

Thick concepts and reasons for action

Rebecca Claire Davis

MPhil

University of York

Department of Philosophy

October 2022

Abstract

This thesis aims to clarify the claim that thick concepts provide reasons for action, and explore whether any version of that claim can be defended. The claim is notably connected with the philosopher Bernard Williams, but here I do not approach it as a defence of Williams or his views. I understand reasons for action to be facts that stand in a reason relation to agents and actions. Reasons for action can be normative reasons *for* an agent to act or reasons *why* she acts. In order to defend any version of the claim that thick concepts are distinctive in providing reasons for action, it is necessary to take a position on thick concepts and defend them as evaluative concepts. My conclusion is that some understandings of the claim can be defended, albeit in a restricted range of applications.

Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as References.

Table of Contents

Abstract	2
Chapter 1: Introduction	5
1.1 Introduction	5
1.2 Twentieth century debates about non-cognitivism	11
1.3 More recent thick concept debates.....	13
1.4 Reasons for action	14
1.5 Thick concepts and reasons for action - further arguments	15
1.6 Conclusion	18
Chapter 2: Twentieth century debates: cognitivism, non-cognitivism and thick concepts	19
2.1 Introduction	19
2.2 Expressivism, projectivism and quasi-realism	21
2.3 Objections to non-cognitivism and expressivism	27
2.4 Blackburn and Dancy on thick concepts	33
2.5 Conclusion	39
Chapter 3: More recent debates: thick concepts and evaluation.....	41
3.1 Introduction	41
3.2 Thick and thin	42
3.3 Disentangling arguments: shapelessness and graspability	44
3.4 Daniel Elstein and Thomas Hurka: a reductive analysis.....	46
3.5 Pekka Väyrynen: The Pragmatic View.....	51
3.6 Debbie Roberts: Irreducible Thickness.....	56
3.7 Simon Kirchin: The Liberal View	60
3.8 Conclusion	62
Chapter 4: Reasons for action and thick concepts	63
4.1 Introduction	63
4.2 Actions	65
4.3 Reasons	67
4.4 Thick concepts and normative reasons on an externalist view.....	71
4.5 Thick concepts and normative reasons on an internalist view.	75
4.6 Thick concepts and motivating reasons.....	80
4.7 Thick concepts and explanatory reasons.....	83
4.8 Conclusion	84
Chapter 5: Reasons for action and thick concepts: further arguments	87
5.1 Introduction	87
5.2 Programme explanation	89
5.3 Programme explanation, thick concepts and normative reasons	93
5.4 The evaluation of actions	96
5.5 Virtue concepts and the virtuous person	97
5.5 Particularism, generalism and virtue concepts	100
5.6 Conclusion	103
Chapter 6: Conclusion	104

Chapter 1: Introduction

1.1 Introduction

‘Thick concepts provide reasons for action.’

This is the claim that I want to explore and ultimately defend in this thesis. In this introductory chapter I will introduce thick concepts and reasons for action and say something about the philosophical issues involved in making the claim that thick concepts provide reasons for action. I will then outline the contents of the remaining chapters. To begin with though, let me say that I think the claim has intuitive appeal. Allow for a moment that *KIND*, *CRUEL*, and *TACTLESS* are all examples of thick concepts. I think the following examples of the application of those thick concepts in reason statements are straightforward, if dull:

The fact that it is *KIND* to walk my friend’s dog for her is a reason for me to walk my friend’s dog.

The fact that it is *CRUEL* to kick the cat is a reason for me not to kick the cat.

The fact that it is *TACTLESS* to remind my colleague that she failed her driving test is a reason not to remind my colleague that she failed her driving test.

These examples may be boring, and one immediately thinks that there could be all sorts of reasons for me to walk my friend’s dog, but the statements by themselves don’t look odd or fantastic. So why are there some philosophical issues?

The first issue is whether there are such things as thick concepts, and if there are indeed such things, whether they are in some way distinctive and interesting in terms of providing reasons for action. Perhaps all concepts provide reasons for action. ‘Blue’ is not a thick concept, but the fact that a dress is blue may be a reason for me to choose it (if I like that particular colour, for example, or if I am trying to match the colour scheme for my friend’s wedding). So if there is something distinctive to the claim that thick concepts provide reasons for action, it ought to involve distinguishing between thick concepts and other concepts that

are not thick. Are thick concepts particularly apt to function as reasons, for example? Might thick concepts always function as reasons, or function as reasons independent of subjective motivations (such as the motivation to match my friend's colour scheme)? The second issue is that when we ask what reasons there are for an agent to act, for her to do something, we could be asking what reasons favour the action or make it a good or right thing to do to some degree, or we could be asking what reasons explain why she acts in that way. In making the claim that thick concepts provide reasons for action, are we talking about reasons that favour an action or reasons that explain it, or both?

Also, the claim that thick concepts provide reasons for action raises the more general question of how it is that concepts can provide reasons for action. For this last issue later on I will just set out the understanding or apparatus that I am working with. I take the view that reasons are facts that stand in a reason relation to an agent and an action. Facts are true propositions, and the atomic building blocks of propositions are *concepts* predicated of an object.

A note on terminology: in the examples just given, I denoted the putative thick concepts in small capital letters - KIND, CRUEL. I shall continue to use small capital letters to indicate thick concepts or putative thick concepts throughout the thesis, except in direct quotation. What are concepts? Following Frege we can say that there are objects and predicates. Predicates refer to *properties* or *concepts*. (Trueman 2021) So, the notion of a property or concept is what can be predicated of an object; objects can have properties or *fall under* concepts. (In this thesis I am using *property* and *concept* interchangeably; for example anything and everything that falls under the concept COURAGEOUS also has the property 'courageous' i.e. they are exactly co-extensive. I am not taking a view on whether concepts are identical with properties.¹) In this thesis, terms² refer to concepts; we can speak of the meaning of terms or the content of concepts. What I will call 'thick terms' are the words we use to refer to thick concepts.

1.1.1 Thick concepts

Thick concepts are typically understood as concepts that have both an evaluative dimension and a non-evaluative (often called descriptive) dimension. David Wiggins (1976) clarifies the use of the word 'value' and distinguishes between evaluations and directives.

¹ See for example Strawson (2000) who thinks we should be cautious in arguing ontological differences on the basis of syntactic differences. I am taking the approach for reasons of economy of exposition; the thick concepts literature talks about *concepts*, much of the rest talks about *properties*.

² In a Fregean treatment terms refer to objects. I need 'term' for a different purpose. in the thick concepts literature some writers talk about 'thick terms and concepts' and some writers think there can be thick terms but not thick concepts, so I need to keep 'term' for that distinction.

I propose that we distinguish between *evaluations* (typically recorded by such forms as ‘x is good’, ‘bad’, ‘beautiful’, ‘ugly’, ‘ignoble’, ‘brave’, ‘just’, ‘mischievous’, ‘malicious’, ‘worthy’, ‘honest’, ‘corrupt’, ‘disgusting’, ‘amusing’, ‘diverting’, ‘boring’ etc. – no restrictions at all on the category of x) and *directive* or *deliberative* (or *practical*) *judgements* (e.g. ‘I must Ψ ’, ‘I ought to Ψ ’, ‘it would be best, all things considered, for me to Ψ ’, etc.). It is true that between these there is an important no-man’s-land (comprising, e.g.,...general or particular statements about actions that it is ignoble or inhuman or unspeakably wicked to do or not to do). (Wiggins 1976 reprinted 1998 p95).

The question of what it is to be evaluative, and how to characterise evaluation, will be a theme throughout the thesis. Likewise whether it is possible to separate the evaluative and non-evaluative. (I use ‘non-evaluative’ here for the contrast with the evaluative; different terms are used for the non-evaluative in the literature including ‘descriptive’, but following Debbie Roberts (2013a) I find non-evaluative the most neutral characterisation.) Let me say from the beginning that I do not think it is possible to give a *definition* of the evaluative, or a clear test as to whether or not a concept is evaluative. There are lots of grey areas. With regard to any fact-value distinction or dichotomy,³ I do not think there is one, and I hope that why I think that will be fleshed out as the thesis proceeds. I might say that as a concept itself, EVALUATIVE does not classify cleanly. To start things off, I prefer to say that the notion of evaluation is a notion we recognise as being about value or disvalue, things being good to some degree or bad to some degree.

The notion of a thick concept is a concept that is held to be *evaluative* but is also more *empirically informative* and specific than ‘thin’ evaluative concepts such as good, bad, right or wrong. Thick concepts are usually introduced by way of examples. The examples of evaluations given by Wiggins, with the exception of ‘good’ and ‘bad’, are candidate thick concepts. Perhaps the most obvious suggestions are virtue concepts: KIND, CRUEL PRUDENT, COURAGEOUS, JUST, LIBERAL. Here are some more suggestions made by Simon Kirchin

MACABRE, MUNDANE, JEALOUS, DULL-WITTED, HUMANE, JEJEUNE, DEFIANT, TANTALIZING, STIMULATING, KITSCH, CONTORTED, DISARMING, SCINTILLATING, SINISTER, HONOURABLE, RIDICULOUS, GENTLE, PICARESQUE, SLOBBISH, ANTAGONISTIC, TANGLED, TRANSCENDANT, STRAIGHTFORWARD, BOLD, MARVELLOUS, TASTY, DIABOLICAL (Kirchin 2018 p 115)

The point about specificity in regard to thick concepts is that there are more things that can fall under the concept ‘bad’ than under the concepts CRUEL, JEALOUS, or SLOBBISH. That is not to say that ‘thick’ (or ‘thin’) classifies, that all concepts are either thick or not. It may be better to think of a continuum of concepts from thicker to thinner.

The term ‘thick concepts’ was introduced by Bernard Williams, although other writers had

³ Putnam (2002)

previously written about the thick.⁴ Williams wrote:

What has happened is that the theorists⁵ have brought the fact-value distinction to language rather than finding it revealed there. What they have found are a lot of those “thicker” or more specific ethical notions I have already referred to, such as *treachery* and *promise* and *brutality* and *courage*, which seem to express a union of fact and value. The way these notions are applied is determined by what the world is like (for instance, by how someone has behaved), and yet, at the same time, their application usually involves a certain valuation of the situation, of persons or actions. Moreover, they usually (although not necessarily directly) provide reasons for action. Terms of this kind certainly do not lay bare the fact-value distinction. Rather, the theorist who wants to defend the distinction has to interpret the working of these terms, and he does so by treating them as a conjunction of a factual and an evaluative element, which can in principal be separated from one another. (Williams 1985 pp 143-144)

Elaborating on his comment about thick concepts functioning as reasons, Williams continues:

They [thick ethical concepts] are characteristically related to reasons for action. If a concept of this kind applies, this often provides someone with a reason for action, though that reason need not be a decisive one and may be outweighed by other reasons...Of course, exactly what reason for action is provided, and to whom, depends on the situation, in ways that may well be governed by this and by other ethical concepts, but some general connection with action is clear enough. We may say, summarily, that such concepts are “action-guiding”.

At the same time, their application is guided by the world. A concept of this sort may be rightly or wrongly applied, and people who have acquired it can agree that it applies or fails to apply to some new situation...We can say, then, that the application of these concepts is at the same time world-guided and action-guiding.’ (Williams 1985 pp 155-156)

Here we see the example of thick concepts being used to push against the fact-value distinction. Here also is the claim that thick concepts provide reasons for action, although the claim is caveated ‘usually (although not necessarily directly)’. Williams talks particularly about ethical thick concepts, but, as Kirchin’s examples show, we can suggest thick concepts that are not ethical or moral but aesthetic or epistemological. (Again, I do not suggest that these are clear or fixed categories.)

While Williams’s assertion that ‘they usually...provide reasons for action’ has been quoted in the literature, and the basic claim acknowledged, particularly by writers who agree that thick concepts are evaluative,⁶ the mechanics of how this is supposed to work have not been extensively examined, apart from research on Williams’ own philosophy. Scholarly interest in thick concepts has lain elsewhere. Claims about reasons for action, or action guiding-ness, are passengers to the main debates about thick concepts as they take place in the literature. This

⁴ Gilbert Ryle wrote of ‘thick descriptions’ - see Kirchin (2013, 2017). Philippa Foot’s discussion in her essay ‘Moral Beliefs’ (1958) concerns examples of what would later be called thick concepts

⁵ Williams’ target is what he calls the ‘linguistic turn’ and the prescriptivism theory of RM Hare.

⁶ For example Simon Kirchin. I will return to Kirchin’s discussion of thick concepts in chapters 4 and 5.

is the research interest of the thesis question.

It is not the main purpose of this thesis to critically appraise or defend Bernard Williams's position on thick concepts, or on reasons for action. For example, Williams thought that thick ethical concepts were *prior* to thin ethical concepts; that it is human reflection on more specific notions such as CRUELTY that gives rise to more abstract notions such as 'good'. He also thought that ethical concepts are specific to local communities and cannot be objectively secured. In that sense he was an irrealist about ethical concepts. Williams also thought that reasons for action must be related to an agent's subjective motivations. Many of these issues are discussed in the thesis. The possibility that any claim that thick concepts provide reasons for action can be made only, or most naturally, by a writer who holds a number of Williams' philosophical commitments, not just the ones about thick concepts, is addressed. However, what this thesis is attempting to do is to examine how a number of claims fit together, independently of adopting Williams' positions.

In more recent literature, a number of positions on thick concepts have been developed, and as a broad schematic they can be categorised in the following way (Roberts 2013b). 'Pragmatic' views hold that evaluation is no part of the meaning of a term, or content of a concept, but must be accounted for in some other way. While there can be seemingly thick words or terms, the notion of a thick concept is highly dubious. 'Content' views hold that evaluation can be part of the meaning of a term or the content of a concept. However, there is a significant divide in content views between 'reductive' positions that claim that the evaluative and non-evaluative content can be separated by analysis, and 'non-reductive' positions that hold that there just is no way of analysing a thick concept to the point where a non-evaluative element can be discerned or specified.

I have already noted that the question of what is to be evaluative will recur throughout the thesis. Other issues about thick concepts that will be touched on are as follows:

Evaluative flexibility: although I am saying that thick concepts are evaluative and, as a quick sketch, this is about value or dis-value, good to some degree, bad to some degree, it is possible for a thick concept to signify 'good for being CONCEPT' in one instance or context and 'bad for being CONCEPT' in another. EXTRAVAGANT is an example perhaps. We may worry about being EXTRAVAGANT with money, but a Christmas pantomime without EXTRAVAGANT costumes would be the poorer for their absence. Those writers who doubt that thick concepts are evaluative as a matter of content, or that their content is irreducible, point to this flexibility in the evaluative valence as evidence in favour of their own views.

Disagreement —essential contestability: the notion of 'essentially contested concepts' was

introduced by WB Gallie (Gallie 1956).⁷ There can be concepts where two competent users of the concept can reasonably disagree whether it is instantiated in a particular case. Some writers highlight this a typical feature of evaluative concepts (for example Wiggins 1987, Roberts 2013a)

Disagreement — objectionable thick concepts: there are some concepts that we can use competently but choose to avoid, because we disagree with their evaluative stance or their evaluative point. BLASPHEMY⁸ is of course a famous example. Others suggested are LEWD and CHASTE. I may well not disapprove of language that is LEWD but if you point out an instance of LEWD language to me, I may not be able to *deny* that the language is LEWD, at least in the sense that I recognise that it is the kind of language that falls under the concept 'LEWD', even though I disagree with the evaluation involved in applying that concept. Conceivably lots of concepts are potentially 'objectionable' in this sense. Again, writers who deny that thick concepts are irreducible evaluative concepts use this difficulty in expressing disagreement as an argument in favour of their own positions.

1.1.2 Reasons for action

'Thick concepts provide reasons for action' is a very condensed claim. How does a concept, or the application of it, provide anything? Here is some of the unpacking:

- A An atomic proposition is a concept predicated of an object.
- B A fact is a true (atomic) proposition
- C A reason for action is a fact that stands in a reason relation to an agent and an action.

With this machinery in place, I will use the claim that 'thick concepts provide reasons for action' as shorthand for the claim that 'facts that are true atomic propositions, and where the predicate of the proposition is a thick concept, provide reasons for action'. In the thesis I will use the terminology 'thick-concept facts' to pick out the facts expressed by atomic propositions that involve thick concepts.

Clearly, in order to make use of thick-concept facts, I need it to be at least plausible that propositions involving thick concepts can be true, that they can be facts. It is for this reason that I spend some time looking at the debates about thick concepts and establishing that in my view some concepts can be both evaluative and irreducible or not separable into

⁷ Gallie's examples included social justice, democracy and adherence to a religion, so perhaps more matters of political philosophy than ethics. His seven marks of a contested concept included that the concept must be appraisive of value, what satisfies it must be open to modification and users recognise that the concept is contested.

⁸ Oscar Wilde ' "Blasphemous" is not a word of mine.' Testimony during his trial 03 April 1895

component parts. I have briefly indicated that some writers think that thick concepts are evaluative as a matter of content, but that the content can be analysed into its evaluative and non-evaluative components. Such positions would allow that a thick-concept proposition is capable of being truth-evaluable. However, such positions would leave open a question, in respect of a claim that thick concepts provide reasons for action, as to whether it is the non-evaluative or evaluative component that is doing the reason-giving work. If the reason-giving role of a thick concept proposition can be reduced to either the reason-giving role of its evaluative or of its non-evaluative component, then from this it would follow that there is nothing special about thick concepts in providing reasons for action: we can give an account in which they either operate like non-evaluative concepts or they operate like thinly evaluative concepts, with their 'thickness' not contributing in any distinctive way to their reason-giving role.

Do thick-concepts facts differ from non-thick concept facts (whether these are purely descriptive facts or purely evaluative facts) in providing reasons for action? On an initial look, there could be at least three versions of the claim that thick concepts provide reasons for action. These are:

- I Thick concept facts *always* provide reasons for action.
- II *Only* thick concept facts provide reasons for action.
- III Thick concept facts are *more apt* to provide reasons for action than non-thick concept facts. (This is clearly a weaker claim than I and II and 'more apt' may require explanation.)

What is meant by 'a reason for action' also requires elaboration. When we ask what reasons there are for an agent to act, for her to do something, we could be asking what reasons favour the action or make it a good or right thing to do to some degree, or we could be asking what reasons explain why she acts in that way. Can the claim that thick concepts provide reasons for action be maintained about any reason for action question or statement? If not, which avenues are most promising?

In order to explore the question of the thesis, that is, of whether there is a distinctive and interesting sense in which thick concepts provide reasons for action, and hopefully to answer it, this thesis will proceed as follows.

1.2 Twentieth century debates about non-cognitivism

Chapter two of the thesis concerns the historical debate in the literature between

cognitivists and non-cognitivists, with a view to addressing the challenge that there are no interesting, irreducibly thick concepts in the sense of concepts with irreducibly evaluative content. There are many philosophical positions that draw distinctions between the evaluative and non-evaluative and are grouped under the heading of 'non-cognitivism'. Some of the key arguments about thick concepts in the literature arose from debates between cognitivist and non-cognitivist writers about moral realism. The position I begin with in chapter two is the non-cognitivist, expressivist position of Simon Blackburn (1984), known as Projectivism. Blackburn's position holds that only descriptive, non-evaluative features of the world are real. Any sense of good or bad, right or wrong, pro or con, is an attitude *projected* on to the world. Blackburn presents his account as a Humean position. An upshot of the position is that there is a sharp underlying distinction between evaluative and non-evaluative language. Blackburn acknowledges however that in our everyday language we speak as though evaluative propositions can be true. His stance on that is not to argue that we are in error in so speaking. Rather, in an enterprise that he calls 'quasi-realism', he aims to establish that the projectivist has the right to treat such propositions as truth evaluable.

One appeal, I think, of expressivism is that much day-to-day language does involve the expression of attitudes and opinions. The current British Government is presently formed by the Labour Party. That is a matter of fact, but different people will have varying attitudes to it including both strong approval and strong disapproval. That we do express attitudes is no threat to the defender of thick concepts. All that is wanted by her is that there can be evaluative concepts and evaluative facts, not that there are no other ways to communicate approval and disapproval.

In response to Blackburn and other expressivists, writers including John McDowell and David Wiggins have pointed to evaluative concepts or evaluative properties as counterexamples. In order for the projectivist project to go through, it must be possible to identify the extension of a candidate evaluative concept, or the possession of the property, prior to any attitude being projected on to it. They argue - persuasively I think - that this can't always be done. What do all the things that are FUNNY have in common, except that they are evaluated as FUNNY? As these arguments have been influential in the literature I spend a little time explaining them.

Writing specifically about thick concepts, Simon Blackburn makes an argument that his position is more readily able to deal with evaluative flexibility than accounts that take evaluation as irreducibly part of the content of thick concepts (Blackburn 1992, 2013). As I indicated earlier, evaluative flexibility is the idea that some concepts are used to indicate a positive evaluation in some circumstances and a negative evaluation in others. This is a real issue. Unless we can accommodate evaluative flexibility in a theory, we may be forced into

acknowledging judgements that we don't agree with (objectionable concepts) or forced into accepting that two people are using the same term but referring to different concepts — in effect talking past each other. Jonathan Dancy (1998) resists Blackburn's argument on this, pointing out that 'to reject a concept one has to do more than fail to see its point; one has to know what its point is and reject it for that reason.' (Dancy 1998 p269)

Chapter two surveys this debate between Blackburn and his opponents over the reducibility – or otherwise – of thick concepts to a descriptive component plus an attitude, and concludes that Blackburn has not done enough to establish that thick concepts can be so reduced.

1.3 More recent thick concept debates

Chapter three looks at more recent literature and the development of substantive positive views on thick concepts by writers Pekka Väyrynen, Debbie Roberts and Simon Kirchin. A major theme of these writings concerns what it is for a concept *to be* evaluative. Thick evaluative concepts are contrasted with thin evaluative concepts, but then the question is raised whether any clear distinction can be drawn between these and indeed whether thin evaluation can be taken for granted.

More recent literature has distilled the 'disentangling' remarks made by McDowell and Williams into arguments about 'graspability' and shapelessness. Daniel Elstein and Thomas Hurka grant the point that, in the case of thick concepts, it may not be possible to isolate non-evaluative content without resort to some evaluative resources. However, they say, it does not follow from this that the overall pro or con evaluation associated with a concept cannot be separated from the rest of the content or what drives the extension of the term. They therefore argue for the possibility of a reductive analysis of thick concepts.

Pekka Väyrynen makes a good point that McDowell's arguments secure the shapelessness of thick concepts, but shapelessness in and of itself does not secure or entail that a concept is evaluative. After all the concept 'old' appears to be shapeless except with respect to being old, but it is not typically regarded as a paradigmatic thick concept. Väyrynen argues that only thin evaluative concepts are 'inherently' evaluative as a matter of content. What is going on with 'paradigmatic' thick concepts is that the concepts are plausibly unitary but that is because they are non-evaluative as a matter of content, and they have evaluation normally associated with them as a matter of conversational implication.

Debbie Roberts and Simon Kirchin argue that Väyrynen's position, and other reductive positions, implicitly require a claim or view that all evaluation must involve thin evaluation.

For writers who want to oppose this claim, then a different account of evaluation is indicated and they need to show what that can be. Their positions are in sympathy but different. Roberts discusses Elstein and Hurka's reduction argument, and she uses their pattern of analysis to show that, if thick concepts are held to be both irreducible and evaluative, then their being evaluative cannot be a matter of being pro or con. Thick evaluation must be something else. She develops arguments about 'irreducible thickness' to show that to whatever grounding level we pursue the matter, the particular in view cannot be shown to be unimpeachably non-evaluative.

Simon Kirchin calls his position the 'liberal' view of thick concepts. He contrasts this with a 'conservative' view, in which evaluation can only be a matter of being pro or con. Kirchin thinks that thick concepts are evaluative. Thick concepts are complex and have many functions. A thick concept will have some instantiations where either a pro (positive) or con (negative) judgement is conveyed, but not every instantiation will be pro or con and that's fine. The upshot, I think, is that almost any concept can be a thick evaluative concept in Kirchin's view. While Kirchin's position is the one that I have most sympathy with, it is not a position that emphasises the reason-givingness of thick concepts.

