manner in which she describes her account. But her main achievement is showing that for people who share a background moral literature can act as a guide to nuanced moral philosophy which distinguishes by ways of examples seen through full narratives the ways in which lives can be lead. This is the perfect expansion upon the clumsy thought experiments of the generalist, trying to get us to ascertain which principle to go by. In good literature, as in real lives, we have all the context we need to truly ascertain what the best courses of actions are.

My main gripe with Crary has been over a conception of concepts being held hostage to certain interests which themselves seemed worryingly out of reach of deliberation. The idea that the cold hearted may see no wrong in the actions of Mr Darcy when he is first introduced, while those of us with a better grasp of the sensitivities which underscore a 'better' grasp of moral evaluation find ourselves perplexed and infuriated by his overlooking what he could share with Elizabeth. When Darcy learns this, it is presumably because he has within his motivational set the ability to see what is truly good; he sheds the set of concerns of status and privilege which then gives him direct access to what is truly important. The important point that Crary does not rule out, but which I believe she should explicitly rule in, is that the capacity which Mr Darcy clearly had to see these truths, despite his inability to recognise it at a certain time, marks the same distinction that should be the analysis of the psychological states of most people who seem to have a blind spot in

moral perception in some area. The capacity is constitutive of a set of interests which allow us to understand morality at all. This is an optimistic outlook but I believe the correct one. Those with whom we disagree, and ourselves when we are blinded to moral truths, may be unable to see some specific moral truth here and now, but in many cases we have the capacity to see them if we were to deliberate correctly – and, crucially, that this deliberation may be a matter of showing and not telling – such that only upon moral reflection could we have access to a conceptual repertoire so as to say discursively what it was that we were previously missing in our analysis of the situation. The shared perspectives give us the capacity (not the innate ability) for making correct moral descriptions but it does not guarantee them. Rather, there are then various ways of perceiving how these and other concerns are rightly ordered. But someone who does not share our quite particular perspective of a particular situation in this manner still shares the capacity to be able to be brought to do so. Again, most of the work here is being done in the permissible; someone (an outsider) who shared almost none of our more firmly held deontic ideals might not be able to be brought to do so. But this only proves that they, like the psychopaths from earlier, do not have the capacity. Like a colour-blind person, they simply lack the eyes to see the domain at all.

If concept and interest are necessarily tied together then there is little hope of the cold-hearted moving past their original position. McDowell takes a similar position it seems. This in turn only adds to the idea of quietist retreat. There is nothing left to be said of those who grasp a set of concepts other than that their interests were so aligned as to be able to grasp them. For McDowell moral education ‘good upbringing’ and assort of conversion are the best chances one has of going from not getting morality to getting it. This, I believe, is a mistake and places far too much emphasis on the motivational set one starts with. This a very broad idea of internal realism which seems shared by Crary. Those who grasp the insidiousness of sexism have access to conclusions regarding the wrongs of various practices (which those who do not share that perspective simply do not). The idea of a standard moves us beyond this. When we are saying with Wittgenstein ‘this is just what we do’ we have to make clear just who is included within this we. Is it us locally, here and now, or is it all people who could be seen as open to these reasons? Surely the former forces us to collide head on with an intuitive notion of objectivity? I believe that the latter is the

grounding for a defensible realism. It is a much larger project with more to explain than the internalism which invites quietism. It also satisfies more of our intuitions about what objectivity might consist of us, however.

It is also a question which moves beyond the scope of this thesis. As a quick reminder, this thesis is looking to understand the form good ethical discussion should take, and see how this can be implemented into the medical ethicist literature. Defending against the twin pillars of generalist theories which abstract too much from our everyday moral lives, and non-cognitivist iterations, which makes us lose faith in the objectivity of the concepts we use, is the aim of the thesis. Evading the former required we allow there to be moral knowledge which is neither principled nor commensurate. From this we put so much weight on the utilisation of thick concepts. The latter requires that these thick concepts are not best analysed as a matter of description plus a non-cognitivist, appetitive extra. Whether the argument for objectivism can be made out without an appeal to quietism at some point shouldn't affect the character of moral discourse which I am suggesting would give a more apt rendering of medical ethical discourse than we usually see today.

I will now turn back to thick concepts and look at some of the more recent literature which has arisen from the McDowellian analysis of inseparability we countenanced earlier on. This will encompass various attempts to show that the non-reductive analysis of moral concepts fails, and to try to both rebuff these attempts and utilise them to indicate how we should best conceive of moral reasoning and moral concepts.

Section 6: A Return to Thick Concepts:

The vast majority of what I have wanted to make out as a positive case for the metaphysical concerns about morality has been said. However, as we started this Wittgensteinian analysis with thick concepts and their important position in understanding morality, it makes sense to return to them for two reasons. Firstly, the McDowellian explanation of thick concepts I proffered ought to be able to be bolstered by what we have discussed since then. Secondly, the disentangling argument has returned as a focal point of discussion in the past decade. There are new challenges we have to meet before we can feel confident that we can move towards Medical Ethics with a non-reductive rendering of morality and see what sorts of corrections we can make there. The newer challenges will also be supplemented by older ones. I previously utilised the argument as an entry point into a conception of Wittgensteinian inspired cognitivism. Now we will see if it can rise to the challenge.

The first thing that should be said is that the disentangling argument is the strongest attack that the cognitivist has against the non-cognitivist. The non-cognitivist who wishes to separate the two components (of description and evaluation) can look for times at which we do seem to utilise only the descriptive component as a way to challenge the account. But as cognitivists it is not enough to point to examples of when the two components seem to work together, as this too can be given the non-cognitivist rendering of each component working independently. Much of what will come then will be reactive. There are some positive points to be made however. A further methodological point is whether such discussions give us the resources to determine the debate. After what will be said here I am highly sceptical of this. We have two different conceptions of the terrain, which are born of a disagreement about the direction of analysis, which in turn is born of different conceptions of objectivity and meaning. My own diagnosis is that this is a reflection of the fact that we have two modes of investigation with anything at all: first and third personal. The former represents normative reasons, and the latter represents explanatory reasons. Each has a claim to priority in various ways. I do not need to settle this debate however, I simply need to defend the idea that a non-reductive morality is the least distorting

way to represent morality and show it can sustain any immediate threats from subjectivists.

It is important to state what will achieve this, before we move onto defending the position. It seems to me that McDowell's disentangling argument is sometimes conceived as making much stronger claims than it actually is. What it does not claim is that an outsider can never grasp some thick concepts with ease. This is actually just assumed in the fact that thick concepts are world guided due to their specificity. It is the 'world guidedness' that makes them so interesting. From this it should be clear that there are indeed descriptive extensions to these terms which an outsider could grasp in many cases. What *is* claimed is that mastery of a thick concept would be impossible for an outsider without them sharing some of our *evaluative* interests; that some terms and how they are used in some specific instances (yet to be specified) would evade explanation, except for someone who has been inculcated in the right way.

Secondly, I am defending my own account and not McDowell's, despite agreement with his argument. Mastery of a concept is less tough to grasp from my point of view. It does not imply the *Phronimos* for instance. There is no perfect agent for whom all other reasons are silenced when thick concepts are present. While the lower demand on mastery may seem to be an advantage for the non-cognitivist, it will turn out that a more realistic view on morality and the defeasibility of the reasons therein are an important methodological point for denying their attacks.

Thirdly, and most importantly, we must really decide what we have in mind when we say 'thick concept'. As noted by Dancy, there is a tendency to demarcate them as two components stuck together in such a way that they cannot be prised apart. I would say that this is an unfair advantage for the non-cognitivist.; although I think in some cases this is a fair assessment. In certain cases it seems to me that the descriptive element is so intertwined with the evaluative interests caught up in a form of life that it cannot easily be separated out, and may in fact provide the very differences between those cultures. Such norms would sit as an arbiter, in the deontic position of eliciting necessity from those within the society. I have in mind such concepts as honour in an honour culture such as Japan pre-twentieth century, work ethic in protestant nations and (if one can take an outside perspective for a

moment on our own culture) personal freedom and autonomy in most modern day Western cultures. It is important to remember this does not rule out the idea that if we saw one of these cultures as flawed for holding these norms in such high regard that there would not be the conceptual resources available to allow them to deliberate to the conclusion it was indeed impeding their overall moral growth.

The above also invites a second rendering which Matti Eklund (2011) has also made – that the idea of a unified account of thick concepts may be no more than a Chimera. This would stop the simple sort of three-level metaphysical analysis of the non-evaluative, the thick and the thin. The thick as we see it now would be an umbrella term for various concepts which could be variously analysed. There are further reasons for thinking this three part analysis is simply a mistake, and it will act as my first positive attack on the non-cognitivist rendering as a whole.

Much of the thesis has attacked what might be thought as the descriptive side of the non-cognitivist ideal. We have seen various arguments about why it dangerous to characterise the non-evaluative in such a way as to make in non-perspectival, such that it follows the format 'P because P' because of how the world is and not because of how we are in some way - as though the 'because' does not hide an amalgam of pragmatically driven normative assumptions and so on. But the other side of the non-cognitivist description of thick concepts also needs to be addressed. As Harcourt and Thomas (2013) have pointed out, if thick concepts exist at all, then their two components ought to be separated into two distinct *concepts*. But this precisely what the non-cognitivist says isn't the case. The descriptive side can be, but the evaluative side is merely a psychological addition on the part of the agent. Rather, there is a weird shared tendency for us to attribute similar attitudes towards some descriptive concepts. When we consider what we have said in the thesis so far, this really shows the gap in the analysis of these concepts between the two parties.

When we discussed the work of Judith Jarvis Thomson (1997) **3.4.2**, the conclusion was that she made the externalist's mistake of trying to make morality align with facts by claiming that the interests which were served by our having a certain morality could directly be brought to bear on the evaluation of the concepts which arose from those interests. The correction about motivation that we have made is recognising the importance of those interests, while simultaneously barring their

involvement from meddling within the morality that emerges from them. Rather, they are the basis for the presupposed subjectivity in our account. They are a necessary pillar for an independent explanatory basis for morality to emerge. The good point of Thomson's was her analysis of the word good, where everything was seen as good in a particular way. There was then, no such property as just 'the good'. This move can be utilised to close the door on Moorean analyses of goodness. But it brings with a secondary point crucial for any cognitivist rendering of thick concepts. They are not merely descriptions plus thin concepts. Rather, they tell us the way in which they are good by the position they hold in our moral lives. We needn't ask why courage is good as though the goodness could be seen as an independent property, the conditions of which courage needs to satisfy. Being courageous just is one way of being good.

This has implications for the demarcation point of the thick and thin. On Thomson's account, thin judgments are world guided as well, it is just that it has been left unspecified by the speaker the way in which this is good (for making cheese cakes, for my psychological welfare or whatever). This means, however, that the thick/thin distinction is not so much a distinction as it is a spectrum of abstraction and specificity regarding how much world-guidedness is present within the descriptions given. This is at one with argument that Scheffler (1987) makes in his review of William's work. I wish to make two points about this.

The first is to do with an analysis that most non-cognitivists take up in the centrist debate (Hurley: 1989) namely, 'thin-centrism'. This gives analytical priority to the category of the good and bad, which are then meant to be fleshed out with thick concepts which are relevant to the specific concerns we have. But again, if Scheffler and Thomson are right, thin centrism is in serious trouble. There is no good and bad without there also being something for the good and bad to attach to. This underpins a strong idea that there can conceptually be no such thing as a category without concepts that fall under it and no such thing as concepts which cannot be placed into a category (which provides them with meaning). If neither demands priority then it follows the two can be thought of as interdependent.

The second point is much simpler. The spectrum implies thicker and thinner properties. Thicker properties, that is, the more specific properties that we have, will

be far more likely to take into account the more specific features of our form of life and may be the very thing which causes disagreements regarding ourselves with other cultures. Genealogical explanations of interest and concept often reveal that the most dramatic sorts of moral disagreements that we have with other cultures are born of traditions within the two cultures which do not mesh and which infect the moral standard with ideas that can be separated from it. Beliefs in certain sorts of deities and the sorts of narratives which get told within religions and the sorts of values that the religions highlights can be a prime example of this. If one sheds the religious beliefs (a decision that can be made on grounds quite separate from the moral sphere) one can shed the sorts of ideals which the religion inputs into the moral terrain. This aligns with the idea of an impartial standard quite nicely. What makes the standard impartial is that it brings to bear as little other concerns as possible. The concepts that we find again and again (despite the variance of how they are expressed) will be the ones which pay real heed to central human interests, rather than the interests of people entrenched in highly specific institutions which can interact with moral concerns in odd ways.

6.1: Variability of Valence in Thick Concepts

With all that said about the various ways to conceive the terrain, let us move now to some attacks upon the non-reductive cognitivist position. The two most common are those which threaten the disentangling argument by way of example. As Väyrynen (2011) rightly points out, one needn't be a non-cognitivist to be a 'separabilist' but in most cases the trend holds. I will only address separabilists who are arguing against the cognitivism I have developed. Further, whatever the reasons for separabilism, it is not something that I can accept. I am committed to the thesis that - as far as concepts can be thought of as providing any semantic meaning when extracted from a wider form of life - the meaning of thick moral concepts provide us with a reason for action. This formulation seems to imply that thick moral concepts have an invariant direction. If I assess, correctly, that the concept barbaric applies here, then I am given, by my lights, a (defeasible) reason to in some way or another disarm or lessen that barbarism. This chimes with what Blackburn has made the case for (Blackburn:1992), namely, that if we both claim to see barbarism, but you see

nothing which gives you reason to act, then if we are not to be talking past each other in our use of this concept, the reason giving aspect of the concept cannot be doing the work. This is the non-cognitivist case. We can isolate the very content of barbaric minus the evaluative aspect which inspires the reason giving aspect. The cognitivist must show then how the concepts can both be evaluative in their very meaning and that despite this it is possible that recognition of the same concept seems to allow for different evaluations between cases (or agents).

And indeed, there are examples aplenty of times when a thick concept seems rightly to apply, but in which either it is seen as giving us no reason at all, or even the opposite reason – that what we think of as the usual attitude changes valence, and negative thick concepts like barbaric seem to inspire a positive attitude towards them, or usually positive thick concepts seem to inspire a negative attitude. This seems to imply that in each instance, we can extract the world guided properties as supplying the meaning minus the attitude. So barbaric would stand for some descriptive property, say, ‘overt and unnecessary display of violence’. Where this is found in the world, one could apply the concept barbaric, and yet whether this would be good or bad would be open to evaluation post-recognition. Thus, barbaric does not imply ‘bad in a way’; it merely implies the presence of a property which is open for a non-cognitivist element to interact with it. Let us put this on record as a counterexample to inseparability:

(1) it was barbaric what he did – and all the better for it

This, however, seems to jar with our sensibilities in such a way as to require an explanation or context which would allow for this to be the case. Non-cognitivists would likely be accepting of this fact. Their claim is that between contexts the evaluative valence of the concept can change whilst still keeping its meaning. So, keeping in line with this:

(2) The boxing today was better than last week. Last week it just wasn’t barbaric.

Now we have a context which seems to provide the exact sort of counterexamples a non-cognitivist would lay out. We have a thick moral concept which has seemingly

changed valence from how it is usually used. Barbarism then seems to be a concept that can shed its evaluation and keep the same meaning.

Or can it? In (2) is the word barbaric picking out the same conceptual extensions as it does in its usual moral contexts? More specifically, does the use of barbaric in (2) indicate anything like a *moral* judgment? If it does not, does it pick out a *moral concept* at all or something (in this example) closer to an *aesthetic concept*? And if it is the second, can it really be said – given that each use is constitutive of separate domain of conceptual assessment - that the two uses of the term barbaric pick out one and the same concept?