It is at this point then that I depart from the focus of the recent literature. Having endorsed, in chapter three, the arguments of Kirchin and Roberts to the effect that thick concepts are irreducibly evaluative, I turn in the remainder of the thesis to the question whether, given the existence of irreducibly evaluative thick concepts, these concepts play a distinctive role in providing reasons for action.

1.4 Reasons for action

Aside from discussions of Williams' philosophical ideas, much more has been written in the literature about the evaluative (or otherwise) nature of thick concepts, than about the suggestion that thick concepts have some role in providing reasons for actions. Chapter four hopes to contribute to redressing this imbalance in the recent discussion of thick concepts, by addressing their reason-giving role. Turning, then, to reasons for action, I discuss different understandings of what a reason for action is, and the apparatus by which a concept figures in a reason. I set out an account of what constitutes an action following Donald Davidson and Jennifer Hornsby; not every doing of an individual, still less every event or happening, is an action.

Although the connection between thick concepts and reasons for action is a claim put forward by Williams, writers find a tension, indeed inconsistency, in his positions on these two subjects. This is partly owing to Williams' particular view of what it takes for a fact to be

a normative reason for action. Ulrike Heuer (2012) explores this topic and formulates different ways that Williams' thick concepts claim could cash out in his terms, all of them problematic. Heuer concludes that the two positions are not consistent.

If we stand completely aside from Bernard Williams' philosophical framework, how is a suggestion that thick concepts provide reasons for action to be understood? Is it that a thick evaluative concept provides a reason for action whenever its application arises? Or is it that thick evaluative concepts have some capacity to provide reasons that non-evaluative concepts lack? Secondly, there are different aspects of what it is for something to be a reason for action, and which are now current in the literature. To put it another way, there are different questions about agents and actions to which a 'reason' is the answer (Dancy 2000). I set these out following a discussion by Maria Alvarez (2017). In brief, normative reasons justify or favour actions. Some writers think that a normative reason must yet have a link to the agent's motivations to be a reason. (Williams' view of reasons for action is located here.) Others think, on the contrary, that normative reasons stand on their own merits. Reasons for action can also specify the motivation of an agent or explain the agent's action. In some cases motivation and explanation may be the same thing, but in other cases they might not be.

One reading of the claim that thick concepts provide reasons for action is that there is some understanding of what it is to be a reason for action according to which, where a thick concept applies, there is always a reason to act. In other words, the application of a thick concept will always stand in a reason relation to an action, even where no action was previously under consideration. I shall label this the 'Always' thesis. A second reading of the claim could be that there is some understanding of what it is to be a reason for action in which a thick concept, where it applies, provides a reason for action or is 'action-guiding', whereas a non-evaluative concept does not provide a reason for action or is not action guiding. I will label this the 'Only' thesis.

In sections 4.4 to 4.8 I will go through the possibilities of understanding what reasons for action are as I present them in section 4.3, and I will discuss whether either reading of the claim about thick concepts — the Always thesis and the Only thesis — has any purchase. There is very little to be said in favour of the Only thesis. The Always thesis requires some argument to support it, and perhaps some restriction on its application, but by the end of this chapter my thought that an action is CRUEL is a reason for me not to perform that action is still in play.

1.5 Thick concepts and reasons for action - further arguments

In chapter five, I carry forward two strands of discussion from chapter four. The first is to

maintain that thick-concept facts are normative reasons for action in an externalist sense, where there is an action in view to which they are relevant. Thick-concept facts are always reasons, although they may be (easily) defeasible. The second topic for discussion is the suggestion that what is going on, in the most obvious examples of when thick concepts provide reasons for action - 'I will not kick the cat because kicking the cat is CRUEL' - is that the action itself is being evaluated. I will link this to a discussion about virtue concepts and whether it is feasible to make the same reasons-for-action claim about all thick concepts. CRUEL and KIND are virtue concepts. TACTFUL and MACABRE are thick concepts but not virtue concepts. Are CRUEL, KIND, TACTFUL and MACABRE all reason-giving in the same way and to the same extent?

The main challenge as I see it to the claim that thick-concept facts are normative reasons for action, is to say that they are superfluous and provide no additional reasons to the other reasons that there are. The thought is that if a thick-concept fact provides a normative reason at all, it is at best a 'higher-level' reason, or an abstraction from one or more of the reasons that there are. The normative reasons that there are in any case are facts specific to the situation that may or may not have evaluative content. If a reason for my saying something complimentary to Thomas is that my so saying will be GENEROUS, then, according to this picture, there will be something that underpins the fact that the action is GENEROUS — perhaps Thomas has low self-confidence, or has had an upset — and the underpinning fact is the actual normative reason in this case. The fact that the action is GENEROUS is not an addition to the reasons for acting, and I can do the right thing without even having the concept of GENEROSITY. I recognise that he is need of a confidence boost that I am able to give, and I respond accordingly by offering the compliment. The fact that so acting is GENEROUS may add to the reasons in an explanatory sense, by saying something about the reasons why an agent acts, but that is a separate matter.

Why would I want to resist this? Well one plausible thought is that thin-concept facts or evaluations, such as *x* is right or *x* is good, are not normative reasons in the sense that I am understanding that term (Dancy 2000). They have the appropriate form — they are facts — but can only be an outcome or verdict of all the other considerations in favour of or against an action. The thought is that there must be more to say in favour of an action than just to say that it is a good thing to do. There must be something that supports the judgement or grounds it. To say that thick concept facts are superfluous as normative reasons, that they add nothing to the reasons, is to put them on the same footing as thinly evaluative concepts: that is that if an action satisfies a thick concept there must be something further to say about it, in virtue of which the thick concept applies, and which favours the action. Those of us who think that thick concepts are evaluative and irreducible concepts and therefore interesting should find this result somewhat troubling.

In chapter four when initially outlining the objection of thick concepts add nothing to the reasons, I suggested some examples where a thick concept applies and yet it is not so clear that an underpinning or grounding fact is available. Suppose I contemplate painting a wall in my house in two different colours and then realise that the combination is GARISH. Is there anything more to say than that? If there are any plausible examples here, then at least there is a counter-argument to any blanket denial that thick concepts provide (normative) reasons for action.

I then go on to consider whether the charge that thick concepts add nothing to the reasons amounts to irrelevance, that thick-concept facts are not relevant as reasons. I put forward the discussion of Jackson and Pettit (1990) about programme explanation, which draws a distinction between relevance and efficacy. Jackson and Pettit argue that properties can be causally relevant without being causally efficacious. I claim that analogously thick-concept facts are relevant in providing normative reasons for or against an action although there may never be an instance where there is a thick concept fact and no other fact that speaks in favour of or against an action. This again is enough to show that there is a response to the challenge that thick concept facts are superfluous as reasons. To support this, I will show by examples that concepts are important to recognising a similarity of reasons, and that they function where there are unavoidable constraints on information.

In section 5.4 I return again to Bernard Williams, and some observations he makes in the essay 'Acting as the Virtuous Person Acts' (Williams 2006). Williams remarks that thick concepts are particularly important as negative considerations that prevent an agent from performing an action that she otherwise has reason to do. Imagine that the cat has brought a baby rabbit into the house and is subjecting it to a tortured killing, and so some intervention is needed. Kicking the cat might startle it into dropping the rabbit, (and also express my feelings in the moment), but nonetheless it is CRUEL and that persuades me to find some other means of rescuing the rabbit. I think that what is happening in this case is that the action or prospective action is being evaluated. If I am right about this, then this point may do one of two things. On the one hand it may provide an account of why I have a strong intuition that facts such as 'kicking the cat is CRUEL' are normative reasons, but my intuition is mistaken. 'Kicking the cat is CRUEL' is an evaluation and therefore apt to be an explanatory reason, a reason that explains what I, the agent, find salient in the situation. On the other hand it may support my claim that thick-concept facts provide normative reasons for action. Perhaps though, I have to restrict the claim to instances where the application of a thick concept is to the action or prospective action rather than any other feature of the situation, (for example the agent, or the beneficiary of the action).

CRUEL is a virtue concept, as are KIND, GENEROUS, COURAGEOUS, JUST, PRUDENT. These are the most paradigmatic thick concepts, in that all writers accept them as examples for discussion. Is there a possibility that writers who would otherwise reject the reasons for action claim can accept it for a limited range of concepts? In section 5.5 I begin to look at what is distinctive about virtue concepts. I start with briefly explaining the notion of the virtuous person and then look at a discussion by Williams go how other considerations — including of course the application of thick concepts — relate to the virtues.

In section 5.5 I extend the discussion to the claim that, where a virtue concept applies it will always be a reason, because virtues are ‘ultimate reasons’ (Crisp 2000). It turns out that virtue concepts are the examples of thick concepts where my claim that thick concepts provide reasons for action is most vulnerable to the objection that they do not add to the reasons.

1.6 Conclusion

Chapter six summarises what has been said and concludes.

Chapter 2: Twentieth century debates: cognitivism, non-cognitivism and thick concepts

2.1 Introduction

I said in the introduction that, in order to argue for the claim that thick concepts provide reasons for action, I need to make use of or rely on the following apparatus:

- A An atomic proposition is a concept predicated of an object.
- B A fact is a true proposition (where this is either atomic or a complex proposition whose building blocks are atomic propositional functions)
- C A reason for action is a fact that stands in a reason relation to an agent and an action.

In this thesis I am not going to argue for all of these claims. I do however need to deal with one issue, which is central to the debate on thick concepts. This issue is the thought that where the concept contained within a proposition is an evaluative concept - i.e. one of our paradigmatic thick concepts - then the proposition is not capable of being true, where true means true of the world in a realist sense. Consequently, if we take this scepticism about evaluative facts seriously, it is not possible to have a 'thick-concept fact' and it follows therefore that a thick concept cannot provide a reason for action. Such a position might be held by those who wish to separate out the domains of fact and value, and to combine realism about the realm of fact with a form of anti-realism about the realm of value. On such a position, insofar as thick concepts are evaluative, propositions involving thick concepts cannot be true.

Realism is not easy to formulate (Wright 2003) but I accept Crispin Wright's sketch that it is a view that takes it that how matters stand in the world is settled independently of actual people's beliefs, while holding that it is yet possible for people to capture or know the truth (Wright 2003 p11-12). What I want to be able to say is that 'kicking the cat is CRUEL' can be true just as well as 'the ink is black'. I need both this, and the claim that thick concepts are not analysable or separable into a non-evaluative core and an expression of an attitude, to support my claim that thick concepts can function as reasons.

This chapter of the thesis then concerns the historical debate in the literature between cognitivists and non-cognitivists, with a view to addressing the challenge that there are no

interesting, irreducibly thick concepts in the sense of concepts with irreducibly evaluative content. There are many philosophical positions that draw distinctions between the evaluative and non-evaluative and are grouped under the heading of 'non-cognitivism'. Some of the key arguments about thick concepts in the literature arose from debates between cognitivist and non-cognitivist writers about moral realism. The position I begin with in this chapter is the non-cognitivist, expressivist position of Simon Blackburn (1984), known as Projectivism. Blackburn's position holds that only descriptive, non-evaluative features of the world are real. Any sense of good or bad, right or wrong, pro or con, is an attitude projected on to the world. Blackburn presents his account as a Humean position. An upshot of the position is that there is a sharp underlying distinction between evaluative and non-evaluative language. Blackburn acknowledges however that in our everyday language we speak as though evaluative propositions can be true. His stance on that is not to argue that we are in error in so speaking. Rather, in an enterprise that he calls 'quasi-realism', he aims to establish that the projectivist has the right to treat such propositions as truth evaluable.

In response to Blackburn and other expressivists, writers including John McDowell (1981, 1987) and David Wiggins (1987) have pointed to evaluative concepts or evaluative properties as counterexamples. In order for the projectivist project to go through, it must be possible to identify the extension of a candidate evaluative concept, or the possession of the property, prior to any attitude being projected on to it. They argue - persuasively I think - that this can't always be done. What do all the things that are FUNNY have in common, except that they are judged to be FUNNY? As these arguments have been influential in the literature I spend a little time explaining them in this chapter.

Writing specifically about thick concepts, Simon Blackburn makes an argument that his position is more readily able to deal with evaluative flexibility than accounts that take evaluation as irreducibly part of the content of thick concepts (Blackburn 1992, 2013). As I indicated in the introductory chapter, evaluative flexibility is the notion that some concepts are used to indicate a positive evaluation in some circumstances and a negative evaluation in others. This is a real issue. Unless we can accommodate evaluative flexibility, we may be forced into accepting judgements that we don't agree with because we can't deny that the concept applies (objectionable concepts) or forced into believing that two people are using the same term but referring to different concepts — in effect talking past each other. Jonathan Dancy (1998) resists Blackburn's argument on this, pointing out that 'to reject a concept one has to do more than fail to see its point; one has to know what its point is and reject it for that reason.' (Dancy 1998 p269)

2.2 Expressivism, projectivism and quasi-realism

The projectivist theory

According to expressivist theories of language, there is a clear distinction between evaluative and non-evaluative discourse. The theoretical account to be given of one is not the same as the theoretical account to be given of the other. To illustrate, here are two example sentences:

- A The ink is black.
- B The comment is malicious.

On the face of it, these two sentences have structural similarities. They both express propositions. They might both be uttered in a context in order to make an assertion. We can imagine someone holding a piece of paper on which there is writing, and saying both of these sentences to her companion, talking about the writing on the paper. Her companion, we can imagine, takes her to be making some statements, asserting some facts, saying things which are true, about what she is holding in her hand.

Nonetheless, on an expressivist account of language, there is a clear difference between an utterance of 'the ink is black' and an utterance of 'the comment is malicious'. The first *could* be an assertion, and could be a statement of fact if this is how the world is, if the ink is indeed black. The second is primarily an expression of an attitude. What is uttered cannot be true or false, and there is no fact of the matter to be had about whether a comment is malicious. According to an expressivist account, evaluative discourse is irrealist or anti-realist. Irrealism and anti-realism stand in contrast to realism.

The position of the British philosopher Simon Blackburn, called 'projectivism', is a form of expressivism. Projectivism is the thesis that some of the features that we 'take' to be in the world are our own creation. There is a world proper, a reality, which 'impinges' on human agents. Humans react to the value-free world and form attitudes (amongst other things). Humans then 'spread on' the world their reactions and talk and think as though the world really contains states of affairs answering to them (Blackburn 1981). It follows that projectivism is a causal theory. We can have beliefs as inputs, and our moral sensibilities (sometimes) process those beliefs and produce attitudes as outputs (Blackburn 1984). Natural or descriptive properties are genuine and real, but 'spread' properties are not real.

In other writings Blackburn uses the metaphor of 'direction of fit' (Blackburn 1998, 2013).

Sentences that, when uttered, primarily express attitudes would have the world conform to them, whereas descriptive (or 'representational') sentences must conform to the world; that is, attitudes have a direction of fit akin to desires and in contrast to beliefs, which must fit the world. Someone who accepts the direction of fit contrast deployed in this way, and accepts that the right way to theorise about a sentence which expresses an attitude when uttered is to theorise about a commitment that is not representing the world in any way, is a projectivist.

It is important then to understanding projectivism to consider what attitudes are taken to be in the projectivist's account. Attitudes are 'states of mind whose function is not to represent anything about the world' (Blackburn 1998 p504). In essence, an attitude is a mental state without propositional content. We have a variety of 'commitments' and commitments are mental states. Some commitments, such as beliefs, can have propositional content. If p is a proposition, I can believe that p , fear that p , hope that p ; these are all propositional attitudes towards the proposition (or thinkable) p . A non-propositional attitude (if I may call it that) does not operate on a thinkable.

To Blackburn, the expression of an attitude in its simplest form is in direct speech, and is a "Boo!" or a "Hooray!". This is of significance; it appears that attitudes can be pro or con (or in some instances the attitude is toleration). So, for example, I can express my positive attitude towards a person by saying her name in an enthusiastic tone; "Sarah!" And the utterance with this tone is equivalent to my saying "Sarah. Hooray!". Blackburn thinks that propositional expression arose principally because of the communicative need to challenge, explain and develop our *past* utterances (Blackburn 1992). It is much easier to work with sentences such as 'Chloe said that Sarah is great' than with sentences such as 'Chloe said "Sarah!" or 'Chloe said "Sarah" with a really positive tone in her voice.'

Why might a projectivist picture seem attractive or plausible? Expressivism remains a live project in philosophy, for example in the work of Huw Price (Price 2013). Here are some suggestions to explain the attraction. The first reason is that projectivists are realists about some things and irrealists about others, and this seems like a good place to be. A projectivist can endorse a commitment that I believe that there is a computer on the desk because there is a computer on the desk, without having to make any realist commitments about anything other than the physical world.

A second reason is that we do, in an everyday sense, accept or support the idea that there are facts and there are opinions and many people think that facts and opinions (attitudes) are different in character. If I say that the current British government is Labour, that's a fact (at the time of writing). If I roll my eyes and sigh while saying it, then you gather that my attitude

towards that fact is negative (or if I smile happily you understand a positive attitude). Even if I verbalise - “The current government is an embarrassment!” - and even if you agree with me, we are inclined to think of the negative evaluation as my attitude, rather than part of the realm of objective fact. Another way of putting this is to say that, again in an everyday sense, we suspect there is a difference between objectivity and subjectivity. How could it be that ‘the current British government is Labour’ and ‘the current British government is an embarrassment’ report or describe states of affairs in just the same way? Blackburn’s projectivism looks like a route to make good on this intuition, and I think any expressivist or non-cognitivist position has an initial attraction in this way.

Thirdly, we might think that there are more things we can say, more potentially available sentences to use, than there are real physical states of the world. After all, there are many, many things that we can say about the Labour government, but still there is just one government. Metaphorically speaking, we might say that there are layers of language, and some things we can say are attitudes towards, or inferences from, prior commitments. Huw Price also argues in support of this point — that there is more language than world — which he says gives rise to the placement problem — the problem of finding a place for some of the more esoteric things we seem to speak of, including evaluative features, in a world of bare ‘facts’. (Price 2013)

Blackburn himself makes a related point when he talks about the economy of projectivism, and its advantage in not overcomplicating the world and asking more of it than we know is there (Blackburn 1984 p182). We typically experience physical properties of the world through our perceptual faculties. We cannot see, hear, feel, taste or touch values. Therefore, if realists claim that values are genuine properties, they must also suggest or postulate a mechanism – presumably some kind of intuition – by which we become aware of them.

A further argument that Blackburn makes in favour of projectivism is that the alternative, moral realism, has to include in its account a claim of supervenience that turns out to be deeply problematic (Blackburn 1973). If evaluative properties and genuine or descriptive properties are distinct, a realist will have to hold that the evaluative supervenes on the descriptive. The nub of the counter-arguments from cognitivists is that, if there is supervenience, there is no reason to think the supervenience base is non-evaluative properties (McDowell 1981). I will say more about the counter arguments in section 2.2

Blackburn supports projectivism with his broader framework, which he calls ‘quasi-realism’. The aim of quasi-realism is to explain the realistic nature of our discourse and to show that projectivism has all the same resources for saying things as realism. Our discourse

is ‘realistic’ because we *talk* as though evaluative properties are real and evaluative propositions are capable of truth and falsity. Quasi-realism aims to establish that expressions of attitude have the right to propositional form, and that notions of truth and falsity can legitimately be applied to expressive sentences, or sentences expressing attitudes, while still maintaining the expressivist or projectivist thesis. According to quasi-realism, how we speak does not require a systematic overhaul; projectivism is not a revisionist thesis by whose lights we are misguided or mistaken in what we take ourselves to be doing when we think and speak.

Blackburn’s strategy for establishing quasi-realism is to show that there is no aspect of discourse that cannot be explained on quasi-realist terms. The success of this strategy will be shown where there is a real difficulty in defining a debate between realists and anti-realists, anti-realists being in this instance projectivists (Blackburn 1980). At the end of this section, I will note two objections that Blackburn foresees and believes he can answer. Before that though, I will briefly note the problem that thick concepts present to expressivist theories, and then I will outline the principal objection to expressivism in the literature, the Frege-Geach problem.

For expressivist theories to hold or be plausible, it must be possible to distinguish between evaluative language and non-evaluative language on a case-by-case basis.⁹ This is why thick

⁹ Huw Price sees this issue, which he calls the bifurcation thesis (Price 2013). His solution is to suggest that *all* language can be theorised as expressivist. Price starts with the question of how it is that we can have all sorts of statements and assertions, that don’t obviously correspond or match with anything in this real physical world. Moral facts are one example of this, but so are mathematical facts. Price calls this failure of matching or correspondence the ‘placement problem’.

There are two conceptions of the origin of the placement problem. One is to start with objects – the physical world – and say how do these generate language. The other is to start with the linguistic or psychological data and ask how it relates to the real world. Any kind of non-cognitivism has to take this linguistic conception as the place to start. Non-cognitivism therefore tends to endorse a bifurcation of language; between descriptive assertions that are genuine facts about the world, i.e. they correspond to something in the real physical world, and assertions that don’t match the real world.

The first difficulty that the linguistic conception encounters is how does talk about language become talk about things? In Price’s terms, the semantic ladder doesn’t go all the way down. At some point there is a change of focus from language to things, which we cannot help ourselves to or take for granted.

Price’s project is to undermine the account of these descriptive assertions or facts, but then to say that this doesn’t matter because we can expand the scope of the other account, the non-factual assertions, to cover everything. This is Global Expressivism.

‘the bifurcation thesis, and in particular the belief that some claims are genuinely descriptive, play no role at all in the positive story, in the case of the commitments the expressivist regards as not genuinely descriptive. In other words, the expressivist’s positive alternative to the matching model doesn’t depend on the claim that the matching model is ever a useful model of the relation between natural language and the natural world. So there’s no evident barrier to abandoning the matching model altogether and endorsing global expressivism. This is the view that I want to recommend.’ (2012 30)

terms and concepts feature in the literature on the debate between cognitivists and non-cognitivists. The thought is this: if expressivists think that some of our sentences are used to express facts – ‘The table is square’ – and some are used to express attitudes – ‘The decision is good’ – what can they say about language that appears to be a combination of both – ‘the table is elegant’ – ‘the decision is brave’. These latter sentences, if uttered, would appear to express positive attitudes, but on the other hand not everything that is evaluated positively is elegant or brave. There is something more specific or restricted about them, some way in which they must be ‘true of the world’; in Bernard Williams’ phraseology, they are ‘world-guided.’ I will come back to the challenge raised by thick concepts for Blackburn’s projectivism, but first I will outline another main challenge to the projectivist picture, the so-called Frege-Geach problem.

The Frege-Geach problem and other objections

The Frege-Geach¹⁰ problem for an expressivist account of language is to explain how a sentence such as ‘It is wrong to tell lies’, can mean the same in an unasserted context as in an assertion, given that its meaning in the assertion is supposedly due to its being the expression of an attitude. An example of an unasserted context could be in the antecedent or consequent of a conditional: ‘If it is wrong to tell lies, it is wrong to get your little brother to tell lies.’ Such a conditional is no problem to a realist: it can be deployed straightforwardly in a valid argument with premises and a conclusion that follows logically.

P1: It is wrong to tell lies.

P2: If it is wrong to tell lies, it is wrong to get your little brother to tell lies.

Therefore:

C: It is wrong to get your little brother to tell lies.

For an expressivist, in this conditional example, there is no commitment to an attitude on little brother and lying in either the antecedent or consequent of P2. Yet P1 *does* express a commitment to a negative attitude to telling lies, when taken by itself rather than occurring as the antecedent of P2. However, if the sentence P1 doesn’t mean the same thing when it occurs by itself as it does when it occurs embedded as the antecedent of P2, (which it would appear not to, given that an attitude of disapproval to lying is involved in P1 but not in the conditional P2) this argument commits a fallacy of equivocation. ‘[I]t is a nice sharp problem’, and may seem to cut out any chance of a successful expressive theory. (Blackburn 1984 p 190).