What is at play here, I posit, is a distinction between a term and concept. This is something we must be wary of because of the complex ways in which we can and do use words. As a (less) ethical example, let's use a non-moral term: miraculous. I am an atheist and do not believe in any way in divine intervention. However, I make use of the word miracle. I use it to denote that something highly unlikely has happened which makes me happy, and in such a way that if I did not know better, it could be thought some deity or guardian angel was 'looking out' for me. This shares much of what is meant by the term when it is said by someone who is religious and truly believes that when, for example, they overcome a serious health complication that they were statistically expected to succumb to, that they were indeed being watched over by a deity who really *did* care about their livelihood and existence. Given the additional supernatural element, are my religious friend and I using the same concept? I would say we are clearly not. We literally *mean* various different things by using the same term. We may not hear it in each other's talk. Perhaps when I tell the religious person I was subject to a miracle, they utter something like 'hallelujah' – taking me to mean precisely what they mean when they use the term. But our uses are not co-extensional. They pick out different items in the explanatory terrain. They indicate different terms for correct application, different sorts of assessment and different connections to other concepts and would give rise to different conclusions (such as their affirmation of that deity and their own blessedness). For example, their miracle and my miracle have distinct relations to 'lucky'. My use is essentially co-extensive with 'lucky', whereas they would have to decide whether they were subject to a miracle or 'merely lucky'. As a final point, it is clear that my use of the term is not

incorrect as such, but parasitical on their use. Miracle couldn't be used as a hyperbole of lucky if it weren't for people who do or did believe in the fuller concept.

I think this problem is underrepresented in the literature. And because of that, once we see it, we see how redundant so many of the supposed examples of valence change in the use of terms actually are.

Another example is nicely adapted by Väyrynen (2011: 8) from an example given by Blackburn (1992: 296)

(3) The carnival was a lot of fun. But something was missing. It just wasn't lewd. I hope it'll be lewd next year

This example falls to the very same explanation as (2). Lewd (if we are to allow it as an active thick concept) is supposed to act as an aesthetic complaint about the carnival and not as a moral assessment of it. One could of course use lewdness as a reason to not go to carnivals at all. But at least one part of the very essence of (many) carnivals is that they normatively offer a space in which one feels licenced to act in unconventional ways without suffering from moral judgment. Boxing matches do the very same for brutality. Strip away the moral aspect of the concept and one is no longer discussing the same conceptual terrain. I am not claiming that I want to be let off with being immoral in this instance. Rather, I wish to be titillated and believe that in this context that titillation is not of moral interest at all. My challenge to someone who made a moral argument against boxing for being too brutal would neither deny their use of the term, nor claim that boxing was morally better for it. The moral argument would likely be on the basis of the participants' personal autonomy. The brutality present in boxing, I would claim, is a non-moral description of the events due to the context in which it is set.

6.1.1: The Concept/Term Distinction and Dancy

Sticking with the Scheffler and Thomson conception of the thick/thin 'distinction' then, 'morally bad' is a less specific way of saying (morally) brutal (and possibly (morally) lewd - but we will come to this later). If one worked from brutal towards goodness in the opposite direction, however, one would not arrive at 'morally good'.

The 'goodness' here would be characterised as entertaining, therapeutic, aesthetic or what have you. This, it seems puts me at odds with what has been a hallmark of particularists in overcoming the valence issue in the past couple of decades. One form of particularism - but neither my own, nor I believe that of Wiggins or McDowell - is the one that has been advanced and defended by Jonathan Dancy (e.g. 2004). From the shapelessness criteria of the evaluative, which states that evaluative concepts can't be given a non-evaluative shape, Dancy gleans the thesis known as the 'holism in the theory of reasons' (ibid: 7). This thesis states that a 'feature which is a reason in one case, can be no reason at all, or even an opposite reason in another'. I want to explore how his idea about the holism of reasons works in terms of thick moral concepts.

Dancy clearly thinks that thick moral concepts *are* reasons. They can be used to *describe* events or agents, but even in that description they are supposed to supply us with a reason for celebrating or condemning an action. They are more clearly reasons when they are given in a first-personal explanation for action: "I did it because it was fair". Following his holistic thesis about things which are reasons, Dancy says that thick concepts, like features, will be able to work as reasons for, against, or not be reasons at all dependent upon the context. The changing of polarity is rare in thick concepts, says Dancy. In most cases they are invariant, but there is still room for variance. 'Sometimes lewdness is just what is called for' (2004:122). So his holism does stretch to thick concepts, even if in many cases there is 'interesting invariance' in thicker concepts. I believe that this possibility of invariance, if it is not only held in our thickest of concepts ('dishonesty') is dangerous in that it threatens the intelligibility which generality secures.

I do think we can diagnose and remedy the flaw here, which is most evident in his account of conceiving non-evaluative features as reasons (i.e. 'he kicked him' or 'breaking a promise'). The flaw is that there is a distinction between 'is a reason' and 'can be cited as a reason'. Almost anything can be *cited as a reason* in a case where enough is known about the situation, such that one's interlocutor needs no more information as to ascertain the full reason. In *Ethics Without Principles* (2004), Dancy comes up with a complex explanatory structure which is meant to mirror cases of deliberation such that things that would usually be counted as reasons against ('kicking someone') come out as reasons for – or not reasons at all. Usually

speaking, kicking people would be cited as a reason to hold someone to account, in some cases it could be cited as a reason for praise (the person who got kicked was about to do something truly heinous and the kick stopped that) and sometimes it is no reason for judgment at all (in an arranged kickboxing match between consenting adults kicking one's opponent is just expected so is neither morally good nor bad).

My claim is that it is the uncited part of the reason – the same background which allows the description, 'he kicked him' to act as enough explanatory information for a well-placed interlocutor - that is acting as the actual reason. As *ex hypothesi* there is not enough generality at the non-evaluative level for someone definitively to grasp the evaluative meaning behind the cited description, there must be, if moral reasons are to be communally explanatory, generality instead at the evaluative level. If there is not we are seriously in danger of losing grasp of the explanatory force of reasons. At some point there must be some shared sense of judgment which allows explanations to be intelligible to others, otherwise there would be no stopping point for explanations. 'He kicked him' and 'he did so in a cruel way' could both elicit the question – yes, but was the act morally bad? Only a verdictive judgment would therefore stop the chain of explanation, and at that point the chain seems to be a first personal interpretation, from non-evaluative feature to thick concept, and from thick concept to thin judgment. This only further opens the door to non-cognitivist interpretations.

I agree completely with Dancy that we cannot have principles at the descriptive level of features. Shapelessness is part of my argument as well. The issue I have with Dancy is that his account threatens any generality at all, even between the application of thick moral concepts and verdictive judgments. Interesting invariance is not the same as something's acting as a definitive reason without need for further interpretation. This leads to the problematic position that (variant) thick moral concepts are simultaneously evaluative (the position needed for non-reductive cognitivism) and yet can act as reasons for, against or be no reason at all when they are present (the argument needed for holism in the theory of reasons to apply to thick concepts). Dancy seems to be committed to an awkward situation of saying that thick moral concepts are evaluative alright, but the direction of evaluation, of how they act as reasons, is left undetermined. On this account someone could say that an act ought to be described as morally cruel and it could still be left open as to

whether it was morally okay overall. This seems highly counterintuitive, however. To correct Dancy's account so as to remedy the intelligibility issue I could initially pose that a thick moral concept – not merely a property which shares the name - that is a reason for in one case can *never* be a reason against in another.

When we *give a thick moral assessment* we are explaining how we see the make-up of the ontology of the moral terrain before us as moving us to action (or assessment or belief or utterance or whatever) in the hope that someone else grasps our explanation (if we are justifying our actions). Once we have the eyes to pick out the terrain (think here of Putnam and 'seeing' relation) we have internalised a mode of analysis which moves us to action and which we assume the other person shares at some level. They may conceptualise certain cases differently, but never be completely lost to our intentions. This is what thick moral concepts are; they are both action guiding and explanatory to others. They explain to others why we acted the way we did when a moral justification is available. The terrain that we conceptualise, however, is made up not of concepts, but of properties, and many of those properties are apt to be captured by thick terms. But only in the act of conceptualising those properties in various ways to provide the world with meaning (say aesthetically, humorously or morally) do they take on the form of fully fledged concepts. And if they are moral concepts, I claim, they come armed with an inherent evaluation tied to them by way of the conceptual network we share. Again, if I say a boxing match was brutal and bad for it, it is left open as to whether I am making a moral judgment or an aesthetic one. To meaningfully respond to this, it seems to me, one would have to share not just the some definition which stands for the property 'brutal' in some minimised descriptive sense, but both the aesthetic and moral conceptualisations of the word brutal and some context regarding my views and the object I was evaluating.

Due to this inherent evaluative nature, when a property which is apt to be picked out by a thick term is present and doesn't move us to action this calls for an explanation. And this explanation will be one, by my lights, which makes use of at least one other thick moral concept which is either explicit or implicit in the context. This thick property will be what explains the non-motivational quality of the other thick properties present. It will alter the way we conceptualise the terrain in a way that is explanatorily clear to others. For instance, take the phrase all children hate to hear:

(4) Sometimes you have to be cruel to be kind

No-one who utters this phrase presumably thinks that their action ought to be *morally judged* as cruel, even though they are openly admitting that cruelty is - in some way - present within the action. In *what way* will demarcate the crucial distinction between my own account and Dancy's. Think of an innocent case which would induce this phrase being uttered. Does the cruelty act as a reason against what the person utilising this principle is doing? No. In a way it acts for it – it acts as a means towards correcting bad behaviour. But doesn't this tie in with exactly what Dancy would claim – that cruelty can be a moral reason for in some cases and a moral reason against in others? Or that sometimes cruelty is 'just what is called for'?

I think not - or at least only superficially - because in both instances Dancy must be taking the term 'cruelty' to refer to the same thing: cruelty conceptualised as a moral reason giving concept. If he is not saying this, perhaps by taking cruelty as meaning something like "purposefully causing people to suffer" is sometimes exactly what is called for', then he is taking thick concepts to merely denote descriptive definitional properties. Thick concepts then would become an idle intervening step for Dancy, as we already know that combinations of descriptive features can interact so as to come out as reasons for, against or not at all. But Dancy clearly thinks that thick judgments *do* make a difference other than a mere definitional categorisation. He says, "That an action is obscene makes a difference to how we should act (although not always the same difference) *beyond any made by the features that make the action obscene*" (2004: 17) Thus, as no descriptive features could necessitate my use of the word cruel on pain of breaking the shapelessness thesis, Dancy seems stuck with the thesis that the moral concepts like lewd and (unless he has an argument against my example) cruel can change valence. But this threatens the very generality that we need for moral explanations to remain intelligible.

In my mind, cruelty is not acting as a reason giving property at all in the 'cruel to be kind' case. My contention is that the descriptive feature, say, 'depriving a child of something they want to their immediate distress' is a specification of that which can be captured under some abstracted non-evaluative rendering of cruelty wherever we might see it: 'purposefully acting so as to cause suffering in others'. That is enough for us to recognise some minimal notion of cruelty, but what is happening here is that

we are merely seeing something which is a caricature of cruelty and labelling it as such. Kindness is giving us a reason, not-so-much to overlook the features that denote the cruelty, but to not to see them as amounting to a full version of cruelty.

These features we label as cruel speak for the reason only in this way: they are a structural means to attaining a positive moral result. We applaud the attempt at moral kindness in this case, but we do not condemn the cruelty because it is not a moral feature. Cruelty here fails to motivate at all. This works from the basic intuition that one cannot say that the same action is morally cruel and morally kind, because this is just a more specific way of saying it was simultaneously morally good and morally bad. If cruelty as a property is to be given any meaning at all in this case, it is as a practical (but non-moral) one of a means to an end (that of kindness). But it is tempered by the kindness. How much 'cruelty' is acceptable in various cases is demarcated by the complex notion of the sort of kindness a parent aims for with these sorts of actions. This would count against any idea to the effect that *the moral concept of cruelty* can be either good or bad dependent on context. Cruelty thought of as a vague set of descriptive features – features which act as a the most minimal description there can be of a property of cruelty - can be put to various ends which can only be ascertained once we know the reasons (this is Dancy's shapelessness point) whereas cruelty as an end in itself is necessarily morally bad as part of the moral conception of it. There will be no example that someone can bring forward which will ever show someone being morally cruel and morally good simultaneously. If one's mind leaps to find examples the examples will carry within them a moral value which trumps that of cruelty.

It makes sense to discuss the relation that is left between moral properties and moral concepts. My account obviously claims that even though the descriptive features which underpin moral concepts can be present in a situation those concepts might well not track the relevance of these features into our reasoning due to the way they are arranged in our conceptual network. For example, if my friend has a much better job than me and then we both apply for the same role which is better paid, and he gets the job, I would not call him selfish, however small the pay rise for him. His right to self-betterment outweighs the notion. Although in some descriptive sense it could be said he has taken away goods which were far less valuable to him than they were to me. In cases like this, the moral concept does not track the property even though

the property can be seen to be there. But the separation of property and concept isn't meant to be a full divorce. How could it be? Moral properties do instantiate moral concerns still. It is just that they can be offset by various other moral concerns and values which may be explicit (cruelty and kindness) or implicit (brutality and the moral freedom to box). Only another moral property will be able to be used to explain why a moral property that is present didn't bring one to act upon it as a morally relevant concern. Accordingly, the presence of only a single moral property will necessarily motivate morally for anyone with access to the conceptual network.

6.1.2 – The Position Applied to Non-Cognitivism

At first this may seem like it is quite the coup for the non-cognitivist. After all, we have just admitted that there is a clear regularity between utilising the right evaluative term and the descriptive core of the property which instantiates the use of that term. I would answer - yes, the term; but not the concept. Someone who simply knows that kindness looks like 'helping others', and that cruelty is 'taking something away from someone which they value', will not ever be able to know their way around the various 'cruel to be kind' scenarios and master the complexities of those situations. In fact, given the long term nature of the kindness they would likely expect that we label the act as cruel, as at the level of properties the 'cruelty' is immediately present, whereas the kindness is both delayed and expresses a collection of concerns about raising people well. Such outsiders might have access to the vague descriptive forms of the terms but they will not and cannot have a grasp on the reasons for each move, because they are denied access to the reason giving part; the conceptual network in which those properties become reasons for action. It is access to this conceptual terrain, of what links what property and the results of this - of which properties in tandem can better each other and which worsen each other - which does all the explaining. Add courage to any set of distasteful views you wish and you have a worse situation than one with merely distasteful views alone. Add generosity to weakness of will and the same goes. How could a person with no access to the conceptual network explain these sorts of conglomerations? To put it another way, at what point does meaning enter into the picture at all for the non-cognitivist?

This is crucial, because they were supposed to be answering the question of how thick moral concepts hold their meaning or their semantic content. The non-cognitivist claim was that all of the meaning was in the world-guided aspect of the concepts, and that the reason giving part could be extracted, leaving only the definition equating to the meaning. But this leaves it unclear as to how they would ever know how to judge an ethical concept when it was spotted. Most moral situations are not constituted by one singular thick property (although at least one is needed) but several, each of which stands in potential evaluative reassessment on the basis of the others which were available. Furthermore, we have just seen situations where properties which can be conceptualised morally are present, but do not form the basis of a moral judgement but rather a broadly evaluative one which is given its status by way of its context – say brutality in boxing (or in a less ethically loaded case, brutality in a documentary made to inform people of the actual brutality of some political regime). In either of these cases brutality being present can clearly be a good thing, as it achieves the aesthetic expectation of the viewer or adds to the factual informative qualities respectively. With only the definitions to work from, how could the sort of person the non-cognitivist imagines as gathering the meaning minus the reason giving aspect even tell whether this was a moral use of the term brutality or not? If it is the case that a particular term can be present within various conceptual schemes in various ways, how can the person who pays no heed to such schemes in their applications of the terms ever get it right? It seems like access to a conceptual network is required for anyone to understand what the terms really mean in the various contexts these networks demarcate.