¹⁰ The problem is named after philosopher Peter Geach’s development of a point originally made by Gottlieb Frege.

Blackburn, however, has a response to both of the two issues he sees in the problem. The first issue is to consider what are we up to when we make remarks that are sentences embedded in contexts where they are not asserted: what is our motivation or purpose? Taking the 'little brother lying' conditional as an example, what we are doing when we endorse such a conditional is linking two commitments 'it is wrong to tell lies' and 'it is wrong to get your little brother to tell lies'. Then we are following through the implications and seeing that an admirable sensibility pairs an attitude of disapproval towards the one with an attitude of disapproval towards the other, in order to have a coherent set of attitudes. On this account the whole conditional is an evaluation, and if we endorse this evaluation *and* if we also disapprove of lying, then it follows that we should disapprove of getting one's little brother to tell lies. A different sensibility might pair different attitudes – it is wrong to tell lies but it is not wrong to get your little brother to tell lies – but it would still be a moral point of view. (To give this greater intuitive plausibility, there are lots of examples of "double standards" we could think of in terms of an activity's acceptability for different sections of society. 'It's permissible for boys to fight sometimes, but it's never permissible for girls to fight.')

The second issue in the problem of sentences in contexts where they are not asserted is to give a workable semantics, to show why we use this sentence ('it is wrong to lie') to say what we want to say. This is sketched in Blackburn 1984 and developed further in Blackburn 1998. The 1984 version goes as follows. It proposes a language 'Eex' without evaluative predicates. 'Eex' has 'hooray!' and 'boo!' operators; H!, B!. H! and B! operate on descriptions shown in brackets () , and attitudes are shown inside bars |H! (X)|. Attitudes are shown coupled by semi-colons ; . So the lying example becomes:

B! (lying)
H! (|B! (lying)|; |B! (getting little brother to lie)|)

In longhand (my words not Blackburn's) this is 'Boo to lying. Hurray to combining Boo to lying and Boo to getting little brother to lie'.

Now, if the language Eex wants a way of signalling that not to hold the consequential commitment would be a logical mistake – because it would involve a clash of attitudes and miss the point of why we evaluate -, one route would be to invent predicates answering to attitudes and then use them like descriptive predicates – in order to debate judgements and so on. That is 'our use of indirect contexts...proves us to have adopted a form of expression adequate to our needs. This is what is meant by 'projecting' attitudes onto the world.' Blackburn 1984 (p195).

Some have argued that this ‘clash of attitudes’ falls far short of logical inconsistency and only points to a moral failing. Nonetheless, even if Blackburn’s semantics cannot mimic the logical form of the realist’s modus ponens, there is an intuitive appeal to seeing the difficulty as a misalignment of attitudes.

A second problem that Blackburn sees for the projectivist and quasi-realist is the possibility of being wrong. Suppose I believe something because it is a consequence or output of a theory that I have total confidence in; for example that the gas in the atmosphere of a distant planet is carbon dioxide, or, in a moral case, that capital punishment is always wrong. I might still accept that there is a theoretical possibility that I am wrong in my belief, that something might disprove it. But what could this something be? The only candidates might seem to be empirical facts; a space probe goes to the planet and finds helium, execution is the right outcome in a case of extraordinary evil. In this way, facts are transcendent, they trump theory. However, while both realists and quasi realists could agree in the case of the gas in the atmosphere in the distant planet, only a realist could have the thought that fact trumps theory in an ethical case because only the moral realist accepts that there can be a moral fact of the matter.

Blackburn believes that the quasi-realist can reply that all that is needed to accommodate the possibility of being wrong is the idea of improvement in our position or stance. Of course, the projectivist has no problem with the space-probe example, because this is a physical world belief. In a moral case, willingness to empathise with other people, to see that moral sensibilities differ and to admire people with developed sensibilities (for example) will be quite adequate to give the idea “I might improve.”

Simon Blackburn has written about thick concepts and I turn to his discussion on the topic in section 2.3. Before that, in section 2.2, I will discuss some objections to Blackburn’s theory from John McDowell. These objections have become influential in the thick concepts debate, so it is helpful to turn to them at this point.

2.3 Objections to non-cognitivism and expressivism

John McDowell argues against the non-cognitivist and projectivist accounts of language described in the sections above, including in the two essays ‘Non-Cognitivism and Rule-Following’ (McDowell 1981) and ‘Projectivism and Truth in Ethics’ (McDowell 1987). In the former paper, McDowell presses non-cognitivists to *explain* what sort of propositions count as descriptions of reality, as distinct from propositions whose uses are ascriptive of value, rather than taking for granted that unimpeachably descriptive language must be available. In

the latter paper, and referring specifically to Simon Blackburn's position, McDowell argues that the projectivist account of value language is also flawed in that it regards itself as the *only* alternative to a naive intuitionism about value and reality. In other words, McDowell disputes both sides of the non-cognitivist account of language, the account of descriptive reality and the account of evaluative language. McDowell's own position on language is one that encompasses the application of all forms of concepts and is (in his view) inescapably anthropocentric.

Non-Cognitivism and Rule Following

There are two aspects of what non-cognitivists seem to need from descriptive language. One is the possibility of saying how things really are, how they are in themselves, independent of the point of view of the observer. The second is that value experience tracks this independent reality with some sort of consistency. Absent this reliable tracking, a non-cognitivist cannot claim that the repeated application of the same value concept demonstrates 'going on in the same way'; attitudes may be merely 'sounding off', and their resulting ascriptions of value arbitrary. Non-cognitivists, McDowell complains, just take it for granted that there must be some underlying reality that can be tracked.

One way into the non-cognitivist distinction between reality and value might be to accept an eighteenth-century philosophy of mind, which says that sentiments and passions just are a completely different sort of thing than real physical stuff. A second way, and one that might appeal to those wary of that eighteenth century philosophy of mind picture, says McDowell, is to pursue an analogy with colours and the distinction between primary and secondary qualities.¹¹ When a human observer perceives an object as red, the redness is part of the perception of the object, or a manifestation of the object. It is a secondary quality and a consequence of the observer's particular human faculties of perception being brought to bear on the physical primary qualities of the object. Allowing for light conditions and error, and so on, human observers will consistently perceive the redness of the object. Notwithstanding the lack of any obvious analogy with the faculty of perception, might something similar be going on in the case of value?

For this theory to work it must be possible to disentangle the real (or primary) from the secondary, to be able to distinguish the real features that are being tracked when we attribute secondary qualities to objects. This is far from straightforward. Here is the passage from McDowell:

Now it seems reasonable to be sceptical about whether the disentangling manoeuvre here

¹¹ See also McDowell (1985) 'Values and Secondary Qualities'

envisaged can always be effected; specifically about whether, corresponding to any value concept, one can always isolate a genuine feature of the world — by the appropriate standard of genuineness: that is, a feature that is there anyway, independently of anyone's value experience being as it is — to be that to which competent users of the concept are to be regarded as responding when they use it: that which is left in the world when one peels off the reflection of the appropriate attitude

Consider, for instance, a specific conception of some moral virtue: the conception current in a reasonably cohesive moral community. If the disentangling manoeuvre is always possible, that implies that the extension of the associated term, as it would be used by someone who belonged to the community, could be mastered independently of the special concerns that, in the community, would show themselves in admiration or emulation of actions seen as falling under the concept. That is: one could know which actions the term would be applied to, so that one would be able to predict applications and withholdings of it in new cases — not merely without oneself sharing the community's admiration (there need be no difficulty about that), but without even embarking on an attempt to make sense of their admiration. That would be an attempt to comprehend their special perspective; whereas, according to the position I am considering, the genuine feature to which the term is applied should be graspable without benefit of understanding the special perspective, since sensitivity to it is singled out as an independent ingredient in a purported explanation of why occupants of the perspective see things as they do. But is it at all plausible that this singling out can always be brought off? (McDowell 1981 pp 201-202)

McDowell goes on to say that a claim of supervenience does not secure the distinction. It is not sufficient that value properties supervene on some other physical properties or features, and therefore that if the supervening level changes there are changes in the level supervened on. For consistent tracking to be possible, the items that fall under the extension at the supervening level must have some identifiable shape or kind at the *supervened on* level, so that they can be recognised and tracked. It is very plausible that this will not always be the case - for example the *only* thing funny things have in common may be that they are *funny*. This line of argument from McDowell has been very influential in the thick concepts debate. It introduces the possibility of 'shapelessness', that for some concepts the items that fall under them have no feature in common other than that they fall under that concepts, and so those items as a group are shapeless with respect to any other feature. Shapelessness will be discussed in more detail in chapter three.

In a further line of argument, McDowell suggests that an assumption that it is possible to make judgements about a reality that is independent of any observer rests on some misconceptions.

A succession of judgements or utterances, to be intelligible as applications of a single concept to different objects, must belong to a practice of going on *doing the same thing*. We tend to be tempted to a picture of what that amounts to, on the following lines. What counts as doing the same thing, within the practice in question, is fixed by its rules. The rules mark out rails along

which correct activity within the practice must run. These rails are there anyway, independently of the responses and reactions a propensity to which one acquires when one learns the practice itself;...

The picture comes in two versions. In one, the rules can be formulated, as a codification of the practice in independently accessible terms...[in] the other version of the picture...we appeal to grasp of a universal, conceiving this as a mechanism of an analogous sort' (McDowell 1981 reprinted in McDowell 1998 pp 203-204)

As Wittgenstein showed in his remarks about following a rule, even our confidence that someone is correctly extending a pattern of numbers by adding two each time, and that she will continue to do so, is based on our shared form of life and understanding about what it is to 'add two', and not the observation of some systematic response to the world as it is. We may *hypothesise* that the adder of twos has a mental engagement with something external and independent of her in doing what she is doing, but this is no basis for our confidence that she will continue to add two, as we only have a fragment of the pattern (of the results of her activity) to judge from.

The claim that McDowell makes from these arguments is that we have no access to a point of view that is independent of human faculties, and therefore no basis on which to *assume* we can distinguish the application of non-evaluative concepts from the application of any others. Blackburn, and other expressivists, do not have the right to assume that a separation of language is possible, unless they can argue for that separation independently.

Before moving on from 'Non-Cognitivism and Rule Following', I want to note some remarks that McDowell makes about the relationship of the non-cognitivist project and reasons for action.

'What I have in mind is an argument for non-cognitivism that goes back at least to Hume...It has two premises. The first is to the effect that ascriptions of moral value are action-guiding, in something like this sense: someone who accepts such an ascription may (depending on his opportunities for action) *eo ipso* have a reason for acting in a certain way, independently of anything else being true about him. The second premise is this: to cite a cognitive propositional attitude — an attitude whose content is expressed by the sort of proposition for which acceptability consists in truth - is to give at most a partial specification of a reason for acting: to be fully explicit, one would need to add a mention of something non-cognitive, a state of the will or a volitional event. Clearly, it would follow that ascriptions of value, however acceptable, can be at most in part descriptive of the world. ' (McDowell 1981 pp 213-4)

What McDowell is saying here is the following: if you think that any belief that a proposition is true cannot on its own be sufficient as a reason for action, and you think that an ascription of value, the content of a belief that *p*, where *p* is an evaluative proposition, *does* provide a reason for action, then you have to think that there is more going on in a belief

that p than just a belief that something is true. McDowell complains that people find the second premise obvious, reflecting a misconception of the sort of explanation a reason explanation is, and indeed a 'quasi-hydraulic' view of reasons. In other words, a conception of what a reason is can drive the claim that evaluative language is separable into two parts.

I will return to this in chapter four. The inference that I want to make is that, to be open minded about what a reason is, it helps to think that evaluative concepts can be irreducible just like other concepts.

Projection and Truth in Ethics

In this essay, McDowell ascribes to Simon Blackburn the view that any position of moral realism must be of a blunt intuitionistic variety. A person who holds, say, that a proposition such as 'action x is unkind' can be *true*, can only be claiming that the value 'unkind' can be discerned by some blunt application of human faculties. There are two clear disadvantages to such a position. One is that the 'right to truth' is unearned - the intuitionist realist is simply positing that truth is available. The second is that such a position leaves no room for improvement in ethical stances; it does not appear possible to alter or develop the judgement that 'action x is unkind' or indeed, where there are two unkind actions, to reason which is better or worse.

In contrast to this, Blackburn's projectivism holds that evaluative judgements are projections of attitudes onto the world. These attitudes are shaped by ethical sensibilities that can be shaped, reasoned about, criticised and developed. The upshot is that the projectivist has both earned the right to talk about the 'new creation' of projected attitudes as though it is true of the world - quasi-realism - and has, in his own lights, no difficulty in allowing for the possibility of ethical development.

McDowell's objection to this picture has two main points. One point is that Blackburn appears to think that, because intuitionistic realism, a view that we perceive values in the world as we perceive cats and trees, is unappealing (and McDowell agrees with that), projectivism must be the only satisfactory alternative. This in and of itself must be argued for, and Blackburn does not do this. The second point is that, if the projectivist picture is correct, in the forming of an evaluation, an evaluative judgement, the attitude projected must take explanatory priority. The judgement that 'action x is unkind' is *explained* by the projected attitude. In contrast, intuitionism would have the real feature of the world as prior in explaining the evaluative judgement. The difficulty for the projectivist is that there is nothing to secure the selection of the appropriate attitude; she can't appeal to the fact that 'x is funny' to warrant the attitude of finding something funny being called into play, since in her own

lights 'x is funny' is a projection of her attitude and so this would be circular reasoning.

The obvious third alternative, for McDowell, is a no-priority view. It is the interplay between feature and response that develops the ability to make a judgement that something is comic, and that one thing is more funny than another. David Wiggins (1987) similarly argued that there are property-response pairings where the response cannot be picked out without reference to the property or the property without making use of the response, for example 'comic'. Well-established pairings will get to the point where we can focus on the response and debate the marks of the property. It will therefore become 'essentially contestable' what a thing has to be like for there to be any reason to use the value term, and therefore contestable what the extension of the predicate is.

One might think - and McDowell does not directly address this point in this essay - that Blackburn can defend his position here by saying that the projectivist in making an evaluation *is* responding to a feature, but a descriptive feature of the world, not an evaluative one. In that case, though, he is threatened by the arguments that McDowell makes in *Non-Cognitivism and Rule Following*, that there is no basis for thinking that the 'independent' world can be reliably tracked. Otherwise, Blackburn's position would, I think, be committed to the view that attitudes are innate to us, and not just a few attitudes, but a sufficient number of different attitudes to account for all our different ethical judgements.

Virtue and Reason

To complete the discussion of John McDowell's writing, here is one further topic concerning his ideas on reasons for action; the notion of a virtuous person. This topic will come up again briefly in chapter three and again in chapter five. In the essay 'Virtue and Reason' (1979) McDowell explores the idea that 'the right thing to do' in a given situation is established *necessarily* by reference to what a 'virtuous' person does or would do in that situation, rather than being established by reference to some moral theory or set of principles. This line of thought is due to Aristotle, and one crucial element in McDowell's account is to understand the difference between a virtuous person and a *continent* one. A continent person is capable of understanding all the reasons for and against a course of action, but in order 'to do the right thing' she has to bring something else, some motivating element, to her decision. (An *incontinent* person might also grasp all the reasons but will ultimately act contrary to what the sum of the reasons indicates, perhaps because she lacks the will to make herself do the right thing.) A *virtuous* person grasps the relevant reasons and this grasp issues in her doing the right thing, automatically as it were.

To make this point, McDowell again rehearses Wittgenstein's arguments about following

a rule, as he does in McDowell (1981). This time the point of the argument against codifiability is that the virtuous person is not working through *a practical syllogism* in order to decide on an action: she does not have a rule as a major premise to which she can add the particular circumstances as a minor premise and thereby reach a conclusion. Rather, the virtuous person's cognisance and susceptibility to the right reasons results in right action.

The core explanation [of how a virtuous person reasons how to act] must at least be seen against the background of the agent's conception of how to live; and if the situation is one of those on which any of several concerns might impinge, the conception of how to live must be capable of actually entering our understanding of the action, explaining why it was this concern rather than any other that was drawn into operation. (McDowell 1979 reprinted in McDowell 1998 p148)

In another essay, 'Are Moral Requirements Hypothetical Imperatives?' (McDowell 1978), McDowell expresses this thought as that the salient reason or reasons *silence* all the other reasons or consideration for the truly virtuous. I will return to the idea of the virtuous agent in chapter five.

2.4 Blackburn and Dancy on thick concepts

I briefly noted in section 2.2 that the notion of a thick concept looks like a challenge to the possibility of classifying all predicates as strictly either evaluative or non-evaluative. This classification seems essential for Simon Blackburn's projectivist picture in which evaluative properties are projected onto a purely non-evaluative realm of fact. A projectivist needs a response to the challenge raised by thick concepts. As I discussed in section 2.3, writers such as John McDowell used evaluative concepts, or evaluations, as counter examples to Blackburn's picture and the non-cognitivist positions more generally. McDowell's concern was to make these arguments, not to give an account of thick concepts and I am not aware that he ever used the term 'thick concept'.

Simon Blackburn has two essays on the thick (Blackburn 1995, 2013). His view is that there are very few thick concepts, if any, if by thick concepts we signify concepts that are irreducibly both evaluative and descriptive as a matter of content. Those that may exist are uninteresting ones (for example 'kraut')¹², where an evaluation has been added to a concept with a clearly determined extension (for example German national). He thinks that there are words or terms that we regard as thick, in that they are conventionally associated with an attitude, but the attitude is no part of the meaning of the word. Indeed the attitude is *flexible*, being both

¹² Blackburn draws attention to derogatory terms, but I think his account would allow for the possibility that appreciative terms are also uninteresting

alterable and cancellable in the utterance of the word. We can see this fairly clearly if we consider firstly how readily any attitude we might associate with a word is changed, reversed or subverted by the way that a word is said, and secondly how uncertain it is which attitude is unequivocally associated with a word.

To start with an example; we use the term 'dainty' (referring to the concept DAINTY) in an utterance such as "the cup is dainty". DAINTY is a candidate thick concept, an aesthetic one rather than a moral one. Perhaps we are in a museum looking at a piece of Wedgewood porcelain. On Blackburn's account, in using the utterance we are saying that the cup is delicate and small *and on that account* good. This is a stronger claim than to say that the cup is delicate and small *and* good (in some way). A causal claim is present — the smallness and the daintiness of the cup gives rise to its being good. However, on many occasions we may think that DAINTINESS in cups is in no way positive; — we want morning coffee to come in a big cup and can easily imagine my saying that the cup is dainty in a tone of disappointment. This illustrates Blackburn's point that the attitude that goes with a term or concept is flexible.

To note again is how this gloss concerning thick terms and concepts moderates the notion that utterances of some sentences are *primarily* an expression of an attitude. Unless the word used in expressing an attitude is very thin ('good', 'bad'), then it seems that an evaluative utterance such as "the cup is dainty" must, in projectivist terms, be voicing a belief 'input' as well as expressing an attitude 'output'. Some critics of expressivism have assumed that where an attitude is expressed, an expressivist infers that there is no literal content to the utterance, no conditions of truth or falsity, and this inference is turned on the expressivist as the "speech-act fallacy". (Blackburn 1984 pp 169-170). This charge is rejected on two counts. Firstly, as long as the attitude may give the role of the utterance, there is an available argument that it is the best explanation for whatever ensuing commitments (states of mind) the speaker arrives at. Secondly, it is of no consequence that an utterance is descriptive as well, as long as 'its distinctive meaning – the aspect which separates it from any underlying B-descriptions – is expressive' (p170).

In his 1992 paper, Blackburn sets his remarks in the wider context of the communication of attitude in language. When we utter words, part of the information conveyed comes through intonation and tone. Intonation concerns the communication of information, force and attitude. Tone can vary the attitude expressed. Using a metaphor, Blackburn says that the intonation and tone are the tune, and we can alter what we mean in uttering a word by changing the tune we sing it to. However, it is not that the words and the tune are unrelated. 'Generally speaking, saying something with an attitude in your voice licenses the hearer to suppose that the attitude is expressed because of what is said.' (Blackburn 1992 p133).

Furthermore, Blackburn says that the realisation that an attitude is not endorsed may affect the application of a term or the making of a judgement. As an example, Blackburn suggests the possibility that some people are ‘fattist’ – they consistently find any instance of ‘x is fat’ also to be an instance of ‘x is fat↓’ where fat↓ is ‘fat’ said with a sneery tone. Here the suggestion is that ‘fat’ is a non-evaluative term referring to a non-evaluative concept whereas ‘fat↓’ is a term that pairs the same non-evaluative concept with a negative attitude. Blackburn suggests that these fattists would hesitate to use fat↓ in relation to someone who is undoubtedly, descriptively or objectively, fat, if they find that person admirable in other respects and consequently don’t want to suggest a negative attitude towards him or her. Blackburn’s example of a fat but admirable person is the opera singer Luciano Pavarotti. To summarise, according to Blackburn’s account, when we have something that we may at first pass to be a thick or evaluative term, or an evaluative judgement, what we have is the application of a descriptive property plus an attitude, and the major part of grasping which attitude is the manner of utterance.

Blackburn contrasts his account of thick concepts with the slightly more complex one of another expressivist, Alan Gibbard (Gibbard 1990). It is Gibbard who introduces the example of LEWD. LEWD is one of the thick concepts referred to in the literature as ‘objectionable’ or potentially objectionable. The notion of an objectionable thick concept is one where, although we understand clearly whether it has a positive or negative evaluation (in this case negative), we are unwilling to endorse the evaluation. There is nothing wrong with something being LEWD might be the thought. Other examples often given are CHASTE and BLASPHEMOUS. Writers who doubt that evaluation is part of the meaning of a thick term or concept of the concept point to objectionable thick concepts as evidence that the evaluative and non-evaluative aspects of a thick concept must be separable.

Blackburn says,

On Gibbard’s analysis, we have a basic description of the kind of thing that may be judged lewd: a display of sexuality. We have an attitude expressed, which Gibbard calls “L-censoriousness”. These parallel the two elements I have put in place for ‘fat’. Thirdly, however, we have the acceptance of a presupposition, namely, ‘the general importance of limiting sexual displays’. These limits give a rationale for the fourth element, which is the judgement that the feeling of L-censoriousness is warranted. Since this judgement is in turn given an expressivist analysis in Gibbard’s account, we see that two different states of mind are voiced: the L-censoriousness, and the acceptance of norms permitting L-censoriousness for acts of this type, as being, or in virtue of their being, acts of this type. (Blackburn 1992 p 134-5)

To Blackburn the last two of Gibbard’s elements need not be specified in relation to a particular example, but are backgrounded in our linguistic practice. A ‘norm-governed epistemology’ is in the background of meaning, as is the possibility that our judgement in a

particular case is unwarranted. These may be brought to the fore if I am challenged and am obliged to defend or explain my judgement in a particular case, but in Blackburn's view they are not communicated in the original judgement.

Blackburn says that he and Gibbard both argue that challenge and explanation are the features of communication that lead us into propositional expression rather than sticking to a system of 'Boo-hooray'. There is a good reason to move from intonation to lexical expression. Intonation is deniable afterwards. For example, if I say, "John likes computers" with a sneer in my voice, I can answer any objection to my attitude by innocently saying "I only said that John likes computers." Whereas if I say, "John is a nerd," it is harder for me to deny that my attitude towards John might be taken as negative or critical. The point, I take it, is that it is practically useful to have words that are conventionally associated with an attitude, because the attitude is retained in reported speech when the utterance of the moment has passed. This applies as much to positive words as negative ones. ("I promise you, I was really nice about Jane in the meeting. I said '...'.")

However, the 'stabilising' effect of these lexical conventions is only partly achieved. The intonation of the moment will always be a more reliable guide to what a speaker is thinking, and while the choice of word is significant 'this is not to say that there is just one thing...that it signifies' (Blackburn 1992 p 138). This is where Blackburn's analysis comes apart from Gibbard's. In the example of LEWD Blackburn does not think that censoriousness is compulsory. If we think of a carnival or Shakespeare's fools, or some superstar pop singers – just some of his examples of LEWDNESS – our identifiable attitude or attitudes are not necessarily critical. This is a general point about the thick, Blackburn thinks, it is not that Gibbard hit on a peculiar case in LEWD.