I believe that these questions issue in problems for the non-cognitivist, although I will admit that these ideas are programmatic at best. It will need to be seen just how this concept/term-property division can be utilised in other areas. It seems that this will be a structural principle that should be universal. However, it makes sense that it demarcates the difference between noticing a property is present in each case and being able to place it in a contextual network. This would denote the difference between pointing and naming properties (describing) and actually explaining things and reasoning with concepts. All that seems to be left in the non-cognitivist account is recognition plus approbation or disapprobation. This seems to miss the point that by separating the two components of thick terms one will be left floundering around

those who get the point of the practice which those concepts govern. There is no room for conceptual mastery here. To paraphrase a wonderful example from Alice Crary, when can one say that someone has mastered the concept 'dog'? In young children it may be when all the things that are dogs come under the concept, although she may also include things from outside the category, like wolves and foxes. In older children, it may be when only the things which are actual dogs fall under that concept. In adults, though, it may include an understanding of why the phrase 'Tony Blair was George Bush's lapdog' is apt precisely because it pays reference to dogs.

This pushes us back towards the idea that there are huge pragmatic features of our linguistic capabilities. But there seems little reason why would we have thought this wasn't going to be the case. The idea of 'pure' semantics unaffected by any pragmatic considerations seems like a chimera too. Properties may be vaguely definitional, but concepts seem to come in networks with various normative interactions available between those concepts. If we are trying to abstract from pragmatics to look only at cases of meaning, I challenge anyone to go back to (1) and read this claim without feeling jarred. That jarring can stand as some phenomenological evidence of just how strange it is to think of moral concepts as being untethered by an evaluative direction. And in all the examples I can find, of where evaluative direction changes due to some pragmatic feature that is in play, in the most convincing ones it is usually because the term is no longer denoting anything like an ethical concept.

This analysis of course means that moral properties and moral concepts come apart. What underscores the moral and the aesthetic concepts of brutality are the same property, even though these aren't identical concepts. Non-cognitivists needn't worry about this analysis. To use this analysis as a stick to beat non-cognitivists with would beg the question. After-all, I have claimed that the meaning of a term is dependent upon whether that term is conceptualised in this way or that, which is precisely what they deny in their first premise. The battle with the non-cognitivist is best won at the level of epistemology and how meaning arises at all in the world that they conceive. However, the variability in concept use attack has been stymied. We have shown that for the cognitivist, moral concepts need not be thought of as changing evaluative direction, even if the properties underpinning them can be variously conceptualised

within our language. Further, we have given ourselves more reason not to claim that moral concepts are evaluative but underdetermined. This defends against an intuitive weakness whilst simultaneously giving us a better grasp on how we come to share our moral explanations. It explains how reason giving works and how we can extrapolate from loosely cited reasons to strong moral explanations with ease. All the above is at one with the idea of the impartial moral standard as well. One has to have grasped the conceptual network that constitutes a morality to understand how to get around the moral terrain. If one has different ideas about the morality of boxing than me, that argument can be made, but the disagreement will be umpired by the basic background of the conceptual network that we share.

One final point might be this. Our point in Section 5 was the idea that our coming to reach a shared standard of non-reductive moral concepts was the very thing which allowed us to conceptualise the properties relevant to those concerns as pertaining in the world. The idea that we can still recognise those properties even when they are not relevant to how we are conceiving a specific case hardly seems like a surprise then. The interest which fosters the recognition comes far back in the genealogical story, so the properties are there to see, even when they are not being tracked by our concepts as being relevant here and now. If this is so then this gives a better explanation of the non-cognitivist contention. Without the interest the property could not be perceived. These concerns would fail to be saliently grouped together. The interest secures the property, it is not added on later. Much like my use of miracle then, anyone picking out these properties without a claim to the interest is doing so parasitically and in a diminished way which cannot capture the full meaning the concepts provide to those who utilise them meaningfully within a first order morality.

6.2: Objectionable Thick Concepts

With this in place it is worth quickly looking at the second major threat which separabilists claim should trouble non-reductive cognitivists: objectionable concepts. We saw in (3) above the use of the word lewd in a description of a carnival. Now, lewd is a special concept, claim the separabilists, because it has a historically understandable evaluative direction which most people now object to. It is not that it

has changed its direction of evaluation, but rather, claim many people (we will call them objectors), evaluating people as morally bad (or good) on the basis of their sexual exploits is itself an illegitimate move. If lewd shares an extension with some phrase like 'overt display of sexual behaviour' then objectors believe that this is not apt for moral judgment. However, claim the separabilists, objectors can still use 'lewd' just fine. So the attitude which non-objectors attach to lewd becomes irrelevant to its meaning. This is not an evaluative concept which can change valence, this is an evaluative concept that seems to no longer evaluate at all.

Pekka Väyrynen talks about this at length. The assumption he makes and does not question throughout his work, which I think is key to the whole thing, is that these two sets of people, when they utter the word, are really aiming at the same sense; or to put it another way, are objectors using lewd parasitically? I think there is strong reason to believe that this is begging the question against the inseparabilist if he is a Wittgensteinian non-cognitivist. For it is part of the background theory that if one is parroting a word then one is not really saying it at all. It is empty in their mouth. Lewd stands between generations though, and while it is a rarity these days I can still easily think of places it would be apt. Anything which overtly sexualises children would be one, like some child pageants and so on. Notice here how the claim to lewdness would have built into the idea of 'overt sexual behaviour' something like the term 'inappropriate'. This stops being an acceptable term to utilise for adults in the right settings because we believe that adults have the right to express themselves sexually however they so wish. Other words do not have many modern day settings at all however, and pretty much belong to distinct societies. Perhaps a more 'other-timely' term - and one which is a further example of Väyrynen's - would be 'chaste'. In some vague sense we all know what chaste stands for: it refers to refraining from sexual activity in a positive way, and that that way has with it some connection to moral goodness as derived from holiness. It is a withholding in a way which, to use an outdated phrase, makes one closer to God. This concept will have been used in a variety of ways at a certain time (Victorian Britain in higher society one might have thought). We can still use the word today, certainly. But it seems to me to be open to argument as to whether we would be expressing the same concept as the people who uttered it back then, which is what Väyrynen claims we can do. Are we not thrust into the position of the anthropologist, much like the anthropologist who is

trying to grasp the word taboo? Like her, we could surely grasp the most obvious extensions of the concept, the ones which shared some clear similarity pattern, but would also, if we were thrust into the society for whom the word still had full extensional meaning in an active network of first order moral concepts, be left asking just why certain actions or behaviours were seen as chaste when others were not?

It seems like saying that we can still use the word as though it was a full concept just with the attitude dropped follows the reasoning, 'I have so far used the word without error, therefore I have grasped its meaning, and therefore I have mastery of the word'. If one does not have the resources to use the concept in all the ways that someone who has mastered it can, then one also does not have the resources to know this is the case, unless one were corrected by someone who does have mastery. But in that case, in a world - as the first premise of the separabilist argument goes - where everyone thinks the concept is objectionable, we have lost the last resource for telling us that we have either gone wrong (or right) in our application, or missed what would be considered to be an apt application. One can only conclude that it begs the question that those people who object to a term simply can grasp the sense the term was supposed to convey and are thus employing an empty term well.

Still, Väyrynen thinks this is a serious issue for us. And the seriousness comes from the fact that when we reject a statement of sort:

(5) Isolde is chaste

We are not denying it in a usual way. We are not saying this claim is false, rather we are saying that this statement is 'lacking a truth value'. It 'doesn't refer' in a special way. It is not wrong; it is an inapt mode of assessment in the first place, like most people believe about astrology. 'Paul is a classic Pisces' might be another example of the same sort of phenomena. This, Väyrynen maintains, is important because it shows that the rejection of (5) is not truth-conditional negation, but rather 'meta-linguistic'.

This, he argues, is enough to show that the evaluation the thick term carries is not a part of its semantic content, but a pragmatic feature of how we use the word. But this argument makes a very specific divide between pragmatics and semantics; one

which asks too much of any theorists who sides with a Wittgensteinian view of language acquisition. Words would have atomistic meanings, lexical meanings, or semantic meanings, which could be gathered as being utterly independent of the ways (and contexts) in which the word is used by the people who inhabit a form of life. As Väyrynen puts it, in his discussion of variability, there is a distinction at work between 'speaker meaning' and 'semantic meaning'. And everything that falls under speaker meaning is pragmatically explained and thus does not belong to the semantic content of the term. But Väyrynen also freely admits that no non-evaluative predicate that can be swapped in for chaste can be "...anything like extensionally equivalent to 'chaste' or sufficient for it to apply, or even that it captures the full non-evaluative meaning of 'chaste'" (2009:243).

But now we are left wondering what actually does provide the semantic content. Where is left in 'chaste' for the sense of the term to reside? When people use it evaluatively they are using their own 'speaker meaning' and thus this should be explained pragmatically – but use a non-evaluative predicate to capture it and you've failed to say something extensionally equivalent with the relevant semantic sufficiency conditions. This reminds me of the game example. Game is not an evaluative concept, but it too cannot be grasped by any other non-evaluative predicate other than game. Its semantic conditions are partly carried by being inculcated into a conceptual network. One would be required to be embedded in a form of life where 'game' was applied in variant ways as a 'cluster-concept'. But this explanation is begging for one to say, with Wittgenstein, that there must be agreement in judgments of application that invite a diagnosis of practical concerns on the part of the speaker; of the speaker developing certain propensities and sensitivities. The sensitivity is one of the necessary conditions for using game correctly. And this gives the lie to the fact that there is such a thing as the semantic meaning/speaker meaning and semantic/pragmatic divide conceptualised in the way which Väyrynen wishes, so as to say that some uses of a word are just semantically guided and others are guided by pragmatic concerns. The picture requires that there are pragmatic conditions underpinning the semantics of (at least some) descriptions in order for anyone to mean anything with those words. Now, one can ask, why is this not also true of evaluative judgments? The appeal to pragmatics is gone. What one is concerned about is that the pragmatic underpinnings of the proper semantics

of the concept might include understanding it as having an evaluative direction. The intuition seems to be that evaluative concepts are open to be assessed on a first personal basis and thus pay homage to the attitudinative make-up of the speaker. But this simply ignores the Wiggins and Thomas argument that subjectivity can be conceived as entering into the equation as a pre-supposed set of first-nature concerns; concerns that are necessary for one to grasp the semantic output of the terms generated, and which are one part of the sufficiency conditions for correct application of the concept in the moral domain.

I conclude that moral cognitivism can weather the storm here, and might even provide a mode of analysis which shows not just where the confusion comes from – in mistaking terms for full blown concepts constitutive of a particular domain – but can also from this analysis come to better define itself by ruling out the extreme sort of holism one finds in Dancy and show simultaneously what was right about that analysis and what was wrong about it.

Section 7: Application to Medical Ethics:

In the previous sections I have spent much time in developing a way of considering ethical reasoning, which eschewed the idea of relying upon general principles in such a way as seeing them as foundational to serious ethical thought. In so doing, via a detour through some more broadly epistemological concerns, I also made out a case that ethical reasoning that did not couch itself in such generalist terms - relying instead upon discernment of concept application, which I tried to show was a practical art rather than a theoretical exercise – could still be seen as reasoning towards knowledge. By making the case that all types of knowledge relied heavily upon a practical mastery of normative constructs of a particular practice or form of life, of both the content and system based varieties, I allowed for a different form of moral reasoning which I labelled as Particularist.

This particularist position made various claims, which were quite at odds with the generalist picture and seemed to undermine many of the assumptions about moral theorising which is utilised in the application of morality to practical domains, such as medicine. The theses that emerged as a by-product of overturning the generalist picture to the particularist one, whilst also defending that picture from non-cognitivism, were manifold. Before we start to look at how the acceptance of this picture would change the face of medical ethics, it may be wise to collect these theses together as a point of reference for this section, the purpose of which is to discredit some of the major underlying assumptions that bleed over from that generalist picture into the actual field of medical ethics and best clinical practice.

Firstly, we have the non-definability of ethical and value laden, or thick, concepts in language that is understandable from outside of morality, by way of being grounded in naturalist or rationalist terms.

Secondly, we have the idea of rational moral disagreement. By this we meant that two people might disagree about what the best course of action is without either being able to say that the other's assessment can be traced back to a move that undermines a rational reasoning process. Such a position was bolstered by the ideas of incommensurability, bedrock, and the idea of how potentially various and varied

normative commitments are necessary for discourse in the all areas of thinking, and how they might be applied and interpreted differently by various agents and in various contexts. What was crucial here was a pre-suppositional account of subjectivity, whereby the sense of moral terms was given by way of inculcation, so we might disagree about specific instances and still be singing, as it were, from the same hymn book.

Thirdly, there is an important point about proximity. There is no description of a situation in complete terms that will allow for a full and proper assessment of the ethical parameters of that situation by someone who was not involved within it. The uncodifiability thesis, pointed to in the first of these listed theses, combined with idea that perception is a conceptual exercise, means that when describing a situation, neither evaluative nor descriptive language will allow for someone who is removed from the situation to garner a *full* understanding of it. Either, one aims for a value neutral description - something I would claim is impossible to achieve, as any minor detail can radically alter how one would evaluate the whole - or one aims for a description which includes evaluative language, which would of course colour the description in such a way as to disallow for an appropriate evaluation which differs from the description given. This point raises a serious issue in regards to the relationship between theory and practice. Discussing moral theories as though one could find the one that best represents moral reasoning as it ought to be seen, and then assume one can arm people with this theory and expect uniform answers in specific situations, underestimates the gap that lies between any such theory and practical application; a gap that is necessitated by the fact of interpretation (beyond 'measuring rod' applications) and the overlapping and multi-faceted list of concerns available if we accept language as rising up holistically, rather than being handed down legislatively. Both what counts as the best conception of maximised happiness, or what counts as a law one would will everyone to live by, requires one to look inwards towards one's own conscience, as it were – not just in terms of acceptance, but also in terms of application.

Fourthly, we made an important point about motivations, which showed that shared emotional reactions were inextricably linked to all evaluations but garnered their force at the presupposed level of interest. Rather than being an extraneous factor which hindered rational choices, emotional reactions are necessary for genuine,

rational, moral evaluations. This point is in accordance with the second idea, above, regarding the non-definability of terms, such that one cannot predefine what rationality looks like ('cold and calculating') and then reject modes of thinking that stand outside this given definition. The flexibility of the category of rational thought cannot, without circular reasoning, be predefined by an assessment of those types of reasoning that we mostly, by consensus, agree to be rational, and then utilised against those which do not fit this form. What counts as rational is not handed down legislatively, as though it has some platonic form. To believe it should is itself a normative position. Opposed to this, I have argued that what counts as rational should rise up through the practices themselves and be argued for, or against, on the normative basis that necessarily underpins their usage.

Fifthly, following the work of Alice Crary, I think we can say that the idea of soundbite examples in moral philosophy, of thought experiments which contain minute amounts of information and are supposed to disclose to us intuitions we might be susceptible to from which are meant to draw very large and significant conclusions, are damaging to understanding morality. Crary's quite wonderful analysis of how literature can impart to us a full understanding of ways to live, and how many variant types of behaviours can be classified under the extension of a single concept like pride show just how important full pictures of background and character are for a full and valid moral assessment. We may have disagreed that certain mundane actions were moral actions from a first personal perspective, but I do not disagree that from a third personal perspective or a reflective perspective we can find ourselves or someone else to be severely morally deficient without them having undertaken a classically moral task. In all the talk of what we must and must not do, it seems we can lose sight of the importance of *how* we go about doing the things we already do.