In the concluding paragraphs, Blackburn summarises that there are no thick concepts, and we can see this by the demonstrable detachability of what we identify as the attitude from the descriptive element. He acknowledges that a defender of irreducible, or non-analysable, thick concepts might say that demonstrating how some examples are analysable does not prove that there are no unanalysable thick concepts. However, if concepts are unified as cognitivists claim, that must be because somehow the two separate elements have been formed into a compound or amalgam, each side infusing the other, and neither able to stand alone. In that case, the concept is only stable in that state, and if it 'dissolves' a person loses the concept entirely.

In the essay 'In Defense of Thick Concepts' (1995), which is a discussion of Blackburn (1992), Jonathan Dancy puts the 'standard' cognitive notion of thick concepts that Blackburn is arguing against as follows. Thick concepts are pairings of properties and attitudes. The

attitude is fairly simple (perhaps pro or con) and is fixed (somehow) in the concept. The property is more complex, but crucially is not independently identifiable; evaluation partly determines description. Since a description is not identifiable separately, non-cognitivism (including Blackburn's expressivism) cannot account for or handle these thick concepts, and so is forced to deny their existence.

Dancy summarises Blackburn's case in reply. If we look at typical examples of the thick, the attitudes associated with them are both variable and multiple, and the concepts themselves change over time. If a thick concept is a property and single attitude pairing, these terms are not thick concepts. In fact, expressivism is much better equipped to handle these 'complex, fluid and mobile' terms (Blackburn quoted Dancy 1995 p265). Cognitivists squash them into the straightjacket that is their notion of a thick concept. Finally, the relationship between evaluation and description does not help the cognitivist cause, because a non-cognitivist can show how evaluation partly determines extension. Dancy concludes:

When Blackburn has finished then, the tables have been turned. Instead of the thick constituting a problem for non-cognitivism, it turns out to expose weaknesses in the position of those who were trying to use it as a weapon - a genuinely two-edged sword. (Dancy 1995 p 266).

Dancy says that cognitivists should 'accept with gratitude' Blackburn's insights about the fluidity and variability of the thick. A first thought is whether Blackburn is on the same track as the cognitivist David Wiggins in his remarks about essential contestability. (Wiggins 1987). However essential contestability concerns only our decision whether, for example, an action was lewd or not i.e. whether it has the property lewd. 'The notion of essential contestability is just a different notion from that of the variable and the fluid' (p266).

Fluidity and variability are two different claims, and Dancy thinks that either is sufficient to ground Blackburn's claim that, although we talk about the thick, there are really no such thing as thick concepts. Fluidity is change over time. Cognitivists must defend their position against the thought that it 'concretises' meaning, and greatly increases the risk that two speakers will be 'talking past each other' when they can't agree because really they have two different concepts. "For if you don't respond to lewdness as I do, then on the thickie picture your amalgamated concept is not mine, and we are left in incommunicable solitude." (Blackburn quoted in Dancy p 267)

The thought is that, on the 'thickie' picture, disagreement must mean that person A has a concept of lewdness with a different attitude bolted into it than person B. (We go to a pop concert and I find aspects of the performance highly regrettable and you find the same aspects great fun.) So we are disagreeing systematically on the response case by case. This result follows because Blackburn thinks that the cognitivist thick concept picture must be one

of an amalgamation of two separate 'elements', one descriptive and one evaluative. All the change that can happen to a concept on this account is that it can dissolve back into its separate parts and be lost. A defender of thick concepts, says Dancy, could and would dispute this by arguing against the amalgamation idea and any possibility that the descriptive element can be freestanding.

Furthermore the initial difficulty is very overstated. 'The difficulty is supposed to be that the protester wants to say "Lewdness just isn't important at all, in the sort of way it is generally taken to be," but that someone who thinks this cannot even have the concept of lewdness, though he needs it in order to make his point.' (Dancy p269) This supposedly leads to the talking past each other result; but the whole charge is artificial, indeed 'spurious'. Someone who rejects the common use of the concept LEWD can still perfectly well understand how the concept functions. It is only by understanding a concept that one can reject it, for example, in the case of LEWD, objecting to the social practice that judges behaviour in that way. 'To reject a concept one has to do more than fail to see its point; one has to know what its point is and reject it for that reason.' (p269). A person who rejects a concept perfectly well has the ability to move on and create a successor concept.

Does evaluation shape extension? As we have seen, Blackburn sets up an example where the evaluation fat↓ has a readily graspable descriptive aspect – fat. Fattists are people who consistently take a negative attitude to those who are fat – they evaluate them as fat↓, that is, fat said with a sneery tone. However, since it is plausible to suppose that a fattist would not want to express a negative attitude towards someone she overall terrifically admires, she may withhold the evaluation on occasion. So, evaluation (in this case, attitude) partly determines extension. Dancy's reply here is that there is no issue here between cognitivists and non-cognitivists about whether evaluation (partly) determines extension in the application of the whole concept (or term). What the cognitivists are claiming is that there is no way of getting at the supposed descriptive extension of a concept without evaluation. Blackburn's fattist, on the other hand, won't be able to deny, if pressed, that Pavarotti was fat.

Blackburn is concerned that we should be able to avoid endorsing judgements that we do not agree with, and he thinks that a cognitivist thick concept picture risks this, because if the satisfaction conditions for the concept are met then the evaluative judgement (or attitude) must follow. If this is a problem in the case of *evaluations* it would be a still greater problem if the application of thick concepts gives rise to reasons for action. For then we would be pressed into accepting reasons to act in some way that otherwise we do not think we have. We will see in chapter four that Ulrike Heuer (2012) takes 'Blackburn's challenge' to apply to the reasons for action claim.

To return to the notion that on the cognitivist picture a thick concept must be an indissoluble amalgamation of two separate 'elements', one descriptive and one evaluative, which Blackburn makes and Dancy opposes. Debbie Roberts (2013a) argues that this picture only arises if two separate notions of entanglement found in John McDowell's writings are combined. McDowell accepts a difference between evaluations and practical judgements or directives. (This is David Wiggins distinction, noted above.) Practical judgements are decisions about what one ought to do or it would be best to do in a particular case. In 'Are Moral Requirements Hypothetical Imperatives' and 'Virtue Reason', Roberts says, McDowell discusses directives and holds that for a virtuous person there is an 'entanglement' between cognition and will in her perception of her situation. That is, her cognition of her situation influences her will; her cognition engenders a pro or con attitude to an action.

In 'Non-cognitivism and Rule Following' McDowell discusses evaluations and his aim 'is to argue that competence with an evaluative concept cannot be disentangled (into sensitivity to a non-evaluative feature of the world on the one hand, and a propensity to a certain attitude on the other)' (Roberts 2013 p83). It is these remarks that continue to the 'shapelessness' of evaluations; that there need be nothing non-evaluative, no property or feature, that items falling under an evaluative concept have in common.

In Roberts' view, it is only if the evaluative point of view required to grasp an evaluative concept is equated with the virtuous person's point of view, that the ability to see whether an evaluative concept applies in a certain situation must also contain a thin evaluation or in other words the pro or con attitude to an action. Only then do you get the view that, if a thick concept is irreducible, it is an 'indissoluble amalgam' of the thin and the non-evaluative. What I am drawing attention to here is Roberts' link between the thin and action-guidingness.

2.5 Conclusion

In this chapter I have said something about expressivist theories of language, and that these theories have some appeal. Simon Blackburn, whose own particular expressivist position is called 'projectionism', has deployed some arguments to show how his theory, which denies that there can be any genuinely thick concepts, better explains the behaviour of thick terms than a cognitivist account does. In particular Blackburn argues that the evaluative upshot of the instantiation of these concepts is not fixed or indefeasible.

Opposing this view John McDowell has argued that the distinction between fact and value

which is essential to Blackburn's account, and to non-cognitivism more generally, is not available. Furthermore, Jonathan Dancy has argued that Blackburn's insights about thick concepts can be accommodated by cognitivist accounts. I find McDowell's arguments both plausible and persuasive, so I reject Blackburn's claim that there are no genuinely thick concepts, where evaluation is part of the content of the concept.

Finding the arguments for projectivism wanting was important for the project of this thesis, as it is essential to my project that the application of a thick concept can be truth evaluable, can yield a fact, contra expressivism. This is because the notion of reasons for action that I am using is that reasons are facts. If propositions involving thick concepts cannot be true, cannot be facts, then the claim that thick concepts provide reasons for action falls at the first hurdle.

In the next chapter, I move on to some more recent literature, which has thick concepts as its subject. One issue taken up in this writing is that, while McDowell's arguments are effective against the idea that thick concepts are systematically analysable or reducible, this does not secure that they are *evaluative*. This result might also undermine my project: if 'thick concepts' are non-evaluative, what argument could there be that they are reason-giving where (other) non-evaluative concepts are not.

Chapter 3: More recent debates: thick concepts and evaluation

3.1 Introduction

In chapter two I said that cognitivist writers have pressed non-cognitivists to show what sort of propositions count as descriptions of reality, descriptions involving no evaluative point of view, rather than *assuming* that such propositions are available. Propositions are formed from concepts. In the case of thick concepts, concepts which appear to be both evaluative and to some extent descriptively specific, there are arguments that the extension of the term associated with the concept cannot be identified by non-evaluative features alone. If these arguments succeed, then thick concepts are an important contribution in undermining the non-cognitivist project.

In this chapter, I turn to more recent literature on thick concepts and some detailed explorations of how thick concepts are evaluative, *if* they are evaluative. I start with the contrast between thick and thin (section 3.2). I then show how the disentangling arguments due to John McDowell and Bernard Williams are expressed and used in the more recent literature (section 3.3).

Daniel Elstein and Thomas Hurka grant the point that, in the case of thick concepts, it may not be possible to isolate non-evaluative content without resort to some evaluative resources. However, they say, it does not follow from this that the overarching pro or con evaluation associated with a concept cannot be separated from the rest of the content or what drives the extension of the term. In section 3.4, I discuss their argument for the possibility of a reductive analysis of thick concepts and conclude that it is not convincing.

Pekka Väyrynen argues that only thin evaluative concepts are 'inherently' evaluative as a matter of content. What is on the case with 'paradigmatic' thick concepts is that the concepts are plausibly unitary but that is because they are non-evaluative as a matter of content, and they have evaluation normally associated with them as a matter of conversational implication. In section 3.5 I present his positive view and discuss one of his principal arguments for it.

Debbie Roberts and Simon Kirchin argue that Väyrynen's position, and other reductive positions, implicitly require a claim or view that all evaluation must involve thin evaluation.

For writers who want to oppose this claim, then a different account of evaluation is indicated and they need to show what that can be. Their positions are in sympathy but different. Roberts argues for ‘irreducible thickness’ as a live possibility for evaluation. I present this in section 3.6. Kirchin’s argues for the ‘liberal view’; thick concepts have a number of aspects and when it comes to being evaluative, a concept qualifies if even one instance of the use of the concept is an evaluative use. I discuss this positive view in section 3.7. In section 3.8 I sum up, endorsing Kirchin and Roberts’ picture that thick concepts are evaluative as a matter of content. This then sets the scene for the remainder of the thesis where I will consider, on the basis of the position that thick concepts are irreducibly evaluative, how such concepts can play a role in providing reasons for action.

3.2 Thick and thin

In the literature on thick concepts, the thick is contrasted to the thin. Some words and concepts are accepted as being evaluative on all sides: GOOD, RIGHT, FAIR, BAD, WRONG. One characteristic of these concepts is that they are either positive or negative; they clearly indicate value or disvalue, approbation or disapprobation. A second characteristic is that many and diverse items can fall under them. Hitting another person can be wrong, but so can a team’s tactics in a game of cricket, or an embroidery stitch. ‘The mark of a concept like ‘right’ is that it says nothing about what *other* properties an item falling under it has’. (Elstein and Hurka 2008 p 516 my italics). These concepts are known as thin concepts.

Bernard Williams called these concepts the most ‘abstract’ concepts. He held that we acquire them by reflecting on more culturally specific concepts and abstracting from those. Other writers have argued that abstract or thin evaluative concepts come first, and other more specific concepts are gained as the scope of value or disvalue is restricted. In this thesis, nothing turns on having a view about the priority¹³ or otherwise of the thin. I will assume a no-priority view, where neither thick nor thin concepts are prior. (Kirchin 2017)

To say that the thick is contrasted to the thin, is not to say that there is a distinction or that the distinction between the two is clear or uncontested. Sophie Grace Chappell (2013) has argued that there is no such thing as an absolutely thin evaluative concept. One of her examples is the Japanese concept GIRI. Apparently, says Chappell, it is most natural to translate ‘right’ and ‘wrong’ in English into GIRI terms in Japanese, and yet ‘the concept of giri... holds the highest regard in human relationships...[and] (t)he most basic definition one can give giri is a debt of gratitude and a self-sacrificing pursuit of their happiness’ (Namiko Abe quoted Chappell 2013 p324). Matti Eklund (2011) also argues that it is difficult to

¹³ See Hurley 1989. Hurley called the priority of the thin ‘thin centralism’.

distinguish the thick from the thin because it is difficult to give a characterisation of the thick that is not also satisfied by thin concepts. So for example, Gibbard's initial informal characterisation of the thick is 'A term stands for a thick concept if it praises or condemns an action as having a certain property. (Gibbard 1992 quoted Eklund 2011 p26). However, '(t)here is as much reason to think that 'good' satisfies this condition as that a term like 'courageous' does.' (Ibid p27). I think Eklund is right here.

On a related note, in that they also disagree with a sharp distinction between the thick and the thin, Alan Thomas and Edward Harcourt (2013) say that,

the distinction between thick and thin ethical concepts is best seen as a continuum. If this is right, for a concept to be (relatively) thick is just for it to be more specific and to have narrower satisfaction conditions than one that is (relatively) thin...But it does not follow from any of this that there is nothing interesting going on towards one end of the continuum that is not also going on towards the other. (Harcourt and Thomas 2013 pp 41 - 42)

Debbie Roberts (2013a) argues that furthermore there is a possibility that, in anchoring an understanding of evaluation in the thin, a bias is introduced into the elucidation of thick concepts from the beginning:

Orthodoxy has it that evaluation must be positive or negative; it must have a pro or con flavour. This translates naturally into the thought that any evaluation, to be evaluation, must have a thin evaluative element. If you evaluate something you must, the thought goes, be evaluating it as good in some respect or to some extent, or bad in some respect or to some extent. According to the orthodoxy then, to be evaluative, thick concepts must have a link to thin evaluation. (Roberts 2013 p 78)

This matters, because, as Philippa Foot remarked, if 'a special feature found in all value judgements' is identified, then a contrast between fact and value becomes possible (Foot 1958 p 505). Foot was arguing against this 'special feature' as 'variously identified as expressing feelings, expressing and inducing attitudes, or commending' (ibid) but I think her point generalises to any position where evaluation is taken necessarily to indicate a pro or con stance.

My interest in the thick and thin distinction in this thesis is that a thin-concept fact is not very informative as far as reasons for action go. I have yet to discuss reasons for action, but the thought is something like this: to say 'the fact that it is wrong to kick the cat is a reason for me not to kick the cat' still leaves it wide open just why I should not kick the cat. To say 'the fact that it is CRUEL to kick the cat is a reason for me not to kick the cat,' on the other hand, says quite a lot about what sort of consideration it is that speaks against kicking the cat, suggesting that thick (or thicker) concepts are particularly apt to provide reasons for, or against, action.

3.3 Disentangling arguments: shapelessness and graspability

In chapter two, I outlined why there is reason to think that it is not possible to ‘disentangle’ *all* evaluative concepts into at least two components. This disentangling argument (so called, but in effect an *anti*-disentangling argument) has two lines of attack, which can be called the graspability point and the shapelessness point. Debbie Roberts says the argument as a whole is often attributed to McDowell in its entirety but it is Bernard Williams who makes the graspability point. McDowell by highlighting the issue of shapelessness, underpins the graspability point, which he agrees with, with an explanation of why it is the case.

The first strand of the disentangling argument is graspability. If it is possible to identify the extension of a thick evaluative concept without using that very evaluation, typically put as being able to apply the concept correctly by ‘descriptive’ features alone, then an ‘outsider’ could mimic ‘insiders’ in using the concept without having any insight or appreciation, any *grasp*, of the point of the evaluation that is being made. An outsider is someone from a very separate culture, say. The insiders are the mature and competent users of the concept. Would an outsider, after being shown some examples of KINDNESS, be able to carry on and reliably classify behaviour she encounters from then on in as KIND OR NOT KIND?

Bernard Williams (1985) finds this implausible, although one way he expresses the thought is that an outsider would not *use* the insider’s concept. An outsider would not say ‘x is KIND’ or ‘x is KIND *is true*’, she would not commit to the belief that such is the case:

The observer does not have a term that picks out just the same things as the locals’ term picks out and, at the same time, is entirely independent of the interest that shapes their use. (He has, of course, an expression such as “what they call F,” and the fact that he can use it, although it is not independent on their term, is important: his intelligent use of it shows that he can indeed understand their use of their term, although he cannot use it himself.) (Williams 1985 p 159)

Another way of expressing the point is that the outsider would need to be able to make the evaluations the insiders make, that is, become an insider at least to the extent of using the concept. To generalise a little we could say that the graspability point is that to use a thick concept requires making *some* evaluations, and this is how Daniel Elstein and Thomas Hurka understand the point, as will be seen in section 3.4. This way of putting things leaves open the possibility that the outsider could apply the concept reliably without *endorsing* the evaluation, as Williams appears to acknowledge in the quotation above. This possibility will be returned to in Chapter 4 in the discussion about Williams’ account of reasons for action when taken together with his position on thick concepts.

The second strand of the disentangling argument is shapelessness. In order for people to carry on applying a concept correctly, there needs to be some similarity, some common feature, that everything that falls under the concept possesses. Plausibly, there are some evaluative concepts whose extensions include things that have no other feature to unite them as a group beyond the evaluative feature in question. The term used for this claim is that as a group the totality of the items falling under the concepts are *shapeless* with regard to the non-evaluative.

As I wrote in section 2.3, John McDowell is concerned to argue against any ontological claim that evaluative properties are not genuine features of the world, whereas non-evaluative ones are. An anti-realist, claiming that evaluative features are not real, will have to acknowledge that competence with evaluative concepts requires disentanglement into (1) an ability to recognise some aspect of the world and (2) an attitude or non-cognitive state that is the special perspective from which an item is endowed with value. (McDowell 1981 p 198).

Now, crucially, for this to be possible, the aspects of the world, that fall under e.g. TACTFUL would have to be classifiable. 'A succession of judgements or utterances, to be intelligible as an application of a single concept to different objects, must belong to a practice of going on doing the same thing' (p 203). However, McDowell uses Wittgenstein's remarks on following a rule to illustrate that this cannot be established definitively as codification of rules or grasp of a universal.

Another way of expressing shapelessness is by saying that an evaluative concept *outruns* any attempt to specify or capture its extension completely. This outrunning is both because there may always be new circumstances or sets of features that the concept applies to, and because new, additional, extra features can cancel out features that, in simpler examples, were what merited the application of the concept. As an example of this, Simon Kirchin (2017) describes an adult giving a child chocolate. In a first instance this action is KIND because the child really wants the chocolate. However, in a second otherwise identical instance, if a consideration of the child's longer-term health is added, the effect on her teeth of eating sugar say, then giving the chocolate is NOT KIND. Then again, in a third instance, if the child is very distressed, or this particular chocolate bar has some special emotional association for her, giving her the chocolate is KIND, even though the long-term health implications remain.

To bring together the two strands of the disentangling argument, we could think of Kirchin's example where the central feature of giving a child chocolate can be KIND or NOT KIND depending on the other considerations present. I might extend Kirchin's example further to a case where an adult hands out chocolate to a number of children. The chocolate has been promised to the children if they achieve some goal, so the giving is not best characterised as

KIND. (It is perhaps better to characterise it as FAIR or HONEST.) However, if one child has been having a rough time, and the giver makes sure she gets her favourite chocolate bar from the goody bag, then the giving to this one child is perhaps KIND. Then again, in an otherwise similar example, if any favour like this will expose the child to bullying from the others, then singling her out may be NOT KIND. What could an outsider with no insight into KINDNESS make of this? Her difficulty is that, not only must she pick up on likely relevant features, she must be able to assess the whole combination of features.

Pekka Väyrynen has pointed out that shapelessness is not confined to evaluative concepts. He gives the example of the concept 'old'. 'Old' is unavoidably context specific – an object can be old in some comparisons but young in others; young for an elephant but old for a bumble bee (Väyrynen 2013). Plausibly, all the items that fall under the concept 'old' have no other feature in common. Notwithstanding that the term 'old' is sometimes used negatively, 'old' is not typically thought of as an evaluative concept. Kirchin, in response, acknowledges this, and accepts Väyrynen's conclusion that shapelessness in and of itself does not substantiate that a concept is evaluative, or that paradigmatic examples of thick concepts are evaluative as a matter of content.

3.4 Daniel Elstein and Thomas Hurka: a reductive analysis

If an analysis of thick concepts is possible, what form could it take? The analysis must be complete, it must explain the whole concept, or it isn't an analysis (Kirchin 2017). A reductionist view of thick concepts requires that such concepts can be analysed into at least two components. In its simplest form, this reduction is presented as an evaluative component and a non-evaluative or descriptive component. It is taken in the literature that purely evaluative thin concepts are available to play the first role of being the evaluative component. (I will accept this for the exposition, although I note from section 3.2 that there are some objections to this pure thin view.) However, a thin evaluative concept says nothing about the other features that an item falling under it has. Therefore, in a two-part analysis where one part is a thin concept, all the work of picking out those features must be done by the other component of the concept. In other words, the non-evaluative component must drive the extension of the concept as a whole. This can work for some concepts or putative concepts. The usual examples given in the literature are derogatory terms such as 'kraut' or 'frog', where the extension is co-terminus with a particular nationality. We would not want to give an account of thick concepts based on those.

Daniel Elstein and Thomas Hurka (2009) write that the non-reductionist claim distillable

from the disentangling arguments of John McDowell and Bernard Williams, is that ‘we cannot determine the extension of a thick concept without making some evaluations’ (Elstein and Hurka 2009 p 519). I discussed this in the previous section. Elstein and Hurka grant that the disentangling claim is indeed effective against the simplest available reductive pattern, which analyses a thick concept into two parts; as it were, the evaluative and nothing but the evaluative part, and the non-evaluative and nothing but the non-evaluative part. However, they assert that those philosophers who make the disentangling argument assume that this is the only pattern of analysis available to a reductive account. If we throw off this constraint, they say, a slightly more complex analysis can grant the disentanglement claim and still be unimpeachably reductive.

The following passage outlines their suggestion for a more complex analysis:

(T)here can be concepts whose descriptive component defines an area in conceptual space within which admissible good- or right-making properties must be found, so any use of the concept associating it with properties outside that area is a misuse, but does not identify any specific point within the area as uniquely correct, as a concept like ‘Kraut’ does. The concept therefore has descriptive content, but this content is not completely determinate. The pattern of this analysis is something like ‘x is good, and there are properties X, Y, and Z (not specified) of general type A (specified), such that x has X, Y and Z, and X, Y and Z make anything that has them good.’ This pattern is reductive, because it uses only the thin concept ‘good’ and the descriptive concept ‘A’. But it accommodates the key premise of the disentanglement argument because determining which properties of type A are the good-making ones, which we must do to determine the concept’s extension, requires evaluative judgement. (Elstein and Hurka 2009 p 521)

Elstein and Hurka introduce the terms ‘global evaluation’ and ‘embedded evaluation’. The evaluation represented by the pattern starting “x is good’ is termed a global evaluation by Elstein and Hurka, because it governs the whole concept; anything that falls under the concept is evaluated as good (or bad). This is in contrast to an embedded evaluation, which is a value internal to the concept. For example, were we to think of the concept COURAGE as being ‘good in virtue of accepting the risk of harm in pursuit of a worthwhile end or objective’, then ‘worthwhile’ is an embedded evaluation, because it applies to the end or objective of a COURAGEOUS action but not the COURAGEOUS action itself. (It is useful note this formulation because Väyrynen makes use of it in his position which is discussed in a later section.)