Sixth, finally, and as of yet not fully mentioned, we have to recognise an idea that is starting to emerge from this analysis. What our non-reductive account of morality points towards is precisely what is needed for us to be able to morally improve. The discussions of 'objectionable' concepts, when combined with the idea that recognition of a thick property needn't motivate us unless it our conceptual network of explanation places it in a reason giving position, can go to show this. Our explanation of why we no longer use chaste is that there was no place for it left in our conceptual network once we had come to realise that the traditions which gave it

credence were a useful mode of analysis. The shedding of those religious traditions was inspired by the enlightenment, and one consequence of this was the idea of personal autonomy and freedoms were free to take a strong deliberative position in our evaluative conceptual network. This is the importance of seeing morals as strong reasons, and not silencing ones. Chastity was once a virtue and would thus have silenced other reasons. But morality conceived as a self-appraising conceptual network with various inputs (interests, cultures, traditions, personal outlooks etc.) must be vigilant against becoming too fixed. An independent conceptual network, such as that we are conceiving morality as being, must have some strong internal consistency. Traditions, whether they be beliefs in gods, or accepted hierarchies or what have you do input certain values, and these values have their positives in that they give us a shared sense of meaning. But they are also open to criticism when it becomes clear that they clash with other fundamental values in an unsustainable manner. In the same way, one cannot both chastise on the basis of someone's being unchaste and maintain that what consenting adults do with their bodies should be underscored by the value of personal freedom without internal inconsistency. A more recent version of this might well be the '#MeToo' movement, which (in a quite innocuous but simple example) have questioned such traditions as 'grid girls' in motor-racing competitions. If we are to claim that women are to be valued for who they are and not their aesthetic and sexual qualities for men can we seriously justify such traditions which give the opposite idea and may influence the self-perception of young women. Morality, then, conceived in this manner, is a self-correcting conceptual network. This explains the possibility of moral improvement and betterment in a way that is far more edifying than can be made out by some fixed set of principles, or by a position that equates being moral with being utterly aligned with the most commonly held moral outlooks within that society. And it is this that I believe can be most effective for Medical Ethics. Let us find out what values are at stake within the various traditions of medicine and see which values and which traditions are the ones which we simply cannot do without.

It is with these five major points in mind that I now wish to turn towards the field of medical ethics. There are three major ideas I want to turn this against. The first is the general notion of expertise in ethical concerns. This follows fairly simply from what has gone before but I wish to clarify a few things and address a relatively new move

by the ethicist to match the more particularist outlook of expertise. As we will see, lip service is paid, but in her actions she wishes to hold the same position as before.

I will then move on to the idea of definitional theories; promoted by theorists who still look for a generalist underpinning to their claims. I will consider one case of this, in the form of a dispute between John Harris and Anne MaClean.

I will then address a newer form of discussion in the medical ethics literature, which includes not just that judgements be passed on particular actions or processes. These theories try to give a characterisation of what the practice of health care is all about. As we will see there, any characterisation will fall into the trap of being platitudinous, or otherwise make distinctions which rely on defining everyday terms in restrictive ways to promote a particular characterisation at the cost of a basic understanding of morality or medicine. This is something I believe to be another form of the touchstone methodology. I will critique it thus.

In order to not deflate those who still wish to say something meaningful about clinical practice and ethics, I will then move on to what I see as an example for future discussions in the field.; a paper which makes serious moral points without giving too much away to explanatory moral theorising.

7.1: On Expertise in Medical Ethics.

One of the most obvious conclusions that follows from what I have so far said in the previous sections, is that there is no independently ratified moral standard that would allow for someone to say that they have more authority in making a specific moral decision than anyone else, assuming that all the natural facts were known about the case by both parties in a dispute about what to do. Given the paucity of generalist theories in capturing relevant features in such a way as to determine how one should go on, the seeming necessity of virtue terms in helping guide our way through the moral terrain, and the uncodifiability thesis, there simply seems to be no way that we account for our making judgements about what is morally required of us by way of defining morality against touchstone definitions for verdictive terms, whereby theory creation can be uncontroversial, such that someone who spent more time attending to a theory (in the strict sense) could be said to have greater

knowledge of how to go on. All the way down to bedrock, moral definitions require substantive moves that, such as ordering of certain ends, virtues and values; rather than reveal the very nature of morality itself, can only go to reveal the implicit assumptions (be they moral, ontological, epistemological, methodological, religious or of any other overlapping category of thought) of the theorist who is utilising them. What we are each aiming for is to get things right. And despite difference (Variance in Judgements) we can still be content that we are talking about the same things, even if we see those values as applying differently from someone else here and now.

For all its reliance on norms for one to know how to get around within an empirical practice such as science, however, the practice still puts out a body of knowledge at the general level that certain things are true (upon this conception of science) which makes predictions that either will or will not be falsified. In so doing, we can eventually decide that we have something like a law, or theory which we can say is with some certainty correct. We are not forced to accept a result as absolute proof that a certain proposition can never be questioned. Nor would the laws we state with confidence, due to the results of various experiments, mean anything to someone who hadn't been inculcated to see why the norms that govern scientific enquiry are apt to do so. But for those within the practice they have a strong epistemic appeal.

This is what Wittgenstein means when he says:

It is **as if** we had hardened the empirical proposition into a rule. And now we have, not an hypothesis that gets tested by experience, but a paradigm with which experience is compared and judged... a new kind of judgement
(Wittgenstein, 1978: VI-22)(my emphasis)

If we buy into the 'as if' that I have emphasised here, as it seems to be meant, we need not worry about the idea of the possible interpretability of rules as they are meant to be seen by Wittgenstein's lights. In fact, the 'as if' is meant to direct us between the idea of a hardened proposition and a rule in the Platonic sense. The 'rule' is thus not 'stretching out like rails', but is a propositional end point of study on a *particular conception* that may or may not be discounted at a later date, based on further data gathering, furthering alterations and tweaks to our conceptions or metaphors which we see as forming the best possible explanation of the natural

world. It can be used to bring someone who shares the necessary forms of life against which the practice is meaningful up to speed. While morality mostly works regarding knowing how then, the practical side of things, science, by way of its empirical nature, allows for a highly specific knowledge that. But without being involved or acquainted with the human activity of scientific enterprise these propositions would be meaningless.

'Is' propositions differ from 'ought' propositions in almost exactly the way that the consensus has it, then: one is open to proof, whereas the other is not. What is crucial about this analysis is that it makes the case that proof is not the whole of story for epistemological success. The proof only convinces those *already* accepting that these are the norms that should underpin our judgements about the natural world or this aspect of it (rather than some religion or other mysticism, say) and that these norms should be interpreted this way rather than that.

Herein lays the difference between empirical propositions and moral ones; the difference which allows for expertise in the former and not the latter. An ethical expert cannot be like a scientific expert, a structural expert, or a medical expert. Expertise in these latter senses is always regarding empirical propositions that are testable. We can think back here to Judith Jarvis Thomson's idea of 'good'. Rather than being a type of thing in itself, yet to be defined ('The Good'), it is in most cases used as a contextual evaluation that 'X is good *for* Y' or to show our approval of some ultimate or final end (some things are good without being good in relation to something else). If a medical, structural or scientific expert says that a particular course of action is good, they mean good for keeping a person healthy, making sure a building will not fall down on us, or any other scientific ends we might think of. Before testing, a theory is good because it coheres with the best knowledge we already have and have tested. When we say 'X is good', morally speaking, as we saw, we had no real way to naturalise that expression in an agreed upon set of norms that admit of success or failure, and in doing so become 'hardened' *as if* they were 'rules'. We might have a measuring rod example, say, slavery, which we could then utilise if someone started to make a suggestion which seemed to make some people other people's property in order to overcome some economical issue, say. But even then, there's room to say it's not precisely the same. In the same way if

someone says 'taxation is theft at the threat of imprisonment' we can make the case that there are suitable differences - for all the alarming similarities.

With this in mind, there can be no expertise within ethics in the same sense as there is with scientific expertise. No guidance which does not have built into it, however implicitly, an already established substantive position. And as it is substantive 'all the way down' so to speak, there is also no measure on how one can interpret normative, substantive positions within it as interacting with each other. By this, I mean that, even if we see somebody espouse one particular view for one particular reason, this does not mean that we can pull them up on a contradiction for then assessing a seemingly similar case in, what seems to us, a seemingly distinct way. To believe it does is to buy into the idea that theories come as sets with touchstones that are interpreted in this way rather than that, or ought to be expressed in this way rather than that. One might express a right for a foetus to live in one case, and then cite the negative consequences of not aborting in the next. To believe this is necessarily a contradiction assumes that these ideas are the output of incompatible theories, rather than handy ways of expressing two concerns that could be linked by, say, a specific ideal of compassion to the particularities in each case. I believe that this follows from our discussion of Dancy in 6. Generality occurs at the level of thick concepts, not specific actions which those thick concepts guide. And thick concepts can be spread rather wide and interact in interesting ways.

Any sort of input by the ethicist does not show someone who has some greater insight. It shows a person who has a specific insight, one that is, itself, substantive all the way down. The logical possibility that norms can be arranged in almost any way, be they the sort of norm that act as content or the sort of norm that suggests methodological presumptions, means that there is no map the philosopher - whose licence for authority is in the workings of reasoning itself (logical connections and so on) - has which the non-philosopher does not and thus the philosopher has no special authority here.

7.1.1 The New Claim to Expertise:

This undermines also what I see as a lesser claim towards expertise than that a strict generalist might claim for herself. The strict generalist believes there are truths that can be discovered by explaining the naturalism of moral language, as self-evident if *suis generis* truths, or by *a priori* reasoning. Enough study of these, will, by her lights, arm her with a solid conception of the nature of morality – a conception from which how to go on follows - which would allow for full expertise. We have, I believe, fully renounced this position; although it is fair to say that this was the dream for Medical Ethics when it began.

The lesser claim, however, is often a claim that comes along with the explicit caveat that it *does not* imply expertise in such a way that one ought to defer to their decision at all times, as a layman would to someone with much more experience in a particular empirical matter. It is a claim that is nicely summed up in the following metaphor by Zoloth-Dorfman and Rubin:

“...we suggest, it is the ethicist who is the navigator; knowing the map, being familiar with the terrain and it’s complexities, calling attention to how the desired...course might be changed by the immediacy, temporality and particularity of a given case, and above all guiding but not controlling the course.” Zoloth-Dorfman, (1997)

Again: this account directly means to deny that the medical ethicist is an expert of the type that can give a definitive answer. But the same assumptions are involved with the metaphor of navigation as they would be if the ethicist was said to be the captain. To believe that because one makes a certain move one is therefore committed to various other moves, such that someone who is more in tune with ethics - namely the ‘ethicist’ with the her ‘unique perspective and tools of her profession’ might spot, so as to navigate you away or towards it - speaks of a conception of moral theory as if made out of various possible wholes (or following the metaphor, ‘courses’). If the doctor makes a call, what precisely can the ethicist add to this? Any suggestion of a particularity she may have missed will make the assumption that this particularity is salient. Perhaps the doctor’s not seeing it as such is actually because it is not. To say otherwise would imply general propositions can be made out which demand recognition. If the ethicist claims that she does this, she

seems committed to doing other things. But this too goes against the uncodifiability theory and the rejection of touchstone theories. One starts to wonder whether perhaps a better grasp of the actual ethico-legal process (something which does define rules and is as such rule-based in such a way as to allow expertise) might be what they are referring to. With this I hold no issue in attributing expertise, but it is not ethical expertise so much as administrative. In the drawing up of any code, there will of course be massive problems, if what one is aiming for is to draw a code which truly does pay homage to the idea that there can be right and wrong moral decisions.

The mention of 'particularity' in the above quote only further confuses the issue. For if particulars are seen as an exercise in 'seeing', or perception - where that means the ability to discern minute differences in a way that cannot be placed under a generalised notion - then what the ethicist is claiming over the person she navigates for, is a sort of practical wisdom in concept use and having adopted the right sorts of concerns and learned to apply them and appeal to them in the right sorts of ways. This thesis maintains that that is precisely what it takes to get things right in ethical situations. In which case, mentioning her 'tools' and 'training' seems to imply an assumed method to train people towards practical wisdom. If one were to have such a thing as virtue and practical wisdom in order to open up an honest dialect in such a way as to heighten the perceptions in others of such things, as I have no doubt wise people do, then one can certainly claim to be an expert in the sense that they might be deferred to in hard cases. After all, this is expertise that does not speak of a body of knowledge; and this sort of wisdom is almost presupposed by virtue ethics under various guises.

How one would truly be able to show they have such abilities though, what 'tools' would be necessary for it and what 'training' one might undergo so as to acquire them, seems to be beyond our grasp to say. Wisdom is yet another thick term, and the norms of success and failure that are available for empirical experts are not available for analysis of such terms, especially when hard cases appear. One cannot train to be wise in any strict manner. Wisdom, when contrasted with intelligence, is precisely beyond the scope of working out connections that are explicitly laid out. Rather, it refers to the sorts of understanding that the uncodifiability theory and universalizability-generality distinction tries to point us towards: a practical mastery of understanding when certain concepts apply and the discernment and understanding

of the salient concerns embedded with shared practices that goes along with that skill. This understanding is tied up with the Wittgensteinian account that our concepts gain their meaning in the range of cases in which they are already used seamlessly, their meaning acquired by way of its centrality to a form of life.

One struggles to see what ‘tools’ and ‘training’ could really be being used here. These two phrases seem to be more suited to the previous conception of expert: an expert as it applies to a practice with a given suitable end that is amenable to hardened propositions. If, on the other hand, ‘particulars’ simply refers to how complex cases can be seen as altering the way that a case can be subsumed under some general rule, then what we have is the wrong sort of expertise, as the rule one chooses will reflect one’s own substantive idea of that case – making one merely an expert in one’s own view of the case. If the ethicist changes the mind of the professional making the judgement, it need not be because they have greater insight. They just have a different substantive conception of what doing the right thing means in this case. They may be able to give a more comprehensive sounding defence of their position, but this does not necessarily mean it is a better position over all. Here they might complain we are mixing up two categories of thinking:

The clinical ethics consultant’s expertise is an expertise not in morals, but in ethics. In other words, hers is a discipline that functions not by offering declarative normative judgments, but rather by raising critical questions and focusing conversation and deliberation (ibid)

This has the interesting effect of seeming to make the whole point of ethical discussion to be not about the outcome of a decision, but the process of making a decision. If it is not the results that matter, as this implies it is not, then it seems to be the actual act of raising critical questions. Critical questions are only of use if they push us towards something like the right answer, however. Being asked to justify the patently obvious is of no help to anybody. Indeed, with tough situations also, reflecting too much can result in both inaction and psychological tension. Furthermore, if the right answer isn’t something that the ethicist – as opposed to anyone else – can be shown to know, it seems academic that it be an ethicist who asks them. Again, I wish to make it clear - that there are people who can ask just the right question at just the right moment in order to help one perceive some salient

feature they have missed is not something I dispute. That these people can be professionalised and hired to do so in all situations that job entails very much is. The intuition is that to prove expertise one requires a body of factual or propositional knowledge and yet this conception of knowing lines up poorly with that which we have said characterises ethical knowledge: know how.

On this basis I rule out the idea that expertise can be something which just anyone can see as existing in another agent without having seen them come to varied decisions, live their lives, deal with unexpected situations and reason in certain directions in certain ways that are impossible to generalise but attain our respect. These people are the ones we call wise, perceptive, measured and assured. They deal with situations calmly and with aplomb. They show discernment. One wonders then how these skills could be discerned in a person or trained into a person who wanted to become a 'medical ethicist' in the requisite sense. The lesser appeal seems to concede ground to the particularist idea only by an empty suggestion. In practice, however, the same idea that (somehow) the ethicist has a greater propositional knowledge seems to be in play.

I will follow up this brief discussion by examining how these same generalist traits are still prevalent in more mainstream generalist Medical Ethics; the sort that considers big moral issues that deal with life and death, animal rights, moral status and so on. I first wish to look at an attempt in Medical Ethics which tries to resurrect a metaphysical scheme for the particular tasks thrown up by the de facto condition of limited resources as well as various issues surrounding abortion, unwanted pregnancy and euthanasia; valuing lives. This theory is definitional in that it proceeds by way of taking something that is a moral term and trying to illicitly find a way of making it appear, by way of definitional analysis, a technical one, which when applied to cases ushers in a particular moral standpoint.

7.2 Harris vs MacClean: Metaphysics vs. Inculcation

Harris (1984), our philosopher who utilises the method of definition – a method I aim to show as resting on a fundamental flaw within normative ethics – bases his account around the notion of which lives we might find valuable. There is

no doubting the idea that his position is made out in order to alleviate some concerns we might have about treating some beings with moral respect, while not doing so with others. This giving of respect, he says, can be seen as being given due to something the beings or objects that we ought to treat well possess, which others do not: moral status. Having moral status, he aims to show, implies that a being has a valuable life as such; not merely because it is valued by someone but because it is valuable in and of itself.

The question thus hinges on him giving us reasons as to why some beings ought to be seen as having moral status, while others simply do not. Harris believes that he can, and furthermore, that such a theoretical position is needed. Those beings which he will show to have moral status, such that their lives are valuable, will be termed *persons*. And as he puts it, we all 'need' some 'justification'. It ought to be the case that 'people are not people because they are accepted, but rather they are accepted because they are people' (Harris, 1985: 14). One can certainly see the purpose of this position, especially as it relates to the idea that if we simply rely upon our practices as they are - if we only look - we may not make progress or see similarities quick enough to stop tragedies such as those of slavery or the holocaust - where persons were grouped as not worthy of moral status precisely because the communities in which those injustices happened took them as slaves, as undesirables, and as such overlooked features which we now see as making them worthy of being moral equals.

And indeed, this is precisely the direction of Harris's argument. In deciding what does have moral status he takes something 'we all agree with' – the idea that a fully functioning adult human being has moral status and is thus a person.

If we look at creatures we are sure are persons if any are –normal human beings – and can find features of their lives or capacities which, unlike differences of the 'featherless biped' type, incline us to judge their lives of more significance and value than lives which lack such features, we might come closer to the concept of the person. (ibid)

His analysis follows that of Locke - who he recommends to us in trying to distinguish why persons are morally special - in that it takes as morally pertinent those features which are special to the group that we almost agree about – features such as:

consciousness, language, reason and the ability to 'consider [oneself] the same thinking thing, in different times and places'. (ibid p.15)

So, we have isolated a feature, self-consciousness, that seems to belong to those sorts of beings who we all agree have this thing, *moral status*. But this does not really take us any closer to an account of what makes a person, or a being with moral status, unless we also decide that it is the only feature which is relevant in that equation. This is one of MaClean's reservations regarding Harris's account. He appears to move from the claim that self-consciousness is valuable in all beings that have it, enough for us to afford them the status of persons, to the claim that *only* those beings with self-consciousness count as persons. As such, MaClean claims, he has begged the question about all those beings of which there is dispute. There is contention between the two theorists as to whether this begging the question pertains to Harris's argument. Harris claims it does not (Harris, 1995: 221). His reason for saying this is that there is an argument placed between the assumption (his words) that full grown adults are persons and his conclusion that foetuses and babies are not. (It is fair to say that both Harris's account and MaClean's concerns about it are based largely around the moral status of certain beings, but mostly around children and foetuses. While these arguments pertain to all beings that lack self-consciousness, I will refer to foetuses and babies).

Harris does indeed place an argument between the initial assumption and the conclusion but I cannot see it as a sound one. As he puts it in his reply to MaClean,

I have treated normal adults as a paradigm... [and] *produced an account which justifies their paradigm status*. If for example the only feature of alleged moral relevance that distinguished humans and chickens had been that the former are featherless bipeds, *the moral relevance of feathers would have required demonstration or the paradigm extended to embrace chickens*" (1995: 221 leading onto note 6)

Whether or not his argument meets the criteria that Harris sets himself, that the feature of moral relevance (self-consciousness) is demonstrated as being relevant, remains to be seen. I would claim that it does not. Harris has two points to make on it. The first, as we have seen, is that we do seem to hold it self-consciousness in moral regard. This is merely a consensus point about our actual reactions to things.

It acts as no proof, especially on Harris's account, which is precisely meant to lead us to conclusions that are true regardless of our sentiments upon such issues ('we value them because they are people'). Here, however, all we learn is that there is general (although not necessary) agreement in human moral life that fully grown human beings have moral status – one I can agree with - coupled with a less obvious one: they have it because they are self-conscious, and only because they are self-conscious. What Harris requires, then, is an argument which equates moral status (being valuable in oneself) to having self-consciousness.

So what is the argument which settles that self-consciousness itself is the marker of a life having value? It comes from an assessment that is developed in 'Strategy 2' (1985: 15) - which is meant to move us beyond the simple observation that we value beings that have self-consciousness and in many cases this seems to be a good reason as to why we do so value them - that we cannot say that any particular sort of self-conscious life is more valuable than the next. Here, there is an interesting move from trying to see what makes for a valuable life (something Harris claims is a relative matter of the values of the individual being asked – no ends we give for life could ever be said to be definitively better or absolutely appropriate for all conscious beings) to what makes a living being valuable (not relative to anyone, but absolutely set). Here we see the intuition that crops up from his early assumptions:

- 1) If there is to be a reason as to why we treat certain beings as absolutely valuable (in that they have a special moral status), this implies there is such a thing as an absolute value - that stands outside of our particular or relative evaluations.
- 2) All of our valuing is otherwise relative ('X is valuable plus adjunct')
- 3) Without an absolute sense of value, our evaluations of what is valuable in every sense would be relative, including which sorts of lives. X's life would only be valuable if it is valued by X or others.
- 4) As there is disagreement on what sorts of beings have value there would be no way of determining what the right thing to do is with each life, and each life would only be as valuable as it is relatively valued (as lettuces are)
- 5) Without a theory which gives us an example of absolute value therefore, no life is valuable except in relation to the preferences of that group or individual at that time, with each reason being as relative to that position as the next

- 6) Therefore, we need a theory of what makes something absolutely valuable which makes no reference to ways of life.

Here we can see a similarity to Wright's concerns **4.3**, to the extent that we need a way to ratify independently of our own actual judgements that something is valuable as such. With the relative/absolute distinction, Harris seems committed to the idea of trying to find a way beyond our actual judgements. Otherwise he would assumedly agree with Wright that the whole thing is matter of cultural artifice – but unlike Wright, take that to be the end of meaningful normative declarations on value.

One can see also, perhaps, the similarity to the Kantian construction of the Categorical Imperative that was shown by Schopenhauer to be an 'ingenious trick'. In lieu of being able to give an account of what makes any particular lives themselves valuable, but in need of an absolute universal of what makes a life valuable, Harris is left only with the *capacity of valuing* itself (thus shorn of any pre-decided notions of what is relatively valuable) in order to fill in the absoluteness of value he needs. This mirrors Kant's attempt to garner content for moral law from only the *form of the law*. Both stripped away our anthropological senses of good, and are left the concept under analysis to carry the whole weight of the account. In Kant's case it is the law. In Harris's it is value.

Thus, what is valuable absolutely is the *capacity to value*, as every other thing that is valued (albeit relatively) is dependent upon the capacity to value. Everything else you value makes your life more or less valuable to you, relatively speaking. What makes life valuable as such is that you do value (something, it doesn't matter what). If we accept this argument then we can say that the only beings that have absolute value are those who can value things. Very early on, Harris makes the point that 'valuing is a conscious process...to [both] know what we value and to be conscious of our attitude towards it.' (1985: 15)

But the reminder of Kant should make us wary here, as should the notion of absolute value - of something that is beyond our scope of assessment. We have already seen how appeals to metaphysics seem to beg the question when answering moral issues. Whichever metaphysical position one takes, one runs into the danger of begging the question. The attempted definition of 'person' as, 'possessor of the

capacity of self-consciousness' does precisely this. The magic trick happens in the disguising of a *moral* term as a *technical* one. If we feel that we ought to agree with Harris's account, but that we do not agree with the conclusions – that babies might be killed without any moral residue if they are not relatively valuable to an already functioning person – it might be that one has made the mistake of not seeing the evaluative elements that run through the supposedly technical definitions of Harris's account.

The majority of the evaluations that permeate that account are not to be found in the parts that he would claim are the moral equation. Rather, they are hidden away in the various assumptions within the three parts of his account: Firstly the selection and analysis of a paradigm feature; secondly, the account of valuing relatively and absolutely; and thirdly the assumptions in the idea of requiring a theory to distinguish between the two types of valuing.

To accept Harris's theory, we need to accept the three steps. I will now show that the whole process lies on a fallacy about theory and moral judgements.

The first pillar is the one that states that we just do value some beings as persons. These beings therefore have something that we see as valuable, if our evaluation is to be rational. That value can be attributed to a characteristic that they have, namely, that they are self-conscious. While we can agree with Harris that we do find fully functioning humans as having valuable lives, we need not assume that it is merely because they have the proposed characteristic - self-consciousness - in such a way for this to be the only reason we find them valuable. It may be a *reason* to say they have value, but it need not be the only one.

Furthermore, the value that attaches need not be because of some characteristic that they *have*, or even many characteristics that they have; it may not be, that is, a natural characteristic of theirs. It may be in how they are related to us in various roles that they perform, in the sorts of lives that they live, or in what they represent.

Indeed, there are various ways to see that looking at the natural features of what makes grown human beings, taken as a species of object that are being considered as apt for moral concern misses how we do treat 'human being' as a thick term, one which carries with it certain assumptions. For one, as Cora Diamond has pointed out (Diamond: 1978), once we detract from the self-consciousness, we still seem to give

a special status to things that are human. We do not eat human meat that is available even without our harming in order to get it, we give human beings funerals, we respect their wishes post death, and so on.

Harris would say that these final points are all irrelevant. We needn't do any of these things, they are mere contingencies, or traditions, that we merely happen to find valuable (relatively so on his account). But this is where our attack on the third point comes in to play. Harris would say that we cannot defend these ideas *rationally* because we have no theory which tells us why these are the sorts of things that we ought to value. Our observations do not add anything to our arguments, for we lack any explanation that the observations of how we do act have anything to say about how we ought to act. Without a theory we cannot be sure we are acting rationally.

This is the point MaClean takes up. She makes the point that 'Harris...argues backwards' (MaClean, 1993: 21) in creating a theory. In order for a theory to have greater explanatory power than 'what [bioethicists] like to call our intuitions' (ibid) – what Harris would call 'a relative value' – that theory would need to precede our particular judgements. The truth there would need to be antecedent of all human judgements. Otherwise we might simply all be mistaken. But his theory does not do this. The first premise makes this clear. Harris starts with something we all *do* find valuable (an intuition common to our community) and then makes the further claim that this is due not to the other reasons that people within that community would cite when giving a description of the value of their lives, but because they have this shared characteristic: self-consciousness. This move from 'we value them and they all share this characteristic', to 'we value them because of this characteristic' is unwarranted. The further move 'we can only absolutely value lives with that characteristic' relies on the same unwarranted inference. And getting there depends on the second point: that there is a distinction that can be made out theoretically between what is valuable as such and what is valuable only relatively.

MaClean seems to believe Harris fallaciously but accidentally conflates the two types of value, relative and absolute. I believe that this move is certainly supposed to be a move to the metaphysical. Harris is looking not for what is valuable to some group at some time, but what is valuable from *any* position (Harris, 1985: 16) and such a position requires that there be at least one fixed *absolute* value. MaClean's refusal to

accept this is likely because of her Wittgensteinian epistemology, one I share, which counts out the move to the metaphysical. But in not taking seriously the intuition that guides Harris to look for a metaphysical position, she fails to connect with the ordinary reader why 'relative value' – or values that do not necessarily enforce consensus from the point of view of a shared rationality - being all there is to discuss does not necessarily make the whole enterprise of discussing morality a merely culturally-relativistic phenomenon.

Harris, in his reply to her, does not mention her attribution to him of this contentious conflation of the two sorts of value, except in passing, in the quote above about his methods of demonstrating why self-consciousness is indeed a morally relevant feature. It is clear from his reply that he has not left behind the idea at all. In fact, he builds a defence of his position from it (Harris, 1995: 220-222) by attacking the 'positive account' that MaClean has; claiming that from her position we cannot say anything meaningful about what we ought to do that is not tied up in what we already happen to do, and can thus make no progress, or give solid (to be read unassailable) reasons as to why people should change their minds.

This he believes will be due to the fact she eschews any theoretical position. But her so eschewing it is precisely because any theoretical position she takes, much like Harris's, would itself have so many evaluative assumptions attached to it that it would hold no greater explanatory power than any which does not appeal to a theoretical position but to the values we do have.

Here are some excerpts of Harris's reply to MaClean which show that he still firmly believes that a theory is needed and has thus without argument simply rejected her most important point:

...her book is not a thesis on bioethics but an antithesis (1995: 220)

Maclean has two main objections to my theory of personhood. The first is that it is false and the second that it is a theory.... such an objection from a philosopher might be described (to borrow an analogy from pugilism) as leading with your chin (ibid)

[MaClean states] 'no-one who sets aside all his particular judgment concerning life and death can have anything at all to say about what makes for a valuable

life'... This is true; what is false is that I (or any other ethicist, so far as I am aware) claimed the contrary (1995: 221)

I am puzzled by Maclean's complaint that: 'bioethicists present the theory they put forward as showing that their judgements and theirs alone are correct'.

Anyone who puts forward a theory puts it forward as correct. (ibid)

That Harris remains puzzled means that he fails to grasp that MacClean's point, that the sort of explanatory theory he posits becomes an idle step which ought to give us no extra sense of security about our judgements. He still seems to believe that his giving a theory at all is something that an ethical philosopher ought to do to give us clearer insight into the problems at hand. The desire of Harris, to make a theory which explains our judgements as being either rational or irrational, is made manifest in his reply to MacClean's anti-theoretical position. She states that we can defend the value of children by simply saying that they are the sorts of beings whose lives are valuable and listing why we do find them valuable. She puts it like this. 'That they are babies is all the reason in the world [to think of them as valuable]'. This, she says, is an attitude that we know – from being involved in a full moral life – is apt for babies. She here borrows Wittgenstein's aphorism, "My attitude towards him is an attitude to a soul; I am not of the opinion that he has a soul" (Wittgenstein, 1953: II iv 22).

In reply, Harris gives an argument by making these comparisons, followed by a statement:

My attitude towards him is an attitude to a slave

My attitude towards her is an attitude to a woman

All of the classes of objects of these attitudes are vulnerable to the prejudices of the attitude-holder, and all, if Maclean's suggestions are to be accepted, are not only beyond the reach of the bioethicist, but beyond any and all rational criticism. (1995: 222)

Harris, then, takes her suggestions as being *beyond all rational criticism* because she fails to give a reason as to why we ought to hold them that makes reference to a theory of understanding morality which somehow bypasses the actual intuitions of morality that we have. And there is a great precedent in medical ethics literature of

the need to overtly theorise about things in order to move us past intuitions and into the supposedly rational. We have already shown that this cannot be done. Any supposedly metaphysical truth about value from which we can deduce particular truths will actually stand as a moral sentiment that has been taken from the values we all have and share, and disguised as somehow being more sacred than all the others; more pertinent; more central. But this is not because it really reveals some great truth of value itself; it is because we do indeed happen to value this thing (in this case self-consciousness) very highly. Indeed, we do value beings which have the ability to value in a special way. That we need to subordinate all other values beneath that, or claim that it is the only thing we can value absolutely, takes the philosophical trick we have just exposed, and leads to the touchstone type of theorising with which we are well accustomed: tell me if it has feature X and I will tell you if it is valuable.