Clearly Elstein and Hurka’s pattern of analysis accepts or assumes codifiability, that a rule or formula can be made. They see no difficulty with this, as any argument that evaluative concepts are uncodifiable would apply equally to thin evaluative concepts and thus not be a distinction between the thin and the thick. The first example that Elstein and Hurka give is the putative thick evaluative concept DISTRIBUTIVELY JUST applied to the sharing out of happiness amongst a population. Putting DISTRIBUTIVELY JUST into their reductive pattern yields:

‘x is good, and there are properties X, Y and Z (not specified) that distributions have as distributions, or in virtue of their distributive shape, such that x has X, Y and Z, and X, Y and Z make any distribution that has them good.’ (Ibid P 522)

The ‘general type’ of property here is being distributive according to a particular property or yielding a particular distributive shape. In the example given, the specific distributive property that is good making requires a choice between whether a JUST DISTRIBUTION of happiness is one of equal shares for everyone or one proportioned to merit. Making this choice here is a matter of evaluative judgement. Therefore, some evaluation has been made in determining the extension of the concept. Therefore we have a thick concept, DISTRIBUTIVELY JUST, and the concept is analysable, and the analysis is consistent with the disentangling argument claim that determining the extension of a thick concept requires making some evaluations.

I don’t think Elstein and Hurka’s proposals can be right. According to the disentangling argument, as I understand it, the implausibility of it always being possible to determine what items fall under an evaluative concept by non-evaluative means alone, is that, were it possible, no insight into the significance of the evaluation would be required on the part of the person making the evaluative judgement. The evaluative point, as Williams called it, must have something to do with the saliency of the concept, why it matters or is relevant to this community. Otherwise an ‘outsider’ to the community and to the community’s evaluative understanding could mimic an ‘insider’ in making judgements. Determining whether the DISTRIBUTIVE JUSTICE of happiness is a matter of equality or merit does not help with this. Once an outsider has observed the distribution of one of the possible properties giving rise to a judgement of DISTRIBUTIVE JUSTICE, (and assuming, as Elstein and Hurka must do, that both the property and happiness are observable) then she could continue in future cases to apply the judgement of DISTRIBUTIVE JUSTICE.

For any example of DISTRIBUTIVE JUSTICE to go through the suggested reductive pattern successfully, we also have to state what item (what good or benefit in the economic sense of good) is being distributed. The writers seemed to have overlooked this point in giving DISTRIBUTIVE JUSTICE as their example concept to show their pattern of analysis. To repeat the formulation; ‘x is good, and there are properties X, Y and Z (not specified) that distributions have as distributions, or in virtue of their distributive shape, such that x has X, Y and Z, and X, Y and Z make any distribution that has them good.’ Taken at face value, this indicates that *some* properties, the X, Y and Z properties, would make *any* distribution good.

Suppose we talk about distributions of opportunity, rather than happiness. I may think that [the distribution of] the opportunity to go to school is fair or just only if it is equally had

by all, but that [the distribution of] the opportunity to compete in the Olympics is just only if the selection process is based on merit. I do think both of these things and I think I am using the same concept in both cases. Imagine however, that I am presented with the case of the opportunity to go to school first. I see a distribution that is entirely equal, and, consistent with Elstein and Hurka's proposal, make an evaluative judgement that the equal distribution property is the good-making one. I am then presented with a case where everyone's opportunity to compete in the Olympics is equal (we just draw lots, say). On Elstein and Hurka's analysis 'there are properties X, Y and Z ...that distributions have as distributions, or in virtue of their distributive shape, such that x has X, Y and Z, and X, Y and Z make any distribution that has them good'. I have already judged that, for me, equality of distribution is good making. So, to be consistent in my application of the concept, I would have to approve the equal distribution of opportunity to compete at the Olympics. However, as I have said, this does not accord with my intuition about how I would judge the two cases. As the shapelessness and outrunning arguments showed in the previous section, it seems better to judge each case independently on a case-by-case basis.

Indeed, the reductive pattern as given might yield results for things we wouldn't dream of evaluating as good or not. A distribution is '(t)he way in which a particular measurement or characteristic is spread over the members of a class.' (Oxford English Dictionary accessed 03.01.2020.) So, we could have the age at death of Dartmoor ponies, for example, or the number of reusable shopping bags owned by current York University students. If Dartmoor ponies all or mostly die at the same age, or York students all own two 'bags for life', then the reductive analysis of distributive justice proposed would yield that these distributions are good, in a moral sense, to anyone who thinks equality is the good making property to merit the application of the concept DISTRIBUTIVE JUSTICE. Matti Eklund calls this 'the puzzle of seeming sufficiency' (Eklund 2011). If we think that we have a formulation of what it takes for an evaluative thick concept to apply, then that formulation appears to be *sufficient* to warrant the application of the concept.

I am labouring these points to show just how difficult it is to get a usable and complete reductive formulation. My view is that DISTRIBUTIVE JUSTICE - an intellectual choice of example it has to be said¹⁴ - is not a discrete evaluative concept at all. JUST and JUSTICE are evaluative concepts and distributions are a domain in which those concepts are sometimes salient. Possibly, the only evaluations that can sensibly be applied to distributions are matters of JUSTICE, and then only when some matter of ethical debate or matter of social policy is in question. (Elstein and Hurka's second example concept of SELFISH does not, in their discussion of it, avoid the objections I have raised, although it is a more commonly used concept, because

¹⁴ Obviously DISTRIBUTIVE JUSTICE has a pedigree in the debate between Nozick and Rawls, and is thereby very plausibly an essentially contestable concept

they give the evaluative judgement required again as a choice between ethical theories.)

Edward Harcourt and Alan Thomas (2013) put the objection to Elstein and Hurka's account in a different way. They say that what the writers are doing with their pattern is to quantify over a range of non-evaluative properties or features. Quantifying is not the same as evaluating. Presumably, say Harcourt and Thomas, in the distributive justice example, both the person who thinks that distributions should be equal and the person who thinks that distributions should be based on merit thinks that she is correct. There is no evaluation to be made by either of them. The property, or properties, that each think is relevant is, or are, non-evaluative.

A second point made by Harcourt and Thomas (and in line with their own view of the thin and the thick as being organised in a determinable / determinate relation,) is that a non-reductivist account is compatible with Elstein and Hurka's pattern. For example (in my words)

x is good, and there is a property courage, and x has the property courage, and the property courage makes anything that has it good.

Compare:

x is coloured, and there is a property red, and x has the property red, and the property red makes anything that has it coloured.

There is, according to Harcourt and Thomas, no way that the property of red can be analysed to yield a sub-component that can be specified independently of the property red and can be added to coloured to yield red. Although of course there are lots of other colour properties besides red that would make anything that had them coloured. Analogously, there is no way of analysing COURAGE to yield an independently specifiable component that can be added to good to yield COURAGE. This is the non-reductivist's claim. This thought could be set out in Elstein and Hurka's pattern:

x is good, and there is a property courage [of general type courage], such that x has the property courage and courage makes anything that has it good.

Nothing about Elstein and Hurka's pattern of analysis rules out the non-reductivist's position, they say; the pattern just puts the possibility of a reduction on the table.

To reiterate, for the purposes of my argument I want to cast doubt on the idea that a reductive analysis of thick concepts is always possible. I want it to be plausible that at least

some thick concepts are evaluative and irreducible. If Elstein and Hurka's suggestion is a live possibility, then the claim that thick concepts provide reasons for action could be met by the claim that thin concepts are action-guiding, and it is the thin evaluative content of a thick concept that is doing the work in providing reasons. Hence it has been important to my argument to cast doubt on Elstein and Hurka's proposed reductive analysis. The idea that thick concepts are irreducibly evaluative remains on the table.

3.5 Pekka Väyrynen: The Pragmatic View

Pekka Väyrynen holds that paradigmatic thick terms and concepts can be accounted for as having entirely non-evaluative meaning and content (Väyrynen 2013). Any evaluation associated with thick concepts is a matter of the pragmatics of language, the everyday use of the thick terms. He calls it the Pragmatic View. This is a deflationary account of thick concepts, in that according to Väyrynen's account thick concepts have no deep and distinctive significance in ethics and do not speak against a fact-value distinction. Many of Väyrynen's arguments make use of the philosophy of language and linguistics.

Väyrynen's 'working definition' of thick concepts uses, as a starting point, the pattern of analysis introduced by Daniel Elstein and Thomas Hurka (2009) and discussed in section 3.4, as well as their notions of global and embedded evaluation. It is as follows:

(TC) A term T stands for a thick concept if (i) x is T entails, as a conceptual matter, that x falls under some distinct concept or concepts, N1,...,Nn (not specified) of a certain (specified) general type A (where A and N1,...,Nn may be purely non-evaluative descriptions or involve embedded evaluations) and (ii) literal uses of x is T in normal contexts somehow convey global evaluations to the effect that x is good, or bad, in a certain way. (Väyrynen 2013 p42)

Väyrynen is aiming to set up an apparatus for discussion, over and above giving a list of common or paradigmatic thick concept examples and saying that this is the sort of thing we are discussing. However it seems to me that this set-up, by assuming two components from the outset, is hardly neutral between the various positions on evaluative language on offer. Certainly Väyrynen appears to hold what Simon Kirchin (2016) calls the 'conservative' view of evaluation, where evaluation must be pro or con. What Väyrynen himself says,

My suggestion for characterising evaluation...is to understand it as information that is somehow positive or negative in flavor. This needn't mean...an expression of a "boo" or a "hurrah". Evaluation might rather be understood as information to the effect that something has a positive or negative standing — merit or demerit, worth or unworth — relative to a certain kind of standard. (Väyrynen 2013 p29)

While Väyrynen appears to accept that some concepts can be evaluative as a matter of content, these are the thin concepts and there are relatively few of them (and as Kirchin notes, Väyrynen does not make any arguments to support how or why thin concepts are evaluative). Väyrynen in turn argues that some writers are apt to confuse shapelessness with evaluativeness. He presses writers with non-reductivist positions to say what evaluation is, if it is not a matter of being pro or con.¹⁵ (Väyrynen 2013 p209)

Given that we do use language to express evaluation Väyrynen's project is to explain how evaluation is connected to terms and concepts, without being part of the meanings of the terms or content of the concepts (except for thin evaluative terms and concepts). His positive view is expressed as follows:

Pragmatic View: Global T-evaluations [thick evaluations] are implications of T-utterances [thick utterances] which are normally not at issue in their literal uses in normal contexts, and which arise conversationally. (Väyrynen 2013 p122)¹⁶

He thinks that, were T-evaluations part of the meaning of thick terms, we would expect the T-evaluations to be semantic entailments, semantic presuppositions or conventional implicatures (Roberts 2015). He can show by linguistic tests that these explanations fail, leaving the pragmatic view as the best explanation.

To note here is that the act of usage is all-important; in order to 'read-off' the evaluation, there has to be an utterance in a context. This is not to say that if I see written on an otherwise blank sheet of paper that, "the term 'beautiful' refers to the concept BEAUTIFUL", I will have no inkling that 'beautiful' is associated with a positive evaluation. Clearly, as competent English speakers, we are likely to understand that. Väyrynen allows that, through repeated and long usage, the evaluative stance can become conventionalised to the term. However, he thinks that *in the beginning* the evaluations were conversational implications produced when the words were used.

In order to argue for his position, Väyrynen uses examples of so-called 'objectionable' thick terms and concepts. Objectionable thick terms and concepts are ones where a speaker is competent with the term or concept, but does not endorse the evaluative stance associated with it. Examples are CHASTE, LEWD, and BLASPHEMOUS. Now a speaker who doesn't like these terms can avoid using them, but how can she respond if someone else brings them up in

¹⁵ This point is particularly addressed to irreducible thickness views, which I will come on to in the next section.

¹⁶ Väyrynen uses 'implication' as a general term. [There are] 'different mechanisms by which speakers can communicate information through their linguistic behaviour...typically I'll use "implication" as an umbrella term for the full range of mechanisms of information transfer from semantic entailments to various merely conversational implications. This is a diverse range.' Väyrynen (2013 p 51)

conversation and forces her to agree or disagree whether the concept is instantiated.

Väyrynen (in common with some other writers) thinks that objectionable thick terms and concepts are useful to the thick concepts debate because they are apt to be disagreed over, and thereby demonstrate that the evaluative and non-evaluative features of the meanings of terms can come apart. The thought is this; suppose you say that something is LEWD and in saying that you are clearly indicating disapproval or a negative evaluation. I want to disagree with you, or at least not endorse what you have said. However, this is not because I think the term is misapplied — I can see that its satisfaction conditions are met. My disagreement with your statement is a disagreement that anything should be evaluated *negatively* just for being sexually explicit beyond normally acceptable boundaries (say). The very fact that I can have this thought shows that the non-evaluative and evaluative aspects of the meaning of a thick term can come apart.

In Väyrynen's project, a denial of an objectionable thick term "x is not lewd" is naturally heard as denying that x falls under LEWD (leaving the content of the concept open), not denying that to fall under LEWD is to be 'bad in a certain way'. Väyrynen's position is that this is well explained because the negative evaluation associated with LEWD is an implication. How can we test that an evaluation is an implication of an utterance, not part of its meaning? Väyrynen uses the 'projection test'. Debbie Roberts (2015), in her review of the Väyrynen 2013, calls this the 'linchpin' of Väyrynen's positive view.

'Projection: An implication projects if and only if it survives as an utterance implication when the expression that triggers the implication occurs under the syntactic scope of an entailment-cancelling operator.' (Simons et al. 2010, 309 quoted in Väyrynen 2013 p60)

Väyrynen claims that,

'If an evaluation conveyed by literal uses of a thick term or concept satisfies Projection, then that evaluation isn't a semantic entailment of the sentence. Thus Projection provides a way to test whether T-sentences semantically entail global evaluation as good or bad in a certain way.' (ibid p60)

I do not think this achieves his purpose and I will try and explain why. To illustrate projection, Väyrynen's uses the following example:

(A) Anna regrets drinking instant coffee this morning.

This has the implication that:

(1) Anna drank instant coffee this morning.

The expression in (A) that triggers the implication is regrets. The implication (1) is a presupposition. This implication, that Anna drank instant coffee this morning is not cancelled in any of the following examples of modifying operators:

- (a) Does Anna regret drinking instant coffee this morning?
- (b) Anna doesn't regret drinking instant coffee this morning.
- (c) Anna might regret drinking instant coffee this morning.
- (d) If Anna regrets drinking instant coffee this morning, then my dislike of instant coffee is vindicated.

However, the assertion (A) also has the following implication (amongst possible others):

(2) Anna has negative feelings about drinking instant coffee this morning.

This implication, that Anna has negative feelings, is cancelled in all the modifying instances a-d, that is, the presupposition that 'Anna has negative feelings' is cancelled just where 'Anna regrets' is not asserted. This is because 'has negative feelings' is a conceptual or semantic entailment of regrets i.e. it is part of the meaning of regrets.

So, how is projection supposed to work for thick terms? What Väyrynen wants, I take it, is to show that the 'global evaluation conveyed by a literal use of a thick term in a normal context' follows the pattern – projection - exemplified here by 'drinking coffee', not the pattern exemplified by 'has negative feelings'. Projection is a phenomenon of presuppositions, and indeed is taken as the hallmark of presuppositions (Beaver and Geurts 2014). Surprisingly, perhaps, Väyrynen doesn't really spell out the example, so the following is my attempt to set it out.

(B) Madonna's show last night was lewd.

This has the implication:

(3) Madonna's show last night was bad.

Now with some modifying operators:

- (e) Madonna's show last night was not lewd.
- (f) Was Madonna's show last night lewd?

- (g) Madonna's show last night might have been lewd.
- (h) If Madonna's show last night was lewd, this vindicates my dislike of pop concerts.

If something is bad for being LEWD, and if the implication (3) is a presupposition, do we agree that it is not cancelled by any of the operators and projects? The example modification (e) specifically says that the show was not LEWD, in which case surely the implication that the show was bad is cancelled?

If we are in favour of Väyrynen's account, or have been persuaded of the pragmatic view independently of this argument, then we might accept that the meaning of lewd is something like 'sexually explicit to the extent of transgressing acceptable bounds' (again Väyrynen's rough sketch) and it carries with it the presupposition that anything that is sexually explicit to the extent of transgressing acceptable bounds is bad or to be disapproved. In that case we would perhaps accept that all the moderators (e-h) work as follows:

- (e') Madonna's show last night was not sexually explicit to the extent of transgressing acceptable bounds.
- (f') Was Madonna's show last night sexually explicit...etc.?
- (g') Madonna's show last night might have been sexually explicit...etc.
- (h') If Madonna's show last night was sexually explicit...etc., this vindicates my dislike of pop concerts.

We would also conclude that the presupposition that anything that is sexually explicit to the extent of transgressing acceptable bounds is bad or to be disapproved is not cancelled by the moderators in the examples.

However, while this may be a demonstration of projection, I don't see it as an argument against any view that the concept LEWD is evaluative as a matter of content. If I hold that the meaning of the term lewd, or the content of the concept LEWD is something like 'bad in virtue of being sexually explicit to the extent of transgressing acceptable bounds', then I can also say that the moderators suggested in (e) to (h) do not cancel the evaluative stance. This is because they do not alter the meaning of the term or content of the concept. What the moderators address is whether particular items fall under the concept, whether Madonna's show classifies as LEWD or not LEWD.

For the purposes of my claim that thick concepts provide reasons for action, I need to cast doubt on Väyrynen's view that thick concepts are not evaluative, and evaluation can be accounted for outside the content of the concepts. If he is right, and his positive view compelling, then there would be no room for a claim that thick concepts in particular provide reasons for action, because 'evaluative' concepts in effect have non-evaluative concepts with

the same extension. Thick concepts cannot yield an interesting class of facts.

3.6 Debbie Roberts: Irreducible Thickness

Debbie Roberts calls her view the ‘inclusive view’ because it takes from ‘pragmatic’ accounts of thick concepts, such as Väyrynen’s, the idea that thick concepts have only one kind of content, and from ‘content’ views of thick concepts that the content is evaluative. Roberts argues that a concept or property can be evaluative without in any sense indicating, or being linked to, a *thin* evaluative property or concept such as ‘good’ or ‘bad’ (Roberts 2011, 2013). Some concepts are evaluative as a matter of content, and are also unanalysable or irreducible. She calls this view ‘irreducible thickness’. Such a view, she thinks, imposes a number of commitments on the holder, including ultimately the denial that all evaluation is like thin evaluation. This is not an easy position to explain, and some writers find it ‘somewhat dark’. (Roberts 2011 p 490)

I discussed the disentangling argument in section 3.3, and how the two strands of the argument, graspability and shapelessness, have been influential in the thick concepts debate. In her essay *Shapelessness and the Thick* (Roberts 2011), Roberts sets out how reductive analyses can meet the challenges of the disentangling argument, thus pushing those who hold a non-reductivist position, such as herself, into a different strategy. Using the property ‘tactful’ as her example, (and where tactful is the property that something has if it falls under the concept TACTFUL), Roberts begins with the simplest form of reductive analysis (RA).

RA1: x is tactful if and only if x is T (where ‘T’ is a non-evaluative predicate referring to a specific set of non-evaluative properties) and x is good in virtue of being T,

Now, as well as being a reductive analysis, this formulation has other features. There is codifiability, in that there is a specific set of non-evaluative properties that give rise to T. There is also universalisability, in that *any* x that is T and is good in virtue of T, is tactful. A non-reductivist can oppose RA1 with the first strand of the disentangling argument, the claim that an outsider will not be able to use a concept competently without making evaluations like an insider. Roberts says that to counter this graspability point, a reductionist view can accommodate that the reductive pattern is not codifiable and does not universalise, by bringing context into the analysis.

RA2: x is tactful if and only if x is T (where ‘T’ is a non-evaluative predicate referring to a specific set of non-evaluative properties) and x is good in virtue of being T, and whether x is good in virtue of being T *is uncodifiably dependent upon context such that*

universalisability with respect to T is not possible. (my emphasis)

In this case, the plausible intuition that an outsider cannot mimic an insider in the use of the concept is respected. Roberts thinks this formulation is a rough sketch of Simon Blackburn's expressivist account of thick concepts in Blackburn 1992. I outlined this in chapter two (section 2.4). Blackburn argues that some people could intelligibly agree that a man is non-evaluatively *fat*, but deny that the man in question is negatively evaluatively *fat* ↓ because the man in question is Pavarotti.

Alternatively, there is another way for a reductive analysis to avoid the graspability point without bringing in the background context. This proposal is to say that, although there is a specific set of non-evaluative properties, these are not specified in the *meaning* of the thick term (and /or are not part of the content of the concept). This is Elstein and Hurka's proposal, discussed in section 3.4. In this analysis, we can think of 'concepts whose descriptive component defines an area in conceptual space within which admissible good – or right-making properties must be found.' (Elstein and Hurka 2009 p 521 quoted in Roberts 2011)

RA3: x is tactful if and only if x is good and there are properties P, Q, & R (unspecified) of non-evaluative sort T (specified) such that x has properties P, Q and R and P, Q & R makes anything that has them good.

This suggestion is also compatible with the graspability point, in that it allows that is very likely that an 'outsider' would have to become an 'insider' to grasp the concept because otherwise it is just too difficult to detect the set of properties that give rise to the evaluation. Elstein and Hurka are unconcerned about arguing against the claim that uncodifiability is a feature of thick concepts. Codifiability, they say, has no bearing on the distinction between the thick and the thin: if thin evaluations are uncodifiable then, on a reductive position, a thick concept will be uncodifiable because of its thin component. RA3 does explicitly claim universalisability, 'P, Q and R makes anything that has them good'. Elstein and Hurka see this as acceptable and compatible with the graspability claim, because some evaluation has to be made in determining what the specific properties P, Q and R are. As I argued in section 3.4, it seems to me possible to find examples that make the universalisation in Elstein and Hurka's proposal problematic. However, Roberts accepts their reductive analysis as a live possibility.

At this point in Roberts' discussion, she says that for a non-reductionist to escape RA3, she has to deny or reject codifiability and universalisability. However those two rejections leave RA2 in play. Something more is needed. Roberts points us back to the second strand of the disentangling argument, and what John McDowell says in *Non-cognitivism and Rule Following* (McDowell 1981) about shapelessness. Shape is what all instances falling under a

concept have in common – the unifying feature of real resemblance. The claim of shapelessness is that evaluative concepts are non-evaluatively shapeless if the real similarity between different instances of the concept's application is not non-evaluative (it is evaluative). (Roberts 2011 p 505). Supervenience is no threat to shapelessness, because supervenience only requires *differences* at the supervened on level, not *classification* at the supervened on level.

So the defender of irreducible thickness has to maintain shapelessness in addition to non-universalisability, uncodifiability and graspability in order to have a set of claims that are not compatible with Elstein and Hurka's reductive account. The second reductive account - RA2 - is compatible with shapelessness, so to answer this she must also claim that no non-evaluative property or set of properties need always be instantiated. Nonetheless, there is yet another alternative reductive account that is compatible with the non-reductionist's set of claims so far:

RA4: x is tactful if and only if x is of non-evaluative sort T (specified) and x has some properties (unspecified) in virtue of being of non-evaluative sort T and x is good in virtue of having those properties where there is no non-evaluative shape to those properties across a range of instances of the predicate's correct application.

All of the reductive analyses RA1 to RA4 are characterised by the conjunction and followed by (in some formulation) 'x is good in virtue of...'. To escape the fourth reductive analysis, Robert's view is that there is only one place left to go, which is to drop 'x is good in virtue of' and with it any notion of a necessary link to thin evaluation. So, the defender of irreducible thickness must also hold that there can be evaluation without thin evaluation, or, in other words, that a thick evaluative concept does not contain thin evaluation.