At this point it is worth noting that just because there is no metaphysical position to say that X or Y are absolutely valuable, there is absolutely nothing wrong with saying that the wholesale killing of babies is 'absolutely wrong'. It cannot be proven from outside our concerns. But it seems like a position that we can hold from inside the practice. When we left behind the idea of Kantianism as a ground for deontological claims, we need not have left behind the intuition that makes his arguments so attractive: that some things we just should not do, no matter how valuable the consequences. Surely the claim that, for there to be certain things no one ought ever to do, there need to be some theoretical device we can't possibly deny, is one that needs further argument. What we were right to leave behind with Kant was the insane premise that there must be a way to prove that I ought to care about the things I care about, and that listing our actual connected concerns by way of accessing our actual moral conceptual repertoire was merely marching on the spot. It is the hidden premise that we can rationally prove to any and everyone, even those who have no care about our values, that they ought to think what we think, which Harris takes up in Kant's stead. In so doing he asks too much both of himself, us, and the limits of rationality and theory building in general.

If Harris claims, as he does about puppies, that this is 'soggy sentimentality' (p.13), that we are not seeing things aright, he can do so. His defence for his position however cannot be given a metaphysical angle without begging the question and

simply ignoring the epistemology that MaClean wishes to impart to us. If he dislikes the idea that ‘their being babies is a reason, all the reason in the world’ (221) he better get used to the fact that his own position is nothing but a laying out of the reasons that already exist in the moral realm for an opposing claim, ‘their being consciousness is the reason, all the reason in the world’. If he believes his position has more sway because it appeals to a theory, then he has been misled by a poor idea of the explanatory power of justification of a moral theory. If he feels any of that soggy sentimentality for babies, he might well want to consider changing his position to be in accord with the source of those feelings or else find a reason as to why this portion of the moral everyday life we live in can so easily be rejected while another part, that consciousness is valuable in all beings who exhibit it, is not so disposable. Furthermore he must do so in a way that does not make an illicit appeal to a Platonic picture.

For anyone worried regarding the point about slavery, we have already given good reason to doubt it. It seems now that had we access to theories of equality back when slavery was around then we wouldn’t have had slavery. Surely then, a theory which denies people can be treated fairly, rather than a conglomeration of mere attitudes, is what would stop such things happening again. But no theory can legitimately stop such injustices happening again. Principles such as ‘treat all men as equals’ were around before the end of slavery. We just chose to not call such people ‘people’ then. It is the broader concerns we share that allow such principles to find their mark. Rather than precluding slavery, it was events like the end of slavery – not due to any theoretical postulations, one might add – that made such theories seem so necessary. Attitudes feed the acceptance of theories; theories cannot change attitudes for those who do not believe in their premises. And such a mention of attitude needs to be addressed. My attitude that you are not a slave, nor ever could be one, is not the same as my mere attitude that I ought to run tomorrow or they should stop making Marvel films. My attitude is really an attitudinative-belief that gains its force from being situated in a balanced web of different ends that support and delineate each other.

I will leave this argument here. I do not want to get into the debate itself, but merely make the point about the forms of the two positions. In Harris’s argument we saw a

perfect example of the generalist tendency to miss the six points this final section started with:

Firstly, the uncodifiability theory was disregarded in his identification of moral status with a single feature: self-consciousness. In so doing he overlooked the varied uses of such terms as 'value', 'human being' and 'potential' that we utilise in everyday lives within context against a set of shared concerns.

Secondly, he derided the idea of rational disagreement. There must, he believed, be a rational way to answer all moral questions, even hard questions, and this would be done by bringing the right theory to bear on the situation. Thus, MacClean's position was represented as irrational, or holding notions that were beyond rational criticism, as opposed to expressing the moral realm in which these ideas find their meaning in a more full and textured manner.

Thirdly, by making 'person' a technical term and thus beyond the scope of moral argument, he put into service the idea that one could tell another the bare facts of the case and have that person describe what could morally be done. One didn't need to see the nuances of the case, as most of these were seen as relative values or soggy sentimentality. As far as absolute valuing goes context doesn't enter into the equation. One is rationally required to value self-consciousness.

Fourthly, the idea that we needed a theory to cut through the 'intuitions' and 'sentimentality' to be rational (as considered 'cold and calculating') meant that Harris promotes the idea that our emotional responses towards things are more of a hindrance to moral knowledge than a helping factor in finding nuanced reasons and in properly instilling an appropriate moral education in people. This would cause his position, when fully developed, problems regarding why people would be motivated towards making the sorts of decisions his theory would see as morally vindicated. One can only assume such motivations would come from the belief that his theory adds explanatory power, and thus make his pronouncements more rational. This I have shown to be false.

Fifthly, in his appeal to basic intuitions about persons who do have consciousness he missed the vital points that those like Cora Diamond and Alice Cray bring to the

table, by showing that in a deeper analysis we value things in ways which manifest not in simple theories but in entire ways of being.

Finally he misses the point about the possibility for moral growth that comes from a nuanced view of a conceptual network. Harris wants moral advancements to come as proofs entailed by universal moral theories. But this is a mischaracterisation of how we come to know morality in general. And in general, Harris just misses the point.

Harris's account makes the mistakes that are common in the more abstract theoretical side of Medical Ethics where sweeping discussions regarding complex ideas like personhood, cloning, animal rights and so on take centre stage. What is clear in this area, in British and American Medical Ethics particularly, is that the sort of rationality that this thesis has set out to reject as misguided, the one that puts touchstone definitions as a central pillar of coming to rational conclusions, still has extreme sway. Indeed, the most famous Medical Ethicists take this sort of line: Harris, Tooley, Rachels, Ruben and Singer being among those that do.

I would argue that a good way to address this sort of suggestion is not to focus on the conclusions as such, but directly find a meeting point of consensus, namely that the intuition that each is striving for (in Harris's case, that self-consciousness does indeed seem like a morally relevant property in various ways) has serious merit, along with the point that Harris makes in one of the passages I quoted above: that to begin thinking at all we must always borrow at least some intuitions from the common forms of life, thought of in those non-technical terms. If we edge out the generalist tendency of an appeal to a metaphysic, Harris's argument can actually get a better grounding for itself and be embellished with further intuitive notions from inside the practice of moral reasoning as it actually stands – that is, fully imbued with our actual way of life.

There is one final point to make. There is a constant threat from those who do seem to persist in these doomed theoretical exercises that those of us who do not share their optimism for rationality to show us a guiding path in complex cases cannot say nearly enough about what we actually need to do in these situations. But, it may be worth remembering there is nothing wrong with not saying anything definitive if one is not sure. On our picture, if someone does have a point to make it will not be

backed by some overarching generalist theory, but this does not mean it cannot be meaningful and effective. Positivity of the sort the generalist is looking for is based on an illusion. Let us stick to the truth and make the head way that allows.

7.3: Organising and Characterising the Practice:

One way of analysing the mistakes that Harris makes then, is by recognising that he tries to smuggle a highly contestable term, 'person', through as something that can be decided on by way of definitional and technical markers only. Once he does this, other parts of his arguments follow more easily, and we end up with dramatic moral conclusions because there was, hidden away in changing the term from an everyday usage to a philosophically loaded one, something at once both reductive and substantive. The same move is employed often in writings on clinical practice. The way Williams (Williams, A. 2012) makes a seemingly plain truth of substantive ideas of value regarding quality and quantity of life and then quickly moves towards a system in which it is applied has the same effect in his much discussed QALY system.

The notion of what factors make up such broad a term as 'quality of life' cannot be sectioned off and then marshalled for some systematic theory in a way where the output of that theory will seem meaningful and correct without us sharing the assumptions which narrowed the scope of the issue to allow for a systematic rendering in the first instance. If only we could know what we really ought to mean when we use these sorts of terms, the thought goes, we could really say something meaningful. The same thing can be said for assumedly technical attempts at deciding when life starts. While physical features can certainly be a major factor here, using them as general markers or rules for applying a term like 'morally valuable human' will have further substantive commitments one may not wish to defend but will feel like, on pain of inconsistency, they have to. Let us admit these are moral issues, with multi-faceted and complex interactions between the values in play that cannot be so easily systematised, as opposed to evaluations that follow directly from the recognition of general features obtaining.

The same battle has also been fought against broader concerns in health care recently, with arguments about what constitutes the best characterisation of the practice of health care. Again, what we have said seems to have serious issues to face in arriving at anything like a resolution. Contenders include Evidence-Based Medicine (EBM), Values Based Practice (VBP) and Person Centred Medicine (PCM).

At first blush it is hard to understand why these conceptions cannot be reconciled. I would like any medical practice I attend to use evidence, place me at the centre of their practice, and utilise values. There is also no immediately obvious reason as to why I ought to have to choose between these options. The problem, of course, is that these terms (evidence, values and persons) have more specialised and action guiding meanings to them in this context than our usual uses of them seem to require. In wanting to say something meaningful, each characterisation of practice places theoretical constructs about what medicine is 'really' all about in such a way as to encroach on other intuitive ideas regarding the other concepts. The assumptions we use to make evidence (specialised to mean the output of our best empirical techniques at data collection in EBM) centre stage seems to substantively go on of its own accord to lessen the trust between doctor and patient and reduce the doctor's role to that of (empirical) evidence collector and patient informer, and the patient's role to being a mechanical object with varying faults to be remedied. Theoretical arguments which fuel what counts as evidence (again, a normal term specified into a technical one in the name of being action-guiding) rule the roost here. I will use this space to try to show that there are certain theoretical assumptions that both EBM and VBP – which arises as a reaction to EBM so as to compliment it (Fulford, 2012) – hold which stops their seemingly sensible and respective suggestions that evidence and values ought to hold prime place in medical contexts.

Then again, are these 'sensible' suggestions? Doesn't that imply some thought has gone into it? The whole point of theorising in a definitional and reductive way, one might think, is precisely to not merely parrot platitudes. Terms like value, evidence, quality and best practice litter the clinical literature. But we all want these things; we are just not sure what they amount to. But this is down to the fact that these terms rise above practicality. They are ideals which we say something has for having

various other features. Once we realise that these features can be inchoate and express various, overlapping and inchoate concerns, we ought not to start with them and work backwards, but start instead at the bottom and work up towards them by addressing those same various, overlapping and inchoate concerns. A strict and principled system will not appear from this, but at least we will be interacting with the subject matter face to face, in the only way we can.

The intuitive idea that EBM seems to pick out is that which we have come to recognise in our thesis as grounding the ‘unfavourable comparison’ between matters of fact (empirically identified) and matters of value (as that which floats contingently upon these factual matters). Unsurprisingly, this technical application of ‘evidence’ will carry with it substantive baggage that makes seemingly innocuous modifications have wildly specific and controversial effects further down the line. If what we have said so far is true, that all judgements must have some grounding in shared commitments – to motivate us towards their acceptance – then this view can only be seen as radically mistaken. When a patient and doctor judge something to be a medical issue the idea of health is central. But unless we are to concede ground to the empiricist thinker and undermine what we have already said, we must view health as a concept that cannot be cashed out without evaluative notions informing our judgements.

Bad health is judged to be so not on some purely diagnostic empirical ruling that only utilises physical traits, then. It likewise takes judgements of contrasts against normal physical workings, the ability to interact with the world in a way that doesn’t have a negative effect upon one’s lifestyle and various other possible markers of what it is to be negatively impacted by the workings of one’s own body. What is healthy and unhealthy, on our picture, will necessarily be a judgement that is contrasted with a set of shared judgements that allow discussions of these concepts in action guiding ways. By EBM’s lights all evidence is made out at the physical level, with clinical epidemiology no longer being a central pillar for medicine, but the very ‘basis’ (EBM Working Group, 1992). The very idea that health had an evaluative aspect to it was simply not taken seriously by defenders of this position. Thus it was without irony that the EBM working group wrote that:

The proof of the pudding of evidence based medicine lies in whether patients cared for in this fashion enjoy *better health*. This proof is no more achievable for the new paradigm than it is for the old, for no long-term randomized trials of traditional and evidence-based medical education are likely to be carried out (ibid) (my emphasis)

The very notion that suggesting that ‘better health’ could be proven by randomised trials - the very thing the other paradigms would resist, and that which was under discussion – would be begging the question was seemingly beyond their imagining. That I ought to reject such a position on the basis of the epistemological account I have proposed is obvious. Evidence of any sort makes sense only against a shared background of values and commitments. That meta-analysis of empirical data is important is not even in dispute. But why we ought to run those trials, what sorts of illnesses we run them on, the design features of such trials and why they have such epistemic pull cannot be made out without the sorts of values that EBM rejects as satisfactory evidence. By the lights of my own position, for those with the eyes to see them, such values can be read off as descriptions of properties that are actually there, and so evidence stretches far beyond empirical data.

The particularism I have developed thus matches up with that which appears to be espoused by Thornton. While his papers (Thornton 2005 and 2011) do not allow the sort of thorough analysis of precisely how it is that we overturn the empiricist and generalist pictures, it does utilise the general conclusions we have come to here, regarding the idea that facts are laden with values even in seemingly empirical contexts – diagnosing illnesses say – and thus we needn’t shy away from the idea that when we are postulating about more obvious evaluative matters that we suddenly need to take on the idea that these are somehow ‘just our opinions’ or tastes on the matter. It is the particulars which decide the judgement, how the facts of the matter are actually laid out, because, as he puts it, we ‘are actually responding to real features of the world’. (Thornton, 2011: 991)

Underlying agreement in the factual are a whole host of evaluative assumptions; we see them as being accurate descriptions of things on evaluation. That a ‘malignant tumour is bad for us’ seems like a description of a state of affairs. But it certainly cannot be made out in purely empirical terms. It is the shared values of privation

from pain, enjoying physical privacy, not having a life disrupted by dramatic treatments, feeling generally fine, longevity of life (and so on) which make the statement 'cancer is bad' make sense in such a way as to feel like a descriptive one. Likewise, if a doctor says he has a good treatment for us, the 'goodness' will be predicated on the qualities the treatment has been shown to have weighed up against risks. These risks will be measured against evaluative notions of what a good outcome is and this will be made against an overarching conception of a good life. Chopping a leg off can be a very effective way to get rid of a possibly dangerous mole situated there, but the life altering effects of the treatment will just not be worth the effort given the other treatments available. These are not mere evaluations as opposed to empirical facts; they are real features of the world; real descriptions of the real properties of the case. If these are the only salient variables and someone disagrees, choosing amputation over having the mole frozen off, say, one does not merely hold a different opinion; they are actually wrong or crazy.

Thornton (2011) uses this sort of particularist analysis (however briefly stated) to assess the merits of one of EBM's rival accounts: Values Based Practice (VBP). Fulford's arguments lead him, at first, to seemingly particularist conclusions about how all judgements are imbued with values: 'all decisions stand on two feet' (Fulford, 2012: 131). His work on showing the centrality of value to decision making deserves praise. But he leaves particularism behind in his methods when, as Thornton points out, he moves away from a codified ethics; not - as the particularist would demand - towards the idea that value properties are there to be assessed in the world (being a part of the same tapestry as empirical claims) but towards a type of 'process' for assessing values, where 'the process [not the decision that results from that processing that] is the end itself' (Thornton, 2011: 991).

We have already seen the oddity of this sort of suggestion above. (7.1.1) Fulford's unwillingness to posit values as actually being there in the world to be considered - grounding agreement in judgements and allowing for understanding people's reasons for disagreement as well - becomes a major issue for his account. Value becomes both central to it and yet meaningless. Again, I believe his failure is predicated upon a sense of truth that is somehow 'out there anyway' and I will aim to show this.

If we get cancer we agree it is bad ('bad' standing as a veridical judgement on top of the multitude of shared judgements listed above) and if there is a quick and generally painless treatment available, over one which is longer or more risky, it is shared judgements based on shared commitments which allow us to reach a quick consensus about which treatment to go for. This is the picture I share with Fulford. I have argued at length that this places values within the fabric of the world above the bedrock level - that it is this picture that demands epistemological primacy - backed up with the claim that meaning and normativity are impossible to explain if perception is seen as evaluatively innocent. But when values clash in dramatic ways, there then seems to be a serious problem for particularism; for who do we listen to, who do we give the authority to in order to say that *this* should be the way to go? By the lights of the unfavourable comparison this places us in a predicament that doesn't (usually) exist for empirical claims. There we turn to an expert, for clarification, say. The comparison pushes up against particularism in a seemingly vicious way. Without the possibility of proof, how can we say to another that, against their seemingly sincere declaration to the contrary – they are simply wrong in their judgements?

Fulford seems to think we simply cannot. He shirks the idea that we ought to try and actually get things right, all thing told; to look for, in his own words: 'pre-set right outcomes' (Fulford, 2012: 24) Cassidy (2013) has argued that Fulford's position dissolves into a form of subjectivism. I agree, if for slightly different reasons. I will now aim to show this, using Thornton's article as a guide but hopefully filling it out as I go.

The way that Fulford sees the picture we seem to share is quite different. Yes we reach a consensus that cancer is bad and this is the best treatment. But the values are not there like the tumour is there. We just happen to share them.

This is still a case of our casting something onto the world. Where I would say we share the values surrounding cancer because we have good reason to be against it, Fulford would seemingly say that the reasons we think we have are predicated upon a local or contingent consensus of values, none of which have no further court of appeal for claiming primacy. Thus, if someone loves their tumour dearly and names it, and thinks all the pain is well worth it to have it around, that person's values merely differ to ours.

Fulford thus pushes his position towards 'answering' ethical questions by way of a process. And the process does not make reference to its conclusion being the right one, decided, as Thornton puts it, by way of 'antecedently real moral properties' (Thornton, 2011: 991). Here is Fulford:

Values-based practice aims to support balanced decision-making within a framework of shared values, based on a premise of mutual respect and relying for its practical effectiveness on good process rather than pre-set right outcomes. (Fulford, 2012: 24)

Fulford, then, gets caught up on the point of proof. It is an epistemological concern which is causing the worry. 'We cannot prove what we ought to do' is thus run together with this claim: 'so we cannot say that our judgement is better than anyone else's'. Fulford does not see this as incorrect, nor any sort of loss. He merely sees it as the fact of disagreement, given that values 'are complex in the sense that they are capable of very different interpretations by different people and in different situations'. (Fulford, 2012: 41) While values are complex and we indeed cannot prove to someone they ought to change their values this does not preclude the idea that one party is not more correct. In our picture this is guided by the normative constraints of the sorts of values in play. Whereas Fulford says:

...that many values, even though shared, are complex in the sense that they are capable of very different interpretations by different people and in different situations (Fulford, 2012: 41)

The mention of many interpretations brings to mind theorists like Wright; and this idea of the endless interpretation of the evaluative will bring similar problems for Fulford. For if values are infinitely interpretable, such that no one should be accorded with being more right, however diverse their interpretation is, then the same goes, by way of our Wittgensteinian account, for empirical judgements as well. Fulford seems to separate the empirical and the evaluative, with only the former being properly true. In fact, this thesis makes part of his account for being able to say that mental health and physical health diagnoses are not so different. Following Hare he makes the claim that all values supervene upon the empirical; that which is really there anyway. It is just the values we share that seem to be intrinsic goods. Whereas the ones we do not wear their contingency of our personal inclinations upon their face. The

uncontroversial health issues of physical ailments are thus so because they are predicated on heavily shared evaluations that supervene upon the empirical world, whereas mental health ones are complex and thus open to vast interpretation.

But this interpretative position falls precisely into the same complaints we posited against Wright. And it is dangerous for Fulford. It again splits into the two levels, things we can say are right and wrong at the empirical level 'the tumour is there or is not', and the level of evaluation which supervenes upon this by way of shared judgements: 'and the tumour being there is bad - and only now with the added evaluation is it a health issue'.

But as we have seen, minus shared commitments, empirical judgements are open to an equal amount of interpretation as the evaluative; because empirical claims are not self-evident without shared commitments guiding our evaluations of how their extensions are to be made out; if they were self-evident we would require the suspect psychological mechanism (4.2.3) to explain them. This realisation, one hopes, might go some way to relieving Fulford's fear that the lack of proof for holding one value over another means that neither party is, in fact, right or wrong.

Thus, it would seem that Fulford's commitment to Hare, when analysed in the Wittgensteinian manner, collapses into something like Wright's analysis. If we can interpret evaluative claims infinitely and there is no way for accounting for shared empirical judgements accept against a background of shared evaluative commitments then empirical judgements also are infinitely interpretable.

On the particularist picture I am endorsing both positions fail. Hare's position illicitly makes the claim that empirical judgements can be made out without shared commitments - which requires the idea of the suspect psychic mechanism - while Wright's ends all hope for actual normativity and the intuition of objective choices even in the empirical claim. But Hare's position fails because it collapses into Wright's. Thus the idea that there can be no antecedent evaluative truths becomes the claim that there can be no truths, just a conglomeration of perspectives. Thornton points out that some theorists who wish to make mental health more normalised have tried to square it off with physical ailments by showing that both can be analysed in purely empirical terms. Fulford instead shows that both are necessarily

imbued with value. With this I agree. What I disagree with is then going on to make the claim that those evaluations are not up for analysis as better or worse.

What I would claim is that the shared commitments we develop to become active members in a form of life – the only background against which we can have normativity at all - normatively commit us (if we have things right) to using these concepts in a way that goes beyond how people happen to use them at any particular time. The community and the form of life can have a wedge driven between them. This is why, for example, Americans can list values they hold dearly – freedom of expression, justice as fairness, diverse political and societal values, innocence before being proven guilty, and freedom of speech – and then act against all of these when a politico-religious state they are unfamiliar with, half the world away, expresses an interest in nuclear armament. Most Americans would not say all these values listed above are trumped by distant concerns for safety, even though this is how many of them happen to act here and now. The universality of the values, inherent in the normative contract picture, extends beyond how they are used by a society. The concepts in play rise above how they function in particular instances and instead demarcate an ideal version of what we believe. Whether we follow them, or give in to weakness of will, is quite another matter.

On my thesis then, we have direct access to the moral properties that are present in a situation. We can say that there are in fact some answers that are better than others, precisely because the people proposing these answers have picked out values which are more readily present and pertinent there. This needn't work against Cavell's point, that there can be a rational disagreement. All I need to show is that there can, in principle, be some cases where rationality requires that one specific path of action is followed, despite the sincere protests of the many. This I have done, by showing that the alternatives to this picture, community ratification and empiricism, fail to give workable epistemological accounts of meaning. The only other possibility is that we are, with McDowell, contractually bound to speak against the majority if there are normative reasons to. Our analysis of VBP has extended Thornton's analysis by showing just why it is that particularism, and not an account of proceduralism, ought to be seen as the main alternative to generalism.

Proceduralism, then, seems to be based on a serious rejection of the idea that we cannot make generalist claims which allow for set rules. In so doing though it drives us towards an account of irrationalism when we bring it to its logical conclusion and show the unfavourable comparison for what it is. This seems like a pattern of assessment that is being quickly established in clinical evaluative practice. That there are no right answers seems to devolve quickly into the claim that all we can do is let everyone be heard in a way that respects their (however wildly) variant opinion. But proof and objectivity come apart. From the fact we cannot prove definitively which opinion is better, it does not follow that all opinions ought to be given equal weight.

7.4: Particularism and Practical Ethical Guidance:

As a further point, it is also fair to say that the real driving force behind Fulford's account is Liberalism. On the basis that we cannot prove any value to be right, we develop a new value: that all values are equal. While this value can make sense in certain contexts – say respecting someone's religious rights to deny treatment, even when we have strong reason to believe their god does not exist, so whatever judgement that a person believes will be cast upon them by this supernatural entity does not exist – it cannot be the default between all judgements; remember our patient wishing to amputate his leg to rid himself of a mole. Whereas Fulford would say this value's being shared is what makes it not need to be questioned, we can say it is not questioned because it actually is a powerful evaluative reason.

Liberalism made out of the case of lack of definitive appeal for values is surely self-defeating. It remains unclear why I ought to respect other's values if they clash with my own. If no value can be made out to definitively trump any other it does not follow that I ought to value all values equally or not speak against ones that I find heinous, ignorant or dangerous (if only for the person who holds the value). If someone makes the case that indeed I ought not to speak out, and I ask why, their answer will necessarily involve an evaluative position: that I am offending the liberalist value of self-legislation. In which case the liberalist value itself stops being equal to my own, it

necessarily demarcates the ways in which I can act upon my own values: but this would be in direct contradiction of itself.

It is crucial to point out, however, that it is a central pillar of my own thesis that, although there are evaluative truths, these truths cannot be formalised into a system which can guide actions before we know the context of the situation they are meant to guide. Even though I disagree with Fulford about the metaphysics of moral reasons, one might ask how particularists can say something meaningful and true about how medical practices ought to be organised and characterised. What principles could we build such a case around?

There is a positive and a negative point here. The positive point is that we can free ourselves from the theoretical ideas that have bewitched philosophers, such that if we are to say why evidence is important we must define evidence just so, or if we are to say autonomy is important, we must define autonomy just so. As such we can hold a whole plethora of values to be available to us without coming to theoretical implausible conclusions. When and what type of evidence ought to be used in certain situations will not be pre-decided upon by theories. We have already seen that such a move works backwards. Those theories will specialise the term in such a way for it to become a new philosophical entity by starting with an intuitive notion and generalise it in such a way as to make certain uses of it - which in everyday language seem perfectly acceptable - seem fallacious. What they call intuitions, the particularist just see as shared, if inchoate universal concerns, which can be specified in various ways depending on context, with no way to abstract from each case to find some general feature.

In each case better or worse specifying from the universal seems possible. And this itself seems to make saying something meaningfully action guiding close to impossible. This picture places us back with the issue of expertise. The requisite sense of expertise for moral, clinical and politicised decision making seems to defy at every turn the use of generality. Being inchoate in this way makes it a practical skill; an art, not a science; know-how, not knowledge-that.

From the back of this comes the negative point. We cannot design a characteristic form for clinical practice or for the imposition of ethical concerns into practice which will give us the right outcome. To say otherwise would be to shirk the commitments

of my own arguments. To be fully aware of which of the universal, inchoate, overlapping and sometimes incommensurable concerns apply here and now I really need to be present and vigilant within the situation that demands it. This does not mean we cannot list appropriate concerns. But if start to systematise them from outside a position where they might actually apply, without the contextual perception of all the salient issues, we will soon be back to fallacious generalising.

Thus, with this picture in place, we can start to see that moral education needs to be practical. Particularism has to be taken seriously if we are to pay due respect to the way the moral terrain actually encroaches upon our reasoning. We can look for universals and see how and where they apply in certain situations. As Thornton's own work has made clear, however, such a picture, based on particularism looks more descriptive than prescriptive. His analysis of the Four Principles account, proffered by Beauchamps (2003) makes this point. Beauchamps pays heed to the idea of incommensurability and also gives, when read properly and supplied with a particularist metaphysic of morality (Thornton, 2005), a theory that gives a full role to judgement. Either specifying or balancing the universal principle they supply there cannot really give us a mechanical answer to problems. Both balancing and specifying are open to judgement. Thornton (2005) seems to give the impression that balancing has a more substantive slant to it, but specification allows the exact same issue to arise. Universal principles do not unpack themselves. Each one has various possible specifications depending on the contextual features that are there. As there is no general feature that can be abstracted from them and in this sense judgement or practical wisdom is called for.

It likewise means that the person who comes to do the specifying or balancing sees the principles. This is a crucial point that cannot be missed. Just because some rules are not open to interpretation, it does not mean that others are not. We have no further place to look for certainty except to our own judgements that have been built up, via inculcation into moral life, and those who we see as wise for advice on such issues. This prompts the question, how do we know we were inculcated correctly and have seen the light in a way that, say, those whose attitudes to some persons was 'an attitude towards a slave'?

Such a fear is warranted upon this account. Here is Onora O'Neill (1996) on the issue:

Particularist strategies...to discern who has ethical standing by inspecting the scope of our community or nation or relationships are often decently clad in the vocabulary of commitment, care, and attachment, but...those who fix the scope of ethical consideration by reference to [these] limits... must exclude from ethical consideration those for whom 'we' care not a whit or... see as outside 'our' community...

Such a fear lingers, but as I have already said before, how we read the sorts of rules that O'Neill is implying can put an end to what she sees as cultural subjectivity in fact depend upon our attitudes to attain any motivity. To believe that rules defeat attitudes is to misread this entire thesis. The rule is made clear to perception as meaningful by way of the attitudes we have. That the rule 'treat all rational people as equals' has such an effect now shows our attitudes are going in the right direction. Again, this is not to say that the attitudes that people had towards slavery then were right by fiat, just by them having that attitude. They were wrong. There were enough conceptual tools present for the moral reality of what they were doing to strike them as being wrong. The same goes for Nazi Germany. If we can learn anything from such examples it is that we have to be careful about honouring our normative commitments to all our virtuous concepts in the full.

As Aurel Kolnai (1938) emphasised so long ago, those who were present at the beginning of the Nazi party did not see themselves as being pure evil. They saw themselves as upholding a very specific set of values. The issue is that these could only be made out at the expense of all others. The sorts of dispositions we all know lead to disaster – stubbornness, pride, unquestioning duty, courage stripped of mercy, obedience, greed – were all present there. This is why moral education is so important. Rules only get us so far if we are not raised with the right eyes to see why they are important. It is this which makes the intuitive claims of each of the touchstone theories so appealing. Autonomy, self-legislation, happiness, flourishing, the continuation of our species, freedom and so on; the list goes on. But it is a mistake to specialise these terms in order for them to be generalised into general principles for action. We feel without the principles we have nothing to appeal to.

This is simply a mistake. The list of concerns goes on; it is simply that none are utterly definitive to the man who is wilfully ignorant. But it is only philosophers who are wilfully ignorant. Such a position is precisely one we put ourselves into to look for some certainty. Having failed to do this with morality, however, can we really turn around and claim that the whole project is therefore dispensable? Are these still not still our reasons and are they not the right reasons? Other than the wilfully ignorant, we have to deal with the actually ignorant if we do not learn the lessons of particularism and educate people in the right way, to see morality not as a theoretical exercise, but as an exercise of practical reason, of nurturing the right sorts of dispositions and perceptive states to be able to see such properties as and when they obtain.

7.5: Getting Moral Discourse Right – Concluding Remarks

I would like to finish this thesis with an analysis of a paper that I believe to really combine all of the concerns that I have outlaid here. It is a paper that pleads, without appeal to any one overarching theoretical position backing up its central claims: that doctors ought not to be those to whom the burden of refusing treatments falls to. The author is Vegard Bruun Wyller (2015)

His concerns are that of a physician who has first-hand experience of bedside rationing. His points are lucid and well considered. The bulk of the argument points out that there is a theoretical clash between, firstly, the everyday ideas and concerns that make up the ethos of the doctor looking after his particular patients and the psychological, societal and practical elements which allow for trust to foster between the two, and the rationalistic and principled sort of theorising that underscores such ideas as distributive justice and health economics. The latter, if taken as something that needs to be done, belongs, claims Wyller, without appealing to any rationalistic grounding about societal identities, to the realm of political decision making. This is because here, a first order evaluation shows us that the trust is not predicated on the belief that the politician to whom it falls has specifically *my* best interest at heart. Instead, we assume that the politician has precisely every member of his state or constituency to look after and cannot be so personalised. The societal roles of doctor and politician, which allowed for this division of labour in such a way as to allow each

role to exhibit the correct sort of concern for the individual – in the case of the doctor as an actual individual with a name a history and a character, and in the case of the politician as one of his herd – has thus been subverted. This subversion brings with it a diminishing of trust in terms of the patient and the doctor, and a diminishing of meaning between the politician and his electorate.

While Wyller takes resources from various philosophical positions including Aristotle and Levinas, this seems to be largely to add some academic clout to his position which is just fine as an argument without it. At no point does Wyller attempt to define any of the key terms that are pertinent to his positive arguments. The concepts of politics, medicine, trust, care and commitment are not defined to suit his purposes, but are left as ideas we can all understand without their identity and comprehensibility being cleaned up to suit rationalist purposes. The way his paper makes use of the idea of ‘by the bedside’ to get across a picture of the actual state of affairs we are discussing, placing us in the position of someone being treated, rather than taking some disinterested stance on the matter. Indeed it is the human element of medicine which shines through in his critique of health economics, of turning people to mere numerical values, which strikes at the heart of the matter there as well.