Roberts summarises the necessary features of an irreducible thickness view of thick concepts as:

- 1) The concept applies to an object in virtue of it having features of a certain sort, of which we may be able to give a rough purely non-evaluative characterisation but where there is no non-evaluative shape for the concept.
 - 2) The object has features of that certain sort that ground the relevant property and this merits the application.
 - 3) Precisely which features ground the property and thus make the application of the concept appropriate, are determined, in a way that cannot be specified in advance, by evaluation on a case by case basis
- (Roberts 2011 p511)

In her 2013 essay, Roberts calls her position the 'inclusive view' because it shares with

pragmatic views the claim that thick concepts are irreducible and shares with content views that they are evaluative as a matter of content.

‘On the Inclusive view competence with an evaluative concept does not *necessarily* involve making a thin evaluation. Recognising that an action is raucous (or lewd, tacky or crude) is a separate question from whether that behaviour is in that respect good to some extent or bad to some extent; to be approved or disapproved of.’ (Roberts 2013 p 85)

One feature of this view is that irreducible thickness sits nicely with Blackburn’s point, acknowledged by Dancy, that thick concepts exhibit flexibility in regards to being pro or con.

In Roberts’ view an evaluative concept is one that ascribes an evaluative property. She says there is no non-circular way of expressing what evaluative properties are, but there are some marks of evaluative concepts, which are ‘constraints on what can count as an adequate account of what makes a property evaluative’ (ibid p 89). These marks include essential contestability and substantive disagreement. Concepts that are essentially contestable *characteristically* (as opposed to possibly or conceivably) admit substantive disagreement about what falls under them. Substantive disagreement about evaluative concepts cannot be resolved by further empirical investigation, and furthermore, even when all the non-evaluative information in a particular instance is recognised and agreed upon, it remains an *open question* whether something really has the evaluative property.

A further mark of an evaluative concept is its anthropocentricity; it is ‘a concept that is intrinsically linked to human concerns and purposes’. Without this feature, we couldn’t see the evaluative similarity across instances. (This is a different point from McDowell’s insistence on anthropocentricity in McDowell 1987. There McDowell is questioning whether it is possible to have a viewpoint and hence language about anything outside of human experience.) Roberts says more:

‘The suggestion now is that what explains this is that the evaluative properties *themselves* are intrinsically linked to human concerns and purposes; the real definition of an evaluative property includes such a concern. One way that this could be the case would be if these properties were response-dependent properties. But we need not tie the Inclusive view to such an account. Another proposal is that we cash out this intrinsic link to human concerns and purposes in terms of importance or *matter*. *We might do this in terms of how evaluative properties characteristically generate reasons* (though without further contextual information we will be unable to say what reasons are generated or for whom).’ (Roberts 2013 p94 - italics in the last sentence mine).

‘We might do this [cash out the intrinsic link to human concerns] in terms of how evaluative properties characteristically generate reasons.’ This is an interesting suggestion. Rather than explain how evaluative concepts (or properties) can provide reasons, we could

take an instance of a concept figuring in a reason as an indication that it is, or might be, an evaluative concept.

Pekka Väyrynen argues against Roberts' characterisation of 'marks' in a similar vein to his arguments against shapelessness: these may be marks of the evaluative, but are they marks of *only* the evaluative? He argues that *ATHLETE* is not standardly regarded as an evaluative concept, yet there can be substantive disagreement about whether, for example, the horse Secretariat is an athlete. He also argues that 'This has surface reflectance property R, but is it yellow?' is open in just the same way as 'This is pleasant, but is it good?'; (Väyrynen 2015 p 211). Simon Kirchin points out that Väyrynen is responding to Roberts' 'marks' as though they were necessary and sufficient conditions of the evaluative, and it is not Robert's intention or strategy to provide such conditions. (Kirchin 2017 pp135-6 footnote 39). Kirchin himself thinks that it is a misguided strategy to give a definition of the evaluative (see the next section) and terminology is difficult; perhaps rather than 'mark' Roberts could say 'trend' or even 'broad trend'. Väyrynen's objection therefore does not succeed in challenging Roberts' arguments for irreducible thickness. If thick concepts are, then, irreducibly evaluative, is there anything further to be said about the nature of the evaluations involved in the application of thick concepts? On this point we will turn to Simon Kirchin's 'liberal' account of evaluation.

3.7 Simon Kirchin: The Liberal View

Simon Kirchin proposes a 'liberal view' of evaluation (Kirchin 2017). Where a conservative view of evaluation (Kirchin's term) regards evaluation as strictly and entirely a matter of being pro or con in some way,

The liberal view of evaluation views evaluation as encompassing pro and con, but it does not confine the nature of evaluation to such matters. It allows that a concept can be an evaluative concept and yet, on certain instances when it is used (sincerely and straightforwardly), there be no positive or negative stance as part of its use and this not threaten the fact that it is an evaluative concept, either overall or in that particular instance of its use.' (Kirchin 2017 pp 133)

As I have noted in section 3.3, Kirchin accepts that the disentangling and shapelessness arguments, as they have been used in the literature, are not sufficient to secure the non-reductivist position that thick concepts are unitary evaluative concepts. Furthermore, he doesn't think it possible to give a 'definition' of what it is to be evaluative. What he does think is that there is an everyday notion of the evaluative that goes beyond pro and con, and that reflecting on this and thinking about examples supports the liberal view. Kirchin's position is substantially in sympathy with the position of Debbie Roberts outlined in the previous section. One difference is that Kirchin holds that pro and con valence must come in to the evaluative

picture somewhere. As we saw, Roberts does not hold this, but acknowledges that putting pro and con (or the thin) to one side makes her position harder to see.

In order to build up his positive view, Kirchin discusses Gilbert Ryle's account of thick descriptions, where thick descriptions – complex descriptions with some degree of specificity – are required to distinguish actions from each other. (Kirchin 2013 and Kirchin 2017 pp 124-129). Ryle's example is a boy winking in a number of different scenarios. Kirchin's point is that the thick descriptions - they have nothing to do with evaluation - cannot be analysed into component parts that both capture the sense of what is going on in each scenario and capture the similarity between them.

Secondly, Kirchin points out that the 'slogan' derived from Bernard Williams' writing that 'thick concepts are both action-guiding and world-guided' obscures the nuances of Williams' position, and the slogan is not correct if taken as a characterisation of all thick concepts. This is surely right. Williams wrote of thick concepts,

'their application usually involves a certain valuation of the situation, of persons or actions. Moreover, they *usually (although not necessarily directly)* provide reasons for action.' (Williams 1985 p 144 my italics).

Moreover, as Kirchin argues, it is more than a little strange to think of aesthetic and epistemic concepts as necessarily action guiding. What action is even in the offing?

To support his position, Kirchin points out that many words and concepts are not hackneyed and much used examples of the thick, but nonetheless probably are thick. As examples, he gives MACABRE, GOSSIP, CAJOLE, CONTORTED, MEDITATION, WARRANTED, JUSTIFIED. I think CAJOLE is a nice example; to persuade someone to do something but with a suggestion that the persuader is using flattery or some other form of appeal, possibly deceitfully.

Central to Kirchin's position are two claims. The first is that we grasp a concept through seeing it used in a range of instances: 'a concept is an abstraction from words and phrases used in all manner of ways in various contexts' (Kirchin 2017 p127). The second is that concepts are multi-functional: 'concepts may do many things: predominantly guide action and judgement, categorise objects, and express positive and negative stances' (Kirchin 2017 p 134). The multi-functionality is why the liberal view can propose that instances of use of evaluative concepts sometimes suggest pro or con, but sometimes don't. The range of instances of use is why the liberal view has to include pro or con as necessarily sometimes being present. In other words, were we to see a concept used, say, a hundred times and it never suggest approval or disapproval, we would not accept that it is evaluative.

The difficulty with this, as I see it, any concept at all is potentially evaluative. Almost all concepts can be thickly evaluative under Kirchin's liberal view. Providing that I witness a concept's being used to indicate a positive or negative stance on an occasion, I can count it. Could we use this thought in the opposite direction; that is, claim that a concept is non-evaluative providing that it is sometimes used in instances that do not indicate a positive or negative stance?

Surely we don't want to say this. We want to keep hold of the distinction between a concept being essentially or intrinsically evaluative or essentially or intrinsically non-evaluative and a concept being used evaluatively or used non-evaluatively in an instance. It is a key part of expressivist and pragmatic accounts of the thick that non-evaluative concepts are used evaluatively by speakers in utterance. However, another item pointed out by Foot, paradigm evaluative concepts such as 'good' are often used descriptively; a good thief. We want to say, following Foot, that all concepts can be used descriptively and evaluatively. This does not speak to their essential nature.

As regards evaluative flexibility, Kirchin argues that the best version of the non-separationist view is just to allow that the same concept can have instances of being pro, con or neutral. Furthermore, the idea of a single attitude again follows from allowing the separationists to introduce the evaluative / non-evaluative distinction in the first place. As Dancy (1995) claims, concepts can perfectly well be multi-attitude.

3.8 Conclusion

In this chapter, I have discussed more recent literature on thick concepts. I have looked at Elstein and Hurka's proposal for a reductive analysis of thick concepts and not found it convincing. Pekka Värynen has an interesting account in which thick concepts may be non-analysable, but that is because, in most cases, their content is not evaluative at all on his view. However, I don't find Värynen's positive account of evaluative language convincing either. Debbie Roberts and Simon Kirchin respond to Värynen's arguments about the over reliance of non-reductivist accounts on the disentangling arguments of John McDowell. They attempt to give alternative positive accounts of evaluative thickness. I conclude that truth-evaluable thick evaluative propositions are still an available option. In the next chapter I turn to reasons for action.

Chapter 4: Reasons for action and thick concepts

4.1 Introduction

In the previous chapters, I have established that there are some concepts, thick concepts, which are evaluative as a matter of content, have satisfaction conditions, and are *irreducible*, which is to say that they are not analysable into other concepts or components (using component as a cover-all term for anything that is not a concept). How to characterise the evaluative is a matter of philosophical debate. My working notion of an evaluative concept is a concept that has an evaluative *point*, and where that evaluative point can be taken to be, on some occasions and however weakly, good in virtue of CONCEPT, bad in virtue of CONCEPT.

Now I turn to a different claim about thick concepts. As I have already quoted in the introductory chapter, Bernard Williams wrote of thick concepts:

‘their application usually involves a certain valuation of the situation, of persons or actions. Moreover, they *usually (although not necessarily directly) provide reasons for action.*’ (Williams 1985 p 144 my italics).

Writers take this quotation from Williams to state his notion of thick concepts. As we can see, it has two aspects; a claim about valuations and a claim about reasons for action. With regard to the claim about reasons, we should not think that Williams himself claimed or thought that the application of a thick concept *always* provides a reason for action (Kirchin 2012, 2017). If nothing else, it is a strain to think that aesthetic concepts provide reasons for action, that because an object is ELEGANT, I have a reason to do something, and Williams recognised this. On reflection though, this strain perhaps arises because I am thinking that the claim is that if an object is ELEGANT I am somehow prompted or *motivated* to act in some way. The claim could be that, in a situation where I already have an action in view, (perhaps I am buying some furniture) that an object is ELEGANT will make some difference to whether I choose one item rather than another. Different ways to understand the claim will be the topic of this chapter.

Aside from discussions of Williams’ philosophical ideas, much more has been written in the literature about the evaluative (or otherwise) nature of thick concepts, than about the suggestion that thick concepts have some role in providing reasons for actions. When I look at the 2012 volume of essays on Thick Concepts edited by Simon Kirchin, only one essay out of twelve, the one by Eric Wiland, is specifically about reasons, and that essay is still about

Williams (Wiland 2012). Of the others, the essays by Kirchin (referenced in the previous paragraph), Smith, Harcourt and Thomas, and Tiberious refer to, or mention, reasons for actions but do not explore the topic. Värynen (2013) addresses six pages towards the end of his book to the question of normative reasons. (What a *normative* reason is will be explained later in this chapter, but in brief it is a reason that to some extent justifies or favours an action.) As far as Värynen's own positive view goes, although thick terms can be used to express reasons 'what the Pragmatic View denies is that the provision of reasons is a function of the meaning or content of thick terms and concepts.'¹⁷ (Värynen 2013 p 237). Simon Kirchin, in his book on thick evaluation (Kirchin 2015) discusses Williams's views, but does not go further in elaborating his own positive view than saying the following:

My whole, positive view of what it is for a concept to be evaluative is that such concepts have multi-various functions: they encompass the ability to praise and blame, and express positive and negative stances, certainly (and *with that they can motivate and provide reasons for action*), but they can also simply offer specific understandings of how the world is that do not have to cast the world in obviously positive and negative ways. (Kirchin 2015 p136-137 my italics)

In her writings, Debbie Roberts make some connections to debates about reasons, notably in embracing radical holism about thick concepts (Roberts 2015) but again reason-givingness is not a feature of her positive view.¹⁸

Although the connection between thick concepts and reasons for action is a claim put forward by Williams, writers find a tension, indeed inconsistency, in his positions on these two subjects. This is partly owing to Williams' particular view of what it takes for a fact to be a normative reason for action. Ulrike Heuer (2012) explores this topic and formulates different ways that Williams' thick concepts claim could cash out in his terms, all problematic. I will come back to her discussion in section 4.6.

If we stand completely aside from Bernard Williams' philosophical framework, how is a suggestion that thick concepts provide reasons for action to be understood? Is it that a thick evaluative concept provides a reason for action whenever its application arises? Or is it that thick evaluative concepts have some capacity to provide reasons that non-evaluative concepts lack? Before I get to a discussion of any of this, I will lay some groundwork and clarify my understanding of certain terms. The first is what constitutes an action; not every doing of or by an individual, still less every event or happening, is an action. This groundwork is section 4.2. Secondly, there are different positions on what it is for something to be a reason for

¹⁷ I don't think Värynen, is casting other writers as making claims here, rather he is saying (again) that there is nothing special about thick concepts.

¹⁸ Van der Linden and Freeman (2017) take on the subject of thick concepts in relation to business ('Profit and Other Values: Thick Evaluation in Decision Making'). However, as their take on thick evaluation is 'the simultaneous consideration of different values' it does not address the questions I am raising.

action, and which are now current in the literature. To put it another way, there are different questions about agents and actions to which a ‘reason’ is the answer. (Dancy 2000). I will set these out in section 4.3, following a discussion by Maria Alvarez. (2017). In brief, *normative* reasons justify or favour actions. Some writers think that a normative reason must yet have a link to the agent’s motivations to *be* a reason. (Williams’ view of reasons for action is located here.) Others think, on the contrary, that normative reasons stand on their own merits. Reasons for action can also cite the motivation of an agent, or explain the agent’s action. In some cases motivation and explanation may be the same thing, but they might not be. Reasons are facts. I will recapitulate from the introduction to the thesis what it is for a thick concept to apply - how a thick concept can be part of what constitutes a fact and therefore figure in a reason.

One reading of the claim that thick concepts provide reasons for action is that there is some understanding of what it is to be a reason for action according to which, where a thick concept applies, there is always a reason to act. In other words, the application of a thick concept will always stand in a reason relation to an action, even where no action was previously under consideration. I shall label this the ‘Always’ thesis. A second reading of the claim could be that there is some understanding of what it is to be a reason for action in which a thick concept, where it applies, provides a reason for action, or is ‘action-guiding’, whereas a non-evaluative concept does not provide a reason for action or is not action guiding; only thick-concepts are action guiding. I will label this the ‘Only’ thesis.

Always thesis:	thick-concept facts always provide reasons for action
Only thesis:	only thick-concept facts provide reasons for action

In sections 4.4 to 4.8 I will go through the possibilities of understanding what reasons for action are as I give them in section 4.3, and I will discuss whether either reading of the claim about thick concepts — the Always thesis and the Only thesis — has any purchase. There is very little to be said in favour of the Only thesis. The Always thesis requires some argument to support it, and perhaps some restriction on its application, but by the end of this chapter my thought that an action is CRUEL is always a reason for me not to perform that action is still in play.

4.2 Actions

An event *X* is an action if there is some description of *X* under which *X*, is or was done intentionally by an agent. (Davidson 1980). Suppose I tip the drink in my hand into a flowerpot. If this happens because you jiggle my arm or just because I am clumsy and twitch

involuntarily, then, although tipping the drink into the flowerpot is something I do, it is not my *action*. On the other hand, if I do it intentionally, perhaps because I think the drink is something I don't like, then tipping the drink into the flowerpot *is* my action. If I think the drink is vodka, when in fact it is gin, then tipping the gin into the flowerpot is something I do unintentionally, but it is still my *action* because it is identical with the act of tipping the drink into the flowerpot and I do this act intentionally.

The word 'action' can be and is used both of a person's doing of a thing and of the thing done in acting (Hornsby 1980). It is important philosophically to be clear about this distinction, and which one of the 'doing' or the 'thing done' is under discussion in the present case. According to Hornsby, the distinction is one between tokens and types of action, and I will follow her in this. A doing of mine (Hornsby herself reserves the word action just for the doings) is a spatio-temporal event. If you and I both tip our drinks into a flowerpot, then to tip a drink into a flowerpot is the *type* or *kind* of *thing done* by each of us. However our doings are two tokens of action, and can hardly be identical as they are done by two different agents.

Now suppose I tip my drink because I think it is vodka, and vodka always makes me unwell. We can say that my doing is a token of two kinds of action; to tip my drink into a flowerpot and to disguise from my host that I have not consumed the drink. Suppose you tip your drink within view of someone else to annoy or offend her. Then your doing is also a token of two kinds of action, to tip your drink into a flowerpot and to annoy someone, and only one of the kinds of action in your case is common to mine. Note that the second kind of thing done by you, the second description of your doing - to annoy someone - is a description of your action in terms of what it causes, annoyance to another person. In the Davidsonian characterisation of an action in the first paragraph of this section, it is where a second description describes the identical token action as a first description, and that first description shows intentionality, that counts the second description as picking out an action.¹⁹ There is only one doing. I do not do two things when I tip my drink and avoid being unwell.

A token action, a doing, is a spatio-temporal event. Until an event *has* happened, it is not possible to pick it out or refer to it as a token. Therefore in deliberating about whether or how to do something in the future, we are always pondering an action type not an instance or token of action. We cannot know in advance of an action or event happening what the consequences of it will be, and therefore what all the descriptions of it will turn out to be, although we may foresee some of them.

As the examples above illustrate, neither all doings, nor any things done, are actions. The

¹⁹ There is no significance to the order of descriptions here: it could be the second, or third or whatever description that demonstrates the intention.

first example I gave was that I tip my drink accidentally, perhaps because I am clumsy. In this instance, my tipping my drink is a doing but not an action. ‘I am clumsy’ could be a *reason why* I do something in an explanatory sense, and more will be said about this in the next section.²⁰ CLUMSY may be a thick concept; I think Simon Kirchin’s Liberal View would accept it. However, and before I go on to discuss what a reason for action is, in this example of an accidental tipping of the drink into a flower pot example, there is no reason for action, because there is no action. In this discussion I am not concerned with reasons in the sense of explanations for ‘things that happen’, I am only with reasons for actions.

While all actions to be actions must have an agent, a token doing will - by definition - involve a particular agent. A kind of thing done will not. (Two different agents can both tip a drink into a flowerpot.)

4.3 Reasons

What is a reason, and what is the relation between a reason and an action? Firstly, here and most simply, a reason is a *fact* that stands in a reason relation to an agent and an action. I understand a fact to be a true proposition.²¹ An agent can have a propositional attitude, such as belief, towards a proposition. Since an agent can be in error, and she can believe something that is not true and so is not a fact, it helps the discussion to have some term for untrue propositions. (For example, ‘Othello killed Desdemona because she was unfaithful’ but, as Othello ultimately realised, Desdemona was not unfaithful at all.) I shall where necessary call untrue propositions ‘putative facts’.

I will proceed on the basis that any claim that a thick concept, or an evaluative concept, is a reason for action, is shorthand for saying that a thick concept is part of the content of a proposition; the proposition is a thick concept predicated of an object. If the proposition is true, then we have a thick or evaluative fact. (Equally, a non-evaluative concept yields a non-evaluative proposition or non-evaluative fact.) Hence, I will regard the phrase ‘the application of thick concepts’ as interchangeable with ‘thick evaluative propositions’ or ‘thick-concept facts’ or ‘putative thick -concept facts’.²² As I said earlier, this is the framework that I am using:

²⁰ Chris Jay has suggested to me that ‘I am clumsy’ could be a reason for a person to take up exercise classes that improve her balance.

²¹ Some writers, for example Jonathan Dancy, take facts to be states of affairs. For a discussion, and a defence of the true proposition view, see Everson 2002.

²² Deflationary views such as Värynen’s need not object to my set-up. They can say that, as thick concepts are not importantly different from other concepts on their view, thick concepts operate in the same way as non-evaluative concepts do.

- A An atomic proposition is a concept predicated of an object.
- B A fact is a true (atomic) proposition
- C A reason for action is a fact that stands in a reason relation to an agent and an action.

To begin fleshing out what a reason for action is:

‘there seem to be at least three distinct questions about the relation between reasons and actions. There are questions about whether there is a reason that favours someone’s action; questions about what reason motivates someone to act; and also questions about what reasons explain his action.’ (Alvarez 2017 p4)²³

If, for the moment, we combine the second and third of Alvarez’s questions, we may say that there are questions about *reasons for* an agent to perform an action and questions about *reasons why* an agent performs an action. Correspondingly, the answers to the questions are the reasons for and the reasons why. Of course, the reason why I do something may very well be a reason that justifies the doing of it; I go for a walk in the afternoon because the weather is fine. Nonetheless there are significant distinctions of view between and within reasons for and reasons why, which I will discuss in the rest of this section.

4.3.1 Reasons For

A reason for an action, a reason that favours an action, is a *normative* or *justificatory* reason. Something that is not true cannot justify anything, and that a normative reason is a fact is a widely held view, and I will take it to be uncontroversial (see Schroeder 2007). What is much more at issue is what the conditions are for a fact to stand in a reason relation to an agent and an action. Very broadly, we can say that positions fall into two groups (Finlay and Schroeder 2017).

The first group of positions are versions of reasons externalism. External reason theorists hold that whether or not a fact is a normative reason for an action has *nothing* to do with the subjectivity of the agent, - her motivations or desires or rationality. Joseph Raz (1999), writes that the classical approach to understanding the connection between action, reason and value holds:

‘that the central type of human action is intentional action, that intentional action is action for a reason, and that reasons are *facts in virtue of which these actions are good in some respect and to some degree.*’ (Raz 1999 p 23 my italics).

²³ Alvarez puts reason singular in the first two questions and reasons plural in the third question but I don’t think there is anything significant in this.

In a Razian picture, it is that the fact picks out something good (or bad) in the action that secures it as a reason for the agent to perform (or not to perform) the action.

The second group of positions are versions of reasons internalism. Internal reasons theorists maintain some connection to the agent, even for normative or justifying reasons. In other words, there must be some connection of the fact to the agent's motives or desires in order for it, the fact, to be a reason: it is this connection that *makes* it a reason. Bernard Williams' view — which I shall return to in section 4.5 — is of this group. Mark Schroeder's 'Hypotheticalism' is another example:

'For R to be a reason for X to do A is for there to be some p such that *X has a desire whose object is p*, and the truth of R is part of what explains why X's doing A promotes p. (Schroeder 2007 p 59 my italics).

4.3.2 Reasons Why

A reason why an agent acts is a reason that she recognises herself for doing what she does, or a reason that makes her action intelligible to someone else. These may not be the same thing. Following Alvarez, I shall separate them out and call them motivating reasons and explanatory reasons. Some writers use the term explanatory reasons for both.

A motivating reason is the reason in the light of which the agent acts. It is the feature of the situation that she picks out in relation to her intentional action. I understand a motivating reason, a reason why in this sense, to be a proposition that is believed by the agent. This formulation includes both facts and putative facts - instances where the agent is mistaken in thinking a proposition is true. Therefore, for a proposition to be both a normative reason and a motivating reason, it must be *both* true and believed²⁴ by the agent. (Of course, motivating reasons may come in combination: I Hoover this morning because the carpet is dirty *and* because there will be no time to do it later.)

We may think that a natural answer to a question about why an agent acted in such and such a way is to state a desire of hers. 'Why did she eat an ice cream? She ate an ice cream because she wanted to.' However, given that the agent had a *desire* to eat an ice cream, she acted in light of her *belief* in a proposition such as 'this is an ice-cream', I shall say. This position is not universally held. Some writers take the view that motivating reasons are psychological entities such as mental states of the agent - believings, desirings etc., or facts about the agent's mental states (and for a long time this was a dominant view (Alvarez 2017 p 17)).

²⁴ Belief is the most straightforward propositional attitude to work with here. I do not mean to rule out, for example, that an agent acts in the hope that something is the case rather than the belief.

However, there is a principled reason for taking motivating reasons to be facts or putative facts and it is that if we do so, a motivating reason is just the same kind of thing as a normative reason. We want this to be the case because if not, it would not be possible to act for a justified reason. (Dancy 2000) There would be normative reasons that are facts and motivating reasons that are something quite different (e.g. believings). While not all reasons *for* are reasons *why*, and not all reasons *why* are reasons *for*, there is a subset of reasons that are both reasons for and reasons why. This is known as the 'identity thesis'. (Heuer 2004).

Now, it seems straightforwardly the case that reasons that motivate an agent can explain actions. I go for a walk in the afternoon because the weather is fine. That the weather is fine justifies walking. The fine weather motivates, prompts or spurs me to walk. That the weather is fine explains why I walk this afternoon. What we need to remember is that a motivating reason may not be true, the agent may be in error. Something that is not true cannot explain anything, nor can it justify anything. Therefore, sometimes the reason that explains an action is, for example, *the fact that I believe* that the weather is fine *rather than* the content of my belief, *the false proposition* that the weather is fine. (So, this example is a case where the explanatory reason is a psychological fact about the agent.)

There is another way in which an agent can be in error, aside from believing something that is false. She might believe a proposition (which might be true or false) and take it to be a justificatory reason for an action when it is no such thing. For example, she takes an umbrella out with her because the weather is fine and dry. In such a case, 'the weather is fine and dry' will not do as an explanatory reason for taking the umbrella because it is not intelligible to a (rational) observer. Therefore an explanatory reason is 'she takes the umbrella with her because she believes that the weather being fine and dry is a reason to take an umbrella'. Again, the explanatory reason cites the agent's belief rather than the content of her belief.

Even in both of these examples, the explanatory reason cites the motivating reason, that the weather is fine. Sometimes, however, explanatory reasons - reasons that make an action intelligible - do not even do that. For example, Jess goes to the hospital. Jess's action can be explained by her goal - to reassure her father, who is in hospital; by her habits - she always goes on Tuesdays; by her character - she is dutiful; or by citing her motivating reason - her father is in intensive care. (Alvarez 2017). I suggest another example of an explanatory reason here - she goes to the hospital because going to the hospital is KIND or COMPASSIONATE.

4.3.3 Interpreting the claim that thick concepts provide reasons for action

With this brief sketch in place, we can see that a claim that applications of thick concepts usually provide reasons for action could be interpreted as:

- (i) a claim that thick-concept facts are apt to be facts in virtue of which actions are good in some respect and degree and therefore apt to be *reasons for* (on the externalist picture),
- (ii) a claim that thick-concept facts are apt to be facts that link an action to an agent's pre-existing motivations and therefore apt to be *reasons for* (on the internalist picture), or
- (iii) a claim that thick-concept facts, where recognised by agents, are apt to be *motivating for* agents and therefore *reasons why* agents perform actions.
- (iv) a claim that thick-concept facts, are apt to be explanatory of actions and therefore *reasons why* agents perform actions.

For any of these, there is the question of whether 'apt to be' is to be understood as the Always thesis or the Only thesis. I don't want the test of the Always thesis to be that thick concepts *never fail* to provide reasons or that the test of the Only thesis is that non-evaluative concepts *never succeed* in providing reasons for action. The discussion is more broad than that. My aim is to indicate what might be required for any of these claims to hold.

4.4 Thick concepts and normative reasons on an externalist view

I outlined in the preceding section that one answer to the question of what reasons there are for an action, is a fact that favours or justifies the action to some extent. This is a normative reason, a reason *for*. A normative reason must be a fact. A normative reason need not be a conclusive, compelling, all things considered reason. Although the literature ordinarily talks about reasons *for* an action, we can equally well talk about reasons *against* an action, reasons not to do something. Reasons against are still normative reasons.

The possible claims I am considering in this section therefore are:

- thick-concept facts are *always* apt to be facts in virtue of which actions are good in some respect and degree and therefore apt to be *reasons for*
- *only* thick-concept facts are apt to be facts in virtue of which actions are good in some respect and degree and therefore apt to be *reasons for*

The external reasons theorist holds that reasons are facts that stand in reason relations to agents and actions. This in itself does not distinguish it from the internalist view. However, on this view, for a fact to be a reason for an agent to act in a certain way, for it to be a fact that favours (disfavours) an action or justifies doing (or justifies not doing) an action, *simply is* that it is a fact in virtue of which this action is good in some respect and to some degree. It

requires no claim at all that the fact motivates the agent, is known or believed by the agent, or explains why the agent acts thus and so. For example, the fact that to walk into town costs no money is a reason for me to *walk* into town, if I want to go to town. (However, actually I walk into town because I like walking and the weather is fine.) The fact that the coat is waterproof is a reason to choose it, if I want a coat. (However, actually I choose the coat because it is blue and I like blue things.) Facts that stand in reason relations can be facts about the agent (I like walking) facts about the thing done (to walk into town costs no money); facts about the experience of performing the action (to walk in fine weather is pleasant); facts about an object involved in the situation (the coat is waterproof, the person is DESTITUTE).

On a first glance, the externalist take on what secures something as a reason - that it is a fact in virtue of which an action is good in some respect and to some degree - looks similar to the notion of a thick-concept fact. A thick-concept fact is a thick concept predicated of an object to form an atomic proposition, where the proposition is true and the fact thereby tells us that the object is good (or bad) to some extent or to some degree. (That is, if we allow for the sake of this argument, that thick evaluation can be characterised in this way. As I explained in chapter three, Debbie Roberts at least would be cautious about a characterisation that always involves pro or con.) However, even though there is this similarity, the examples I have given above show that it is easy to give examples of non-evaluative facts that justify an action; free, in the sense of costs no money, is not obviously an evaluative concept. So, the Only thesis, that *only* the application of evaluative concepts can provide normative reasons for action, and the application of non-evaluative concepts can't, is very hard to motivate.

The Always thesis is that thick evaluative concepts always provide reasons for action when their application yields a true proposition that is relevant to the action in question. The fact that to walk into town is THRIFTY is a reason for me to walk when I go into town. The reason may well be defeasible — rain may overcome THRIFT — but the reason is there. I think the intuitive appeal of the thought that thick-concept facts provide reasons for action is particularly apparent when considering cases where we avoid actions, for reasons such as that that they are CRUEL, or WICKED, or TACTLESS. This is a point that Bernard Williams makes, and I will return to it in chapter five.

There is a countering argument to this view of the Always thesis. This counter or objection is that, while there is an evaluative fact relevant to the situation, this evaluative fact is a result of a feature of the situation, and it *this* feature that is the real reason, properly understood.²⁵ So for example, I may want to say that the fact that walking into town is THRIFTY is a reason for me to *walk* into town, if I want to go to town. The rejoinder is that walking into town is THRIFTY because it costs no money, and that it costs no money is the *real* reason for the agent

²⁵ Stephen Everson has pressed this objection in conversation with me.

to act thus and so. The fact that to walk into town is *THRIFTY* adds nothing to the reasons that there are. We might say that this evaluative fact is superfluous. If this line of argument is right, then a claim or assumption that an evaluative fact or proposition is more apt to be a reason for action than a non-evaluative fact or proposition does not appear warranted.

I think that there are counterexamples to this line of argument, and if there are any plausible examples here, then at least there is a counter-argument to any blanket denial that thick-concept facts provide (normative) reasons for action. Some of the examples I can think of, as it happens, are of the application of aesthetic thick concepts. Suppose an artist is contemplating putting a dash of paint on to a particular part of her painting and then chooses not to for the reason that a dab of paint just there is or will be *INELEGANT*.²⁶ Is there anything more to say than that? It is not clear that there is any other fact that grounds or explains the fact that the dab of paint will be *INELEGANT*. Similarly, if I am shopping for a coat and I choose a particular coat because it is *CHIC*, is there any other fact or facts that is equivalent to the fact that the coat is *CHIC*? I could, of course, say more about the attributes of the coat that appeal to me — the fabric it is made from or perhaps it has a particular kind of collar — but these are not the same as the fact that it is *CHIC*. Suppose a parent sees one of her children taking a book from another and the second child wails. The parent says to the first child, 'You must give the book back because to take it from her is *MEAN*.' Do we really think that some other fact, perhaps that the second child is wailing, is the reason to return the book rather than the fact that it is *MEAN* to take the book in the first place?

There is another response to the objection that thick-concept facts are superfluous or add nothing. This response is to suggest that there can be different levels of facts and therefore of reasons. Thick-concept facts are higher-level or more general facts, thus allowing that there will be other facts that underpin them, but this does not mean that they are not relevant as reasons. This argument is owing to Jackson and Pettit's writing on programme explanation (1990) and I shall return to discuss it in chapter five.

The claim that thick-concept facts provide reasons for action taken on the Always thesis also connects with a debate in the literature about reasons 'atomism' or 'generalism' and reasons 'holism' or 'particularism'. Reasons atomism is a position which holds that some facts that are reasons are invariable reasons. For example, Crisp (2000) makes this argument for judgements that apply virtue concepts; for any case where we can judge that an action is *CRUEL*, the fact that it is *CRUEL* will be a reason not to perform it. Reasons holism is the view that all reasons must be judged on a case-by-case basis. As I discussed in chapter three, Debbie Roberts thinks one of the upshots of the irreducible thickness that she argues for is a 'radical holism'; no similarities can be assumed. Following Crisp then, (and perhaps contra

²⁶ This example is also due to Stephen Everson.

Roberts as I am presenting her) there might be an argument for the Always thesis, if we can argue that what can be said about virtue concepts generalises to a wider selection of thick concepts. I will return to this discussion in chapter five as well.

I noted in section 4.2, where I outline my understanding of what an action is, that a single action (a token action) can have a number of descriptions. Is it possible that the thick concept claim looks more plausible on some descriptions of an action than others? To say that reasons are not linked to the agent's motivations is not to say that facts about the particular agent will be irrelevant to the reason. Compare the case 'that the person is DESTITUTE is a reason for A to help her' and the case 'that the person is Australian is a reason for A to help her'. We may think that the person's destitution is a reason to act for any agent A who is (by definition) in a position to act. The person's Australian identity is less obvious as a reason, and in accepting it we must be supposing that there are some more details about the situation or the agent that contribute to how the fact makes the action good to some degree. Suppose that the agent A is a member of the Australian embassy's staff in the foreign country where the individual is in need of help. Then, for *that* agent A, the person's being Australian really is a justifying reason to help her. We mustn't assume then, that an external interpretation of normative reasons implies that a normative reason is a reason for *any* agent to act in these circumstances, although sometimes it might be.

I also noted in section 4.2 that actions can often be described in more than one way, and token actions (doings) may count under more than one type of action (thing done). In the example of a DESTITUTE person compared with an Australian person example, it may be tempting to suggest that an evaluative fact provides a reason where the description of the action type is less specific and more general compared with the description of the action type where a non-evaluative fact provides a reason; helping a person is the thing done rather than helping a compatriot person or helping a person where there is a pre-existing obligation. However, these descriptions just reflect how I am now discussing the examples. An external reason theorist will point out that in a real case the question for the agent is whether to help *her*, where 'her' points to a particular person.

Another point to note about normative reasons on an external interpretation is that the possibility of the re-description of a token action in terms of what it causes affects the reasons for or against the action. If I make a cake for a charity bake sale, say, and someone is later very ill because she has a rare allergy to some ingredient in my cake, then the one thing I do, my action, is both to make a cake and to cause someone's illness. The person's allergy is a reason against my baking or giving the cake, although not a defeasible one and if I wasn't to know that a person with an allergy might consume my cake I am not necessarily blameworthy. Some of the reasons for and against a token action, then, are relative to the description of

the action. Since some of the descriptions only become apparent after the event has happened, the implication is that some of the normative reasons for an action, on an external reasons view, could not figure in the practical deliberation beforehand about what to do, even for a rational and competent agent. (The possibility of re-description of actions is not unique to an externalist view. However, on internalist views, which are the subject of the next version, reasons are constrained by some connection to the agent.)

4.5 Thick concepts and normative reasons on an internalist view.

To recapitulate, the claims that I am considering here are:

- thick-concept facts are always apt to be facts in virtue of which actions are good in some respect and degree and therefore apt to be reasons for
- only thick-concept facts are apt to be facts in virtue of which actions are good in some respect and degree and therefore apt to be reasons for

An internal reasons theorist holds that to be a normative reason, a fact must yet have some connection to the agent's motivations. A fact without this connection just can't be a reason. There are many accounts of what this constraint is, and how it is formulated.²⁷ There is an obvious and intuitive cause for concern in the constraint: if the agent's motivations must play some role, then it seems to follow that on an internalist account there is no normative reason for any truly monstrous agent, who lacks any motivation at all for the good of others, to abstain from causing harm. I am not arguing for reasons internalism. However, this claim of a link to the agent offers a avenue of thought not open to external reasons theorists about how or why thick concept facts have a special or interesting role in providing reasons for action. Can we say that where someone uses a particular thick evaluative concept and subscribes to its evaluative point that particular thick concept is part of the dispositions or motivations of that person as an agent?

Bernard Williams' position is a case in point. In his view, normative reasons have a connection to the motivations of the agent. His paper 'Internal and External Reasons' (Williams 1979) has been much discussed. In it he argued that reason statements such as 'A has a reason to phi' or 'there is a reason for A to phi' should be understood on an *internal* interpretation. According to Williams, the simplest model of internal interpretation is 'A has a reason to phi iff A has some desire the satisfaction of which will be served by his phi-ing.' (Williams 1979 p 101). By 'desires' of the agent, Williams understands a 'subjective motivational set S' which can contain 'such things as dispositions of evaluation, patterns of emotional reaction, personal loyalties, and various projects, as they may be abstractly called,

²⁷ For example, Schroeder (2007)

embodying commitments of the-agent' (Williams 1979 p105). The subjective motivational set is very influential in (or 'controls') the practical reasoning of the agent. The simple model must also be refined to allow for the shortcomings of the agent. 'To think that a fact is a reason for an agent to act is not to think it is an explanation of an action that she *actually* performs, but rather it is to think it an explanation of an action that she *would* have performed (or would have been somewhat motivated towards performing) if not for her error or ignorance.' (Finlay 2007). Williams puts it that there must be 'a sound deliberative route' from the agent's motivations to a true statement that she has a reason to phi.

Plausibly, for Williams 'the concept of a 'reason for action' just is the concept of an explanation of action, following Davidson.' (Finlay and Schroeder 2017 p16) 'Internal and External Reasons' has accrued an extensive literature and it is not my purpose in this chapter to wade into that debate and evaluate Williams' arguments for reasons internalism. The question here is whether, on this picture of reasons for action, evaluative facts play any different role from non-evaluative facts, and if so, how.

In chapter two, I discussed a challenge from Simon Blackburn about any position in which the evaluative aspect of thick terms is held to be inseparable from the satisfaction conditions that determine their extension. The challenge is that the inseparability of the evaluative and non-evaluative content of a concept appears to commit me to a certain evaluation whenever the satisfaction conditions apply even though on a given occasion I may not endorse the evaluation. This challenge can be extended (and perhaps with even more force) to the reasons that I have if the application of thick concepts is reason giving: the inseparability of the evaluative and non-evaluative content of a concept appears to commit me to reasons that in a given case I may not otherwise accept. In chapter two I showed how Jonathan Dancy resists this challenge, but for the present purposes let us suppose it is a live concern.

Here is a reminder about some of Williams' particular views about the thick. In chapter three, I discussed the graspability point about evaluations, owing to Williams, and mentioned some of the essential features of Williams's accounts of thick concepts. Thick ethical concepts are both world guided and action guiding. They cannot be separated into component parts. They are conceptually prior to thin concepts, are concepts that are generated in cohesive communities, and are *specific to those communities* (Williams 1985). Such a concept cannot be reliably understood or applied by a person who does not appreciate its evaluative point. In Williams 2006 he makes some more remarks about what it is to be an 'insider' user of a thick concept:

'At a general level, those questions [of the truth of the application of a thick concept and its relevance] cannot be separated, since you cannot make the judgements without having the concept, and you have the concept only if you count such considerations as relevant in

deliberation.’ (Williams 2006 p196)

‘It is important that “having the concept” here means “possessing it as one’s own”, where this implies being disposed to use it in judgements with which one is identified. One can understand such a concept, as an anthropologist does, without satisfying this condition.’ (Williams 2006 p196 footnote 8).

We can call those people who have the concept and are disposed to use it, members of the community or insiders. Those who, at best, can understand how the insiders use the concept are the outsiders.

Now, here is the thought: the reasons internalism position held by Williams sketched above, *together* with the insider / outsider distinction about thick concepts provides the explanation for how thick concepts are action-guiding *and* resists Blackburn’s challenge. Some people - insiders - can both competently apply a thick concept *and* have a disposition to be guided by it within their subjective motivational set. Some outsiders can do the former, but no outsider has the latter.

Bernard Williams himself does not write about the connection between thick concepts and internal reasons in his overarching philosophical position; he does not spell out how they are compatible. However, Ulrike Heuer (2012) tackles this question. The set-up given above, of how the combination of the two positions seemingly answers Blackburn, is owing to her. Heuer concludes that Williams’ claim that thick concepts are action-guiding (while by implication non-evaluative concepts are not), and his internalist picture of reasons are inconsistent.²⁸

²⁸ I mentioned that any theory of reasons that links normative or justifying reasons to agent’s motivations (however it does so) must accept or deal with the consequence that wicked agents do not have the reasons for action that virtuous people do, and I also mentioned that the one essay about reasons in the Thick Concepts volume is by Eric Wiland. Wiland uses this example to sketch some thoughts about the connection between a reason for action and a reasonable person. Wiland argues that REASONABLE is a thick concept. It has satisfaction conditions and an evaluative point. The satisfaction conditions are something about taking all considerations into account, not acting on the emotions of the moment, and fairness. He also uses some reflections on the history of the reasonable person test in law to argue that it is culturally specific, which is another dimension of Williams’ position on thick concepts. He then uses Williams’s example of a man who abuses his wife to show a tension between Williams’s views on internal reasons and thick concepts. According to Williams’s reasons internalism, if there is nothing in the man’s motivational set that is served by his treating his wife more nicely, he has no reason to be nice to her. Wiland says that Williams bites the bullet here on the consequence of his own arguments in this scenario. Against this, Wiland argues that the man’s behaviour does not fall under the thick concept REASONABLE and that a REASONABLE person would behave differently. The upshot is that the man behaves according to the reasons he has, and is also not REASONABLE.

As a challenge to Williams it is not powerful. In defence of Williams we can argue that; by his lights, REASONABLE is an abstract concept not an interesting thick one. Williams thought that we arrive at abstract concepts such as ‘right’ by reflecting on our thick concepts. REASONABLE is abstract rather than thick. Also, the reasonable person test substitutes an actual agent for a theoretical agent of ordinary goodwill and cognitive ability. Williams is interested in actual agents not idealised ones and so presumably would reject this move.

The outline of Heuer's argument, as I understand it, is as follows:

- (1) According to Williams, thick concepts are world-guided and action-guiding, that is thick concepts both have satisfaction conditions and provide reasons for action. (Non-evaluative concepts do not provide reasons for action.)
- (2) A distinction between insider and outsider users of thick ethical concepts enables Williams to maintain his claim that thick concepts are action-guiding and at the same time to deny that an otherwise correct application of a thick concept must be a reason for action for an agent who finds that concept objectionable (thus meeting Blackburn's challenge).
- (3) Reasons internalism is crucial to the distinction between insider and outsider users of concepts because it explains why an insider has a reason to act and an outsider does not, in a case where they both judge a thick concept proposition to be true.
- (4) The 'internal reason claim' — the connection between the agent and the fact — cannot be formulated or understood in any way that both maintains a difference between insiders and outsiders in their use of concepts *and* preserves a claim that thick or evaluative concepts are action-guiding in a way that non-evaluative concepts are not.
- (5) Therefore the totality of Williams' position is caught in a dilemma. The thick concepts and internal reasons positions are incompatible. He must accept either that thick concepts are not action-guiding, or that thick concepts in *being* action guiding provide a route for there to be no difference in reasons for two agents even though there is a difference in their actual motivations. If he accepts the second possibility then reasons internalism is false; the agent's motivations are not a constraint on what reasons there are.

The crux of the matter is clearly the internal reason claim issue that I have labelled above as point 4. How is the link between a fact and an agent's motivations to be understood, so that the fact is a normative reason for the agent to perform the action only if the fact is a true proposition whose content is a thick concept *and if* the agent is an insider on that concept, that is a member of a community that possesses the concept. There are, thinks Heuer, three possibilities:

- a) The conditional interpretation: 'an evaluative fact is a reason for a person if and only if and *because* P(T)²⁹ is appropriately linked to her existing motives.' (Heuer 2012 p230)
- b) The belief / desire model of reasons: 'reasons are combinations of (true) beliefs and desires.' (Heuer 2012 P231). This is the broadly Humean model of reasons.
- c) Motivating beliefs: 'an evaluative fact provides a person with a reason because - if she knows it - it also supplies a suitable motivation. P(T) beliefs are motivating for the person who holds them whereas beliefs whose content is merely descriptive are not.' (ibid)

²⁹ P(T): a proposition whose content is a thick concept

The first (a) is the formulation that respects the 'explanatory dimension' of William's account of reasons; a reason must be capable of explaining how an agent acts. The second (b) is the broadly Humean model of reasons.

The difficulty for the thick concept claim in understanding internal reasons on the belief / desire model (b) is fairly obvious. There is no principled restriction on the fact believed by the agent, beyond that it combines with the relevant desire. There is no need for the fact to be evaluative, or for the proposition to make use of a thick concept. If my existing motives and desires include a wish for a new coat and a liking for yellow things, then, when considering a particular coat in a shop, the non-evaluative fact that the coat is yellow and the non-evaluative fact that the coat is my size will be reasons to buy it as much as the evaluative fact that the coat is stylish will be a reason. Thick concepts are not special in their ability to provide reasons for action. Heuer argues against the conditional interpretation of an internal reason claim along the same lines (although she returns to it towards the end of her essay). These are arguments against a claim that *only* thick concepts are action-guiding

In considering the internal reason claim as a motivating belief, option c, it is important to remember that, according to Williams, an outsider can still believe that P(T). She may be in the anthropologist's position. However, Heuer points out,

'If the belief as such is motivating, how can it fail to motivate outsiders? ...If the outsider can acquire and ascertain the same belief as the member, why would it be motivating only for one of them? The answer seems to be: because only one of them is disposed to act accordingly. But if so, the disposition to act must be independent of the belief - but necessary for having a reason,' (p232).

This is an argument against reasons internalism.

Moving away from Heuer's discussion, what if we reject the claim that an outsider can have an evaluative concept, insofar as she understands when it applies, without fully possessing it because she does not use the concept for judgements herself? In other words, we push back at the capabilities of the anthropologist. We might say that someone either has or understands a concept — any concept — or she doesn't. I imagine that there are a huge number of concepts that have never come my way, and I am just incapable of understanding when they apply, or forming beliefs with those concepts as contents. In that case, we can say - on an internalist picture - that the belief of a proposition cannot link up with existing motivations for an outsider, because she just does not understand the proposition and it would never be explanatory of her acting. This, I think, could be a turn that an internalist view could take.