This celebration of this piece of work – which does not stand alone, various authors write with equal compassion and lack of regard for overt-rationalisations (see: Diamond, 1978) – is not overtly academic. The point that stands is how well we can in fact marshal our actual many and distinct concerns into very real, very meaningful expositions of what it is that really matters in ethics without illicitly breaching the parameters in which those ideas can be sensibly discussed and looking for definitions from outside the domain to gain a false sense of clarity or certainty. The inchoate concerns Wyller brings out are not generalised into specific arguments that aim at abstracted principles. Indeed, he sees right through the attempts to aim for such principles when they are applied at the bedside, readily admitting that at the political level, where the trust relationship is precisely set up to not take into account individualistic credentials, generalised principles will be a necessary factor in decision making. He does not take on the voice of someone who is bewitched by the idea of finding the right answer when that is taken as being conceived as a generalised principle. He is an appropriately concerned member of the medical

profession that can see that values that elude the sort of quantification that is so commonplace in the usual discourse are being eroded in such a way as to genuinely distort the sort of relationship that is necessary for clinical care, as abstracted from the multiple specifications that might take given by context.

That this argument was instantly met by an overtly rationalist retort by Magelsson (2016) is testament to what we, as particularists are up against. But it is contributions to the debate such as this one which are crucial to the ongoing battle to depersonalise that which is most personal. His argument gives not only a complaint, that health economics by way of trying to make things 'fairer' (in a theoretically specialised way) is eroding trust; with proof of that erosion and its negative effects being based on first-hand experience and a claim to 'intrinsic, primary value'. Wyller also diagnoses with incredible clarity the reason for this – a movement of power from the political to the bureaucratic which in turn suggests a programmatic remedy: reverse the power transfer and we will regain the trust. In this article, then, we find the embodiment of the particularist position.

It is a piece which exhibits perfectly the five points that I listed at the start of this section as being hallmarks of a particularist argument. Firstly he makes no attempt to define in a way that carries loaded meaning any of the key terms his argument is based around. The notions of trust, care, compassion, power, the division of the political and the medical – each concept is used as we do use them without a specialised definition. The doctor-patient relationship is given an 'ethos' rather than a substantive or loaded definition.

Secondly, Wyller keeps open the possibility for rational disagreement to specific cases being rationed. He does not make, as Magelssen (2016) claims, some overarching claim that rationing *as such* is immoral. His argument could not be used as some sort of definitive proof we ought not to ration here and now. It allows for its concerns to be set aside as and when judgement demands, such as difficult triage situations.

Thirdly, we come to proximity. Wyller's argument is as good as it is, I maintain, precisely because he is a physician. As such he has a more nuanced and discerning grasp of the sorts of values that are crucial to his role. Rather than being observed and analysed into general principles by an outsider such as an ethicist, his day to

day interaction with patients gives him the sort of insight which is indispensable to understanding values in a medical setting.

Fourthly, his arguments for keeping certain values are motivational, in the sense that they make us look at the situation as it is: imbued with the sorts of values we expect within a medical context. The concerns he has are specific to its cause, not abstracted from rational analysis and just applied to a new context. To abstract is to take out the very details that are essential to the appropriate values forming.

Fifthly, his argument is not a philosophical explanatory theory, which we have already made a significant case against. As previously noted, Wyller utilises insights from Aristotle and Levinas, amongst others. But using the insights of philosophers is not enough to warrant the label 'philosophical explanatory theory'. His argument is distinctively moral and cast in moral language and expresses clear moral concerns. He uses these philosophers precisely because their own work eschews the same notion of over-rationalisations. He does not build a theory, as much as he makes a case from first order moral concerns that are our own; concerns that he knows that we can all relate to because we live against a backdrop of shared practices and concerns. These are the sorts of arguments that particularists should not just accept, but the sort they ought to promote. It says something meaningful and true; it directs our attention to a distinct problem and highlights particular areas of concern (trust and doctor/patient relationships) as areas which could meaningfully fix the faults it finds there. In so doing, he shows how a tradition in one set of norms (the desire for algorithmic ethical answers) is causing issues with our most firmly held moral convictions.

If someone disagrees with Wyller, they can do. To do so and be convincing one does not need to step outside of the first-order concerns we share. The point of a non-reductive morality is that we can always make contact with someone else's position. As long as we are not talking past each other – which is simply not possible if we think properly about translatability – then there will always be some point of agreement. If we add to this the idea that we can start with any shared norms and work towards changing an outlook, such that we do not need to share first premises or principles in order to change someone's mind, it seems like this position, properly conducted within a first order moral account, properly defended, could well have

greater positive and effective impact on the moral discourse than any overarching generalist theory could wish to hope.

In this manner it is fair to say that if these sorts of arguments are the way to go then the threat that is posed by generalism is not merely a theoretical one, it is itself a distinctively moral threat. Particularist theorists seem to spend most of their time undoing the issues caused by rash and unthoughtful generalism. We are trading in real human concerns for rough approximations and quantifications and doing so in the name of a serious and distinctive claim to reason, not a caricature of rationality. This generalist academic fashion needs to end if moral discourse is to find its feet once again. Particularist contributors, like Wyller, will hopefully put an end to it.

References

- Anscombe, G.E.M. (1957) *Intention (2nd ed)* Basil Blackwell
- Anscombe, G.E.M. (1989) 'Von Wright on Practical Inference' in *The Philosophy of Georg Henrik von Wright* (eds) Paul Arthur Schilip and Lewis Edwin Hahn, Chicago: Open Court pp. 377-404
- Aristotle (1955) *Nicomachean Ethics* Penguin Classics
- Austin, J. L. (1962) *How to do things with words* Oxford University Press
- Beauchamp TL (2003) 'Methods and principles in biomedical ethics.' *Journal of Medical Ethics* (29) pp. 269–74
- Blackburn, S. (1992) 'Through Thick and Thin' *Proceedings of the Aristotelian Society* (66), pp. 284-9
- Brink, D. (1989) *Moral Realism and the Foundations of Ethics*, Cambridge: Cambridge University Press
- Cassidy, B. (2013) 'Uncovering values-based practice: VBP's implicit commitments to subjectivism and relativism' *Journal of evaluation in clinical practice*, 19 (3), pp. 547-552
- Cavell, S. (1979) *The Claim of Reason: Wittgenstein, Skepticism, Morality, and Tragedy* Oxford University Press
- Cavell, S. (1969) *Must We Mean What We Say?* Charles Scribner's Sons
- Crisp, R. (2014). 'Taking Stock of Utilitarianism'. *Utilitas*, 26(3), pp. 231-249.
- Croom, A.M., (2010) 'Thick concepts, non-cognitivism, and Wittgenstein's rule-following considerations'. *South African Journal of Philosophy* 29 (3), pp.209-232
- Dancy, J. (2004) *Ethics Without Principles*, Oxford University Press

- Diamond, C. (1978) 'Eating Meat and Eating People' *Philosophy* 53 (206) pp. 465-479
- Eklund, M. (2011) 'What Are Thick Concepts?' *Canadian Journal of Philosophy*, 41(1): 25–49
- Evidence-based Medicine Working Group (1992) Evidence-based medicine. A new approach to teaching the practice of medicine. *Journal of the American Medical Association* (268) pp.2420 – 2425
- Freeman, S. (1994). 'Utilitarianism, Deontology, and the Priority of Right'. *Philosophy & Public Affairs*, 23 (4), pp. 313–349.
- Fulford, K.W.M. (2012) *Essential Values-Based Practice: Clinical Stories linking Science with People* Cambridge University Press
- Garfield, J. (2000) 'Particularity and principle: the structure of moral knowledge' *In* Hooker, B. and Little, M.O. (eds) *Moral Particularism*. New York pp. 178 – 205
- Gettier, E. L. (1963) 'Is Justified True Belief Knowledge?' *Analysis* 23 pp. 121-3 [online first]
- Gibbard, A. (1992) 'Thick Concepts and Warrant for Feelings'. *Proceedings of the Aristotelian Society, Supplementary Volumes* 66 pp.267-284 [online first]
- Gibbard, A. (2003) *Thinking How to Live* Cambridge MA: Harvard University Press
- Grice, P.(1991) *The Conception of Value* Oxford: Oxford University Press
- Harcourt, E. and Thomas, A. (2013) 'Thick Concepts, Analysis, and Reductionism', in Kirchin, S. *Thick Concepts* Oxford: Oxford University Press pp. 20–43
- Hare, R.M. (1963) 'Descriptivism' *Proceedings of the British Academy* (49) reprinted in *Essays on the Moral Concepts*, Berkeley: University of California Press pp. 55–75
- Hare, R.M. (1969) *Freedom and Reason* Oxford University Press

- Harris, J. (1985) *The Value of Life* Routledge and Kegan Paul
- Harris, J. (1995) 'The Elimination of Morality' *Journal of Medical Ethics*, Vol. 21, (4) pp. 220-224
- Hurley, S.L. (1989) *Natural Reasons*, Oxford: Oxford University Press
- Hursthouse, R. (1999) *On Virtue Ethics* New York: Oxford University Press
- Irwin, T. H. (1975) 'Aristotle on Reason, Desire, and Virtue.' *The Journal of Philosophy*, Vol. 72 (17) pp. 567-578
- Kagan, S. (1989) *The Limits of Morality* Oxford: Clarendon Press
- Kant, I. (1948) *Groundwork for the Metaphysics of Morals* trans 'The Moral Law' by Paton, H.J. Hutchinson
- Korsgaard, C. (1986). 'The right to lie: Kant on dealing with evil' *Philosophy and Public Affairs* 15 (4) pp. 325-349
- Korsgaard, C (1996) *Creating the Kingdom of Ends* Cambridge University Press
- Kripke, S. (1982) *Wittgenstein: On Rules and Private Language*, Cambridge: Harvard University Press
- Kolnai, A. (1938) *The War Against the West* Gollancz London
- Little, M. O. (2001) 'Lessons on Moral Particularism' in *Slow Cures and Bad Philosophers: Essays on Wittgenstein, Medicine and Bioethics* Elliot, C. (ed) Duke University Press pp. 161-180
- Loughlin, M. (2002) *Ethics, Management and Mythology* Radcliff Medical Press
- MacIntyre, A. (1981) *After Virtue* (2nd ed) University of Notre Dame Press
- Mackie, J.L. (1977) *Ethics: Inventing Right & Wrong* Harmondsworth: Penguin

- **MacClean, A. (1993) *The Elimination of Morality* Routledge**
- **Magelssen M, Nortvedt P, Solbakk JH (2016) ‘Rationing at the bedside: Immoral or unavoidable?’ *Clinical Ethics* [online first]**
- **McDowell, J. (1979) ‘Virtue and Reason’ in McDowell, J.H., 1998. *Mind, Value, and Reality* (1998) Harvard University Press pp. 50-73**
- **McDowell, J. (1981) ‘Non-Cognitivism and Rule Following’ in McDowell, J.H., 1998. *Mind, Value, and Reality* (1998) Harvard University Press pp. 198-218**
- **McDowell, J. (1984) ‘Wittgenstein on Following a Rule’ in McDowell, J.H., 1998. *Mind, Value, and Reality* (1998) Harvard University Press pp. 221-262**
- **McDowell, J. (1996) ‘Two Sorts of Naturalism’ in McDowell, J.H., 1998. *Mind, Value, and Reality* (1998) Harvard University Press pp. 167-197**
- **McDowell, J. (2009) *Wittgensteinian “Quietism”* Common Knowledge vol 15 (3) pp 365 - 372**
- **Mill, J.S. (2001) *Utilitarianism*, Batoche Books**
- **Moore, G.E. (1922) *Principia Ethica* Cambridge University Press**
- **O’Neill, O. (1996) *Towards Justice* Cambridge University Press**
- **Putnam, H. (2002) *The Collapse of the Fact/Value Dichotomy* Harvard University Press**
- **Quine, W.V., (1954), ‘Carnap and Logical Truth’, in his *The Ways of Paradox and Other Essays*, revised edition, Cambridge, MA: Harvard University Press, 1976, pp. 107–32**
- **Rawls, J. (1971) *A Theory of Justice* Cambridge, Mass: Harvard University Press**
- **Richardson, H.S. (1994) *Practical Reasoning About Final Ends* Cambridge University Press**

- Ross, W. D. (1930) *The Right and the Good* Oxford: Oxford University Press
- Scheffler, S. (1987) *Morality Through Thick and Thin: a Critical Notice of Ethics and the Limits of Philosophy [Williams 1985]* *Philosophical Review* 96(3) pp 411–34
- Schopenhauer, A. (1995) *On the Basis of Morality*, (trans) E.F.J. Payne. Providence pp. 72 -73
- Seabright, P. (1987) 'Explaining Cultural Divergence: A Wittgensteinian Paradox.' *The Journal of Philosophy* 84 (1) pp. 11–27
- Sidgwick, H. (1981) *The Methods of Ethics* (7th ed) Hackett
- Simons, Daniel J., and Christopher F. Chabris (1999) 'Gorillas in our midst: Sustained Inattentional blindness for dynamic events.' *Perception* 28 (9) pp. 1059-1074
- Stocker, M. (1979) 'Desiring the Bad: An Essay in Moral Psychology', *The Journal of Philosophy* 76 pp. 738–53
- Svavarsdottir, S. (1999) 'Moral Cognitivism and Motivation' *Philosophical Review* 108: pp. 161–219.
- Thomson, J. (1997) 'The Right and the Good.' *The Journal of Philosophy* 94 (6) pp. 273-298
- Thornton, T. (2011) 'Radical liberal values-based practice' *Journal of Evaluation in Clinical Practice* 17 (5) pp. 988-991
- Thornton, T. (2005) 'Judgement and the role of the metaphysics of values in medical ethics' *Journal of Medical Ethics* 32 (6), pp. 365-370
- Väyrynen, P (2009) 'Objectionable Thick Concepts and Denials' *Philosophical Perspectives* (23) pp. 439 -469
- Väyrynen, P. (2011) "Thick Concepts and Variability", *Philosophers' Imprint* 11 (1) pp 1 - 17

- von Wright (1983) *Practical Reason: Philosophical Papers Vol. 1* Basil Blackwell
- Wiggins, D. (2000) *Needs, Values, Truth* 3rd edn. Oxford: Blackwells
- Wiggins, D. (2006) *Ethics: Twelve Lectures on the Philosophy of Morality* Penguin Books
- Williams, A.H. (2012) 'The value of QALYS' in *Arguing About Bioethics* pp. 423-427
- Williams, B. (1973) *Utilitarianism: For and Against*, with J.J.C. Smart, Cambridge: Cambridge University Press
- Williams, B. (1985) *Ethics and the Limits of Philosophy* Cambridge: Harvard University Press p.141
- Wittgenstein, L. (1953) *Philosophical Investigations* Basil Blackwell
- Wittgenstein, L. (1969) *On Certainty* Basil Blackwell
- Wittgenstein, L. (1978) *Remarks on the Foundations of Mathematics*. Basil Blackwell
- Wright, C. (1980) *Wittgenstein on the Foundations of Mathematics*, Duckworth London
- Wright, C. (1995) 'Truth in Ethics' *Ratio* 8 (3) pp. 209–26
- Wyller, V.B. (2015) 'Give to the Doctor what is Due to the Doctor: Why Fair Bedside Rationing is Impossible.' In *Fair Resource Allocation and Rationing at the Bedside* (eds) Danis, M. (et al) pp. 253-262
- Zoloth-Dorfman, L. and Rubin, S. (1997) 'Navigators and Captains: Expertise in Clinical Consultation' *Theoretical Medicine* 18 (4) pp. 421-432