It is unlikely to be consistent with Williams' picture though. Remember the counterfactual idealisation aspect of Williams' reasons internalism position pointed out by Finlay in the quotation in the third paragraph of this section. We can allow for error and ignorance in the process of practical reasoning. The only constraint is the agent's *actual* 'subjective motivational set'. On Williams' view, the difference between two different agents, when faced with the same facts, must be a difference in their motivations. So we are left needing to explain how the possession or otherwise of a thick concept makes a difference to the agent's motivations.

The upshot is, I think, that if thick concepts themselves form part of the 'subjective motivational set' of the agent, their role³⁰ has to be thought of as something other than simple motivation. Is this coherent? Heuer's first option, the conditional interpretation, was expressed as X is cruel is a reason for agent A, if and only if and because the concept CRUEL is part of A's subjective motivational set. How does this not amount to having a desire e.g. a desire to avoid cruel things, and thus be indistinguishable from the belief / desire model? My suggestion is that a thick concept such as CRUEL is part of the agent's 'subjective motivational set' if the agent has a disposition to find CRUEL salient, or 'relevant in deliberation'. This, I think, is not the same thing as having a desire. It is also a suggestion that Williams himself makes, particularly with regard to considerations that prevent an agent from performing an action that she would otherwise perform. (Williams 2006).

I will return to this discussion picking out the most salient reasons in Chapter 5. This suggestion meets the Always thesis because, if a thick concept is part of the motivational set, and an action is under consideration, and the thick concepts applies, then there will always be a reason, albeit a defeasible one. It could meet the Only thesis if there were some account of how thick concepts have some special link to an agent's motivations that no other concept does. Non-evaluative facts can be reasons, but under the desire / belief model of a reason claim. (I don't see why there should not be a plurality of models). None of this, of course, is an argument for internalism in the first place. In fact my argument will be that in a restricted range of applications, thick concept facts can be normative reasons in the externalist sense.

4.6 Thick concepts and motivating reasons

In this section, the claims that I am considering are:

- thick-concept facts, where recognised by agents, are apt to be always *motivating* for agents and therefore *reasons why* agents perform actions.

³⁰ McDowell (1995) thinks that the subtlety of the SMS is a strength of Williams' position.

- only thick-concept facts, where recognised by agents, are apt to be motivating for agents and therefore reasons why agents perform actions.

In section 4.3, I said that a motivating reason, a consideration in light of which the agent acts, is a proposition believed by the agent. If the proposition is true, then it is a fact, and if false then (here) it is a putative fact. This understanding of a motivating reason rules out the possibility that the motivating reason is a desire or goal of the agent, although we may find it natural to give some aim or want of an agent as the reason for doing something. The stipulation, however, secures that a motivating reason is just *the same kind of thing* as a normative reason and we want this to be the case because, if not, it would not be possible to act for a justified reason, which is a troubling result. (Dancy 2000, Heuer 2004)

The consequence of that understanding for the discussion here is that, if a claim that thick-concept facts provide *normative* reasons can be defended then it would follow that a subset of motivating reasons or considerations — those where the agent acts in the light of getting the reasons rights — are thick-concept facts as well. As we have seen in sections 4.4 and 4.5, there is no compelling argument to hold that all normative reasons are thick concept facts, but I am holding on to the possibility that thick-concept facts always provide reasons, although possibly ones that are easily defeasible or outweighed.

The other possibility, and the questions for this section is whether, of the propositions believed by an agent, propositions whose contents are thick concepts are apt to be the ones in the light of which she acts, whether or not she is correct that the proposition is true.

Suppose I like ice cream and it is a fact that there is an ice cream available to me. This fact is a reason to eat the ice cream, one would normally think, although not an indefeasible one. Suppose also that my eating of the ice-cream is not GREEDY, the ice-cream is not DELICIOUS, there is just no evaluative fact or proposition to be entertained about an action of eating the ice-cream or the ice-cream itself. If the claim that thick concepts are apt to be motivating reasons is understood on the *Only* thesis, that is that *only* beliefs of thick concept propositions can be motivating reasons, then I won't eat the ice cream. I won't eat the ice cream because I don't have a motivating reason at all, in the absence of a thick-concept fact. This really does not sound right or intuitively plausible. I have a general desire or pro-attitude towards ice cream in this example. Everyday in multiple small ways, we just do the things we want to do, unless there are weightier considerations against our so acting.

Now suppose, in the example, that eating the ice cream is GREEDY. This evaluative fact is a reason not to eat it. If the claim that thick concepts are apt to be motivating reasons is understood on the *Always* thesis, that is, that a belief of a thick concept proposition is *always*

a motivating reason, then I won't eat the ice cream. I won't eat the ice-cream because, of all the beliefs available to me about the situation, the evaluative proposition about my acting is this way is the consideration that will strike me as most important. If I have another evaluative belief that favours eating the ice cream, for example that the ice cream is amazingly DELICIOUS and this is my only chance, then I weigh the considerations up. This reading of the thick concept claim is, to my mind, more promising.

I have noted that writers who hold that thick concepts are irreducible and evaluative appear to accept that thick concepts are reason-giving. For example, in the introduction to this chapter I quoted Simon Kirchin as saying that '[thick concepts] encompass the ability to praise and blame, and express positive and negative stances, certainly (*and with that they can motivate and provide reasons for action*)' (Kirchin 2015 p136-137 my italics). At the end of section 2.4 we had the suggestion from Debbie Roberts that it is only if the evaluative point of view required to grasp an evaluative concept is equated with the virtuous person's point of view and therefore the entanglement of the agent's will (on McDowell's interpretation of the virtuous person), that the ability to see whether an evaluative concept applies in a certain situation must also engender a pro or con attitude to an action. My thought is that the claim that thick concepts provide reasons for action is easiest to accept as a possibility if it is seen as statement of what might matter to the agent.

I have taken the position that a motivating reason must be capable of being a normative reason. One upshot of the external reasons theorist view in section 4.5 - or one possibility from it - was that thick concept facts are *never* normative reasons because there is always a more basic reason that underpins the thick concept or evaluative fact. If that turned out to be correct, then in any instance where I, for example, act because the action is KIND or help someone because she is VULNERABLE (taking both kind and vulnerable to be thick evaluative concepts) then my motivating reason is not a normative reason, even though on the face of it the reason picks out some good-making feature of the action. I find this result really odd, and as such a consideration against the position that thick concepts do not function as normative reasons as the true normative reasons are always more basic.

Two more points to note. Firstly, a 'thick concepts are always motivating reasons' claim looks somewhat like the *saliency* claim that I considered in section 4.5 when looking at reasons internalism. The difference between 'always motivating' and 'saliency', if there is one, is the difference between a fact always being acted on unless outweighed, and always being relevant. Secondly, of my example evaluative facts in the ice-cream example, one, 'the ice-cream is delicious' is a fact about an item in the situation, the ice-cream. The other fact 'my eating the ice-cream is greedy' is an evaluative fact about the action itself (or the tokening of

the action that I have under consideration). I shall be returning to this point in Chapter 5.

4.7 Thick concepts and explanatory reasons

Finally, there is another possibility for a *reason why* claim, which is that thick concepts are apt to *explain* agents' actions. In this section, the claims that I am considering are:

- thick-concept facts are apt to be always explanatory of actions by agents and therefore *reasons why* agents perform actions.
- only thick-concept facts are apt to be explanatory of actions by agents and therefore reasons why agents perform actions.

By explain an action, I mean that a reason makes the action intelligible to someone else, either by citing the agent's reason in some form or other, in which case there is a clear link to a motivating reason, or by relating some other aspect of the situation. As I have already noted in section 4.3, an explanation must be true, because something that is untrue cannot explain anything, and so explanatory reasons must be facts, not putative facts. The consequence is that, where the agent is in error, an explanatory reason may be a psychological fact about the agent, for example, what she believes. Whether or not the agent is in error, a rendering of intelligibility may be to relate the agent to the action without directly stating her motivating reason. She goes to the hospital because she is a dutiful daughter. I will call explanatory reasons for action that cite motivating reasons (either directly or as the contents of the agent's mental state) Explanatory 1 or E1 reasons, and I will call explanatory reasons for action that do not do this but are some other fact Explanatory 2 or E2 reasons. (I do not mean to exclude at this point that a fact may be both an E1 and an E2 reason).

Where an explanatory reason cites a motivating reason, an E1 reason, the issues for understanding a thick concept claim are the same as for motivating reasons as discussed in section 4.6. Where the fact that stands in a reason relation to the agent and action does *not* cite a motivating reason, an E2 reason, there may be more to say here.

First of all, there is no obvious plausibility to the thought that all E2 reasons are thick concept or evaluative facts. In Alvarez's example of 'Jess goes to the hospital', introduced in section 4.3, the explanatory reason of 'because she always goes on Tuesdays' is not an evaluative fact. The Only thesis for the thick concept claim is easy to dismiss.

What about the thought that, where a thick concept fact applies, it will *always* provide an explanatory reason for action, the 'Always' thesis? This is more hopeful. Another grouping of E2 reasons are facts about character. Jess goes to the hospital because she is DUTIFUL To my ear, concepts that can be applied to character are evaluative, if not always strongly so:

DUTIFUL, SPENDTHRIFT, HONEST, MALICIOUS, CONSIDERATE, RECKLESS. They exemplify features that Debbie Roberts and Simon Kirchin notice about thick concepts - their anthropocentricity. All of these examples might connect agents to actions - that is be relevant to actions - in certain circumstances. There are myriad descriptive concepts we can apply to people - tall, middle-aged, brown-haired, sporty - but these are not facts about character, nor are they relevant to action.

We can also rearrange the 'Jess goes to the hospital because she always goes on Tuesdays' explanatory reason. 'Jess goes to the hospital *on Tuesday because* it [going to the hospital on Tuesday] is habitual or routine.' Now perhaps 'habitual' and 'routine' are borderline thick concepts, that is, they will strike some people as candidate thick concepts and some not. The point is that here I have rearranged the form of the reason to be a fact about the agent's *doing*, the particular tokening of the action, rather than a fact about the agent or the agent's habits. A fact about the *agent's doing* can hardly fail to be explanatory, to be an explanatory reason, I suggest. If that is the case, then, we have a situation where the Always thesis goes through, because an evaluative fact about a doing will always be a reason for action in the explanatory sense, because any fact will be. In Chapter 5, I will try and argue further for this claim.

Finally, one more thought; Hilary Putnam has written:

'it is not only that entangled concepts necessarily figure in evaluation; to the extent that people's motivations are significantly influenced by their ethical reasoning, we will need to take account of – and to make descriptive use of – a variety of thick ethical concepts in the description of economically relevant behavior.' (Putnam 2002 p 64)

The suggestion here is that, if we consider a range or number of actions, thick concepts facts are apt to make intelligible common factors in the reasons for those actions, and to guide future actions. This is a view that is compatible with the suggestion that thick concept facts are second order reasons; in any single case of an action there will be a more basic fact or reason that warrants the application of the thick concept.

4.8 Conclusion

In this chapter, I have explored various ways that a claim that thick concepts provide reasons for action may be understood. To do this I have put in place my understanding of what an action is. Not all doings are actions. I have outlined how the question of what reasons for action there are can be a question about what reasons justify or favour the action, reasons for, or it can be a question about what motivates the agent or explains her action, reasons

why.

I have suggested that a claim that thick concepts are apt to provide reasons for action could be understood as a claim that thick concepts are special because where propositions involving thick evaluative concepts are relevant to an action there will always be a (defeasible) reason for action. I called this the Always thesis. Alternatively, a claim that thick concepts are apt to provide reasons for action could be understood as a claim that thick concepts are special because only propositions involving thick evaluative concepts can provide reasons for action and propositions involving non-evaluative concepts do not provide reasons. I called this the Only thesis.

The obvious place to look for the connection between thick concepts and reasons for action is in the work of Bernard Williams, who introduced the claim. The literature on the topic of thick concepts and reasons for action is mainly about Williams' writings. However, writers who have studied Williams's views have shown how his two positions have inconsistencies, perhaps fatal ones and that one or other of his positions would have to be adapted. We cannot just look at Williams' writing and say this is how the claim that thick concepts provide reasons for action works.

I have been through the possibilities of understanding what reasons for action are as I categorised them in section 4.3, and in each case I have considered both readings of the claim about thick concepts — the Always thesis and the Only thesis. There is very little to be said in favour of the Only thesis. This is not because thin evaluative concepts can also be action-guiding, but because it is straightforward to see that non-evaluative facts can also provide reasons.

I remain with my intuition that statements such as the following are easy to accept:

- The fact that it is KIND to walk my friend's dog for her is a reason for me to walk my friend's dog.
- The fact that it is CRUEL to kick the cat is a reason for me not to kick the cat.
- The fact that it is TACTLESS to remind my colleague that she failed her driving test is a reason not to remind my colleague that she failed her driving test.

There are two possible avenues, as I see it, to make good on this intuition. The first is to maintain that thick concept facts are normative reasons for action, where there is an action in view to which they are relevant. This is only a claim about pro-tanto, defeasible reasons. (I am thinking of normative reasons on an externalist view, but I think that everything I am going to say would be available to an internalist position as well.) To maintain this claim, I will need

to deal with the argument that thick-concept facts are grounded by other features of the situation, and these features are the real reasons. The second avenue is to explore the suggestion that part of what is going on, in the most intuitive examples of when thick-concept facts provide reasons for action, such as the examples just given, is that the action or prospective action itself is being evaluated. This might connect to the thought that thick-concept facts are best understood as reasons *why*, in that they highlight the feature of the situation that has particular appeal or salience to the agent.

Chapter 5: Reasons for action and thick concepts: further arguments

5.1 Introduction

In chapter four I showed that it is not straightforward to connect a claim that thick concepts provide reasons for action, or are reason-giving, in a special way with any one of a number of perspectives of what a reason for action is, at least when those views are looked at in a high-level, schematic way. At the end of chapter four, I said that I would carry forward two strands of discussion in this chapter. The first is to maintain that thick-concept facts are normative reasons for action in an externalist sense, where there is an action in view to which they are relevant. Thick-concept facts are always reasons, although they may be (easily) defeasible. The second topic for discussion is the suggestion that what is going on, in the most obvious examples of when thick concepts provide reasons for action - 'I will not kick the cat because *kicking the cat* is *CRUEL*' - is that the action itself is being evaluated. I will link this to a discussion about virtue concepts and whether it is feasible to make the same reasons-for-action claim about all thick concepts. *CRUELTY* and *KINDNESS* are virtue concepts. *TACTFUL* and *MACABRE* are thick concepts but not virtue concepts. Are *CRUELTY*, *KINDNESS*, *TACTFUL* and *MACABRE* all reason-giving in the same way and to the same extent?

The main challenge as I see it to the claim that thick-concept facts are normative reasons for action, is to say that they are superfluous and provide no additional reasons to the other reasons that there are. The thought is that if a thick-concept fact provides a normative reason at all, it is at best a 'higher-level' reason, or an abstraction from one or more of the reasons that there are. The normative reasons that there are in any case are facts specific to the situation that may or may not have evaluative content. If a reason for my saying something complimentary to Thomas is that my so saying will be *GENEROUS*, then, according to this picture, there will be something that underpins the fact that the action is *GENEROUS* — perhaps Thomas has low self-confidence, or has had an upset — and the underpinning fact is the actual normative reason in this case. The fact that the action is *GENEROUS* is not an addition to the reasons for acting, and I can do the right thing without even having the concept of *GENEROSITY*. I recognise that he is in need of a confidence boost that I am able to give, and I respond accordingly by offering the compliment. The fact that so acting is *GENEROUS* may add to the reasons in an explanatory sense, by saying something about the reasons why an agent acts, but that is a separate matter.

Why would I want to resist this? Well one plausible thought is that thin-concept facts or evaluations, such as *X* is right or *X* is good, are not normative reasons in the sense that I am understanding that term (Dancy 2000). They have the appropriate form —they are facts — but can only be an outcome or verdict³¹ of all the other considerations in favour of or against an action. The thought is that there *must* be more to say in favour of an action than just to say that it is a good thing to do. There must be something that supports the judgement or grounds it. To say that thick concept facts are superfluous as normative reasons, that they add nothing to the reasons, is to put them on the same footing as thinly evaluative concepts: that is that if an action satisfies a thick concept there must be something further to say about it, in virtue of which the thick concept applies, and which favours the action. Those of us who think that thick concepts are evaluative and irreducible concepts and therefore interesting should find this result somewhat troubling.

In chapter four (section 4.4), when initially outlining the objection of thick concepts add nothing to the reasons, I suggested some examples where a thick concept applies and yet it is not so clear that an underpinning or grounding fact is available. Suppose I contemplate painting a wall in my house in two different colours and then realise that the combination is GARISH. is there anything more to say than that? If there are any plausible examples here, then at least there is a counter-argument to any blanket denial that thick concepts provide (normative) reasons for action.

In this chapter, I consider whether the charge that thick concepts add nothing to the reasons amounts to irrelevance, that thick-concept facts are not relevant as reasons. I put forward the discussion of Pettit and Jackson about programme explanation, which draws a distinction between relevance and efficacy. Pettit and Jackson argue that properties can be causally relevant without being causally efficacious. I claim that analogously thick-concept facts are *relevant* in providing normative reasons for or against an action although there may never be an instance where there is a thick concept fact and no other fact that speaks in favour of or against an action. This again is enough to show that there is a response to the challenge that thick concept facts are superfluous as reasons. To support this, I will show by examples that concepts are important to recognising a similarity of reasons, and that they function where there are unavoidable constraints on information.

In section 5.4 I return again to Bernard Williams, and some observations he makes in the essay 'Acting as the Virtuous Person Acts' (Williams 2006). Williams remarks that thick concepts are particularly important as negative considerations that prevent an agent from performing an action that she otherwise has reason to do. Imagine that the cat has brought a baby rabbit into the house and is subjecting it to a tortured killing, and so some intervention

³¹ Verdict is my word not Dancy's.

is needed. Kicking the cat might startle it into dropping the rabbit, (and also express my feelings in the moment), but nonetheless it is CRUEL and I'll have to find some other remedy. I think that what is happening in this case is that the action or prospective action is being evaluated. If I am right about this, then this point may do one of two things. On the one hand it may provide an account of why I have a strong intuition that facts such as 'kicking the cat is CRUEL' are normative reasons, but my intuition is mistaken. 'Kicking the cat is CRUEL' is an evaluation and therefore apt to be an explanatory reason, a reason that explains what I, the agent, find salient in the situation. On the other hand it may support my claim that thick-concept facts provide normative reasons for action. Perhaps though, I have to restrict the claim to instances where the application of a thick concept is to the action or prospective action rather than any other feature of the situation, (for example the agent, or the beneficiary of the action).

CRUEL is a virtue concept, as are KIND, GENEROUS, COURAGEOUS, JUST, PRUDENT. These are the most paradigmatic thick concepts, in that all writers accept them as examples for discussion. Is there a possibility that writers who would otherwise reject the reasons for action claim can accept it for a limited range of concepts? In section 5.5 I begin to look at what is distinctive about virtue concepts. I start with briefly explaining the notion of the virtuous person and then look at a discussion by Williams on how other considerations — including of course the application of thick concepts — relate to the virtues.

In section 5.5 I extend the discussion to the claim that, where a virtue concept applies it will always be a reason, because virtues are 'ultimate reasons' (Crisp 2000). It turns out that virtue concepts are the examples of thick concepts where my claim that thick concepts provide reasons for action is most vulnerable to the objection that they do not add to the reasons.

In section 5.6 I conclude.

5.2 Programme explanation

The challenge to thick concept facts as reasons for action, as I see it, is to say that they are superfluous. The thought is that if a thick concept fact provides a normative reason, it is at best a 'higher-level' reason, or an abstraction from a normative reason, where the real normative reason in the case is at a 'lower level' or is more specific. Put like this, the problem bears some structural similarity to the problem of causal relevance addressed by Frank

Jackson and Philip Pettit (1990) in their paper 'Program Explanation: A General Perspective'.³² The claim that they argue for is that the causal relevance of a property does not require its causal efficacy; a 'higher level' property can be causally relevant without being causally efficacious. This discussion is clearly about a different branch of philosophy than ethics or language. Alexander Miller (2003) and Mark Nelson (2006) develop the discussion of programme explanation to moral properties.

Jackson and Pettit present three background assumptions:

1. A causal explanation of something must direct us to a causally relevant property as opposed to a causally irrelevant property of the factor it identifies as explanatory: a property relevant to the causal production of the effect explained.

2. One way in which properties are causally relevant is by being causally efficacious. A causally efficacious property with regard to an effect is a property in virtue of whose instantiation, at least in part, the effect occurs; the instance of the property helps to produce the effect and does so because it is an instance of that property.

3. A property F is not causally efficacious in the production of an effect e if these three conditions are fulfilled together.

(i) there is a distinct property G such that F is efficacious in the production of e only if G is efficacious in its production;

(ii) the F-instance does not help to produce the G-instance in the sense in which the G-instance, if G is efficacious, helps to produce e; they are not sequential causal factors;

(iii) the F-instance does not combine with the G-instance, directly or via further effects, to help in the same sense to produce e (nor of course, vice versa): they are not coordinate causal factors (Jackson and Pettit 1990 p108)

The problem arises when these background assumptions are combined with a 'devastating' fourth proposition:

4. The only way for a property to be causally relevant to the production of a certain effect is by being causally efficacious in the process of production. (Ibid p111)

It follows from the combination of all four that the property F is not causally relevant and must be causally irrelevant.

Consequently, to claim that a higher-level property F is causally relevant where the lower-level property G alone is causally efficacious, requires some other account of causal relevance that the F property satisfies. To provide such an account, Jackson and Pettit make use of the possibility of multiple realisability. Using the operation of a computer as a metaphor, the suggestion is that the property F *programmes* for the property G. The property F is causally

³² I use the spelling 'program' when quoting Jackson and Pettit but otherwise the UK English spelling 'programme'.

relevant if its instantiation ensures that there is some instantiation of the lower level property G, or some other lower level property, that is efficacious. Pettit and Jackson give three examples to illustrate their idea, one of which is that water in a closed glass container boils and the container cracks. Boiling is a higher level, temperature property that is being instantiated by the water. Boiling is also an aggregate of a lower level property instantiated by the molecules that make up the water, which is that they are moving or have momentum. The efficacious cause of this particular event of the container cracking is that certain molecules of the water, say, A, B and C, are moving and collide with the container with sufficient impact that they cause a fissure in the glass. (This is put simply, the writers say — there will be some more detailed technical explanation.) The important point is that it need not have been molecules A, B and C that caused the cracking of the glass container. It could have been molecules D, E and F, or molecules G, H and I. The fact that the water is boiling is causally relevant to the cracking of the container, because it encompasses what *could* have happened in any relevantly similar situation, as well as what actually *did*.

Jackson and Pettit see this result as significant because without it, they say, causal explanations can only be given at the most detailed physical level. Since almost every explanation will be grounded by another more specific explanation, we have to keep moving down the levels to get to the actual causal explanation. This rules out any explanations in any science beyond particle physics, but also has the disturbing result that most people have never come by a genuine explanation of things that happen. If causal efficacy is the only way of being causally relevant, I may have observed that the glass container cracked, and I believe that the boiling water caused the cracking, but I am mistaken in attributing the cause of the cracking to the boiling. The real cause was something I can have no access to with ordinary human perceptive and cognitive faculties, and can only indicate in writing by talking about moving and colliding molecules, but the molecules are still not the base explanation. There is a real practical interest in establishing causal explanation at the level of 'boiling': if we are unable to do so, then it appears that almost all of the things we typically count as 'explanations' do not really count as such.

As I have said, Jackson and Pettit do not extend their discussion to moral philosophy. However, this extension is taken up by Alexander Miller (2003). Miller's question is whether the notion of programme explanation can help a moral realist. A moral realist, holding that moral properties are non-reductive, sui generis, natural, properties wants an account of 'how it can be that irreducible but supervenient facts and properties - whether in science or outside it - can figure in best explanations, explanations whose unavailability would result in some explanatory impoverishment' (Miller 2003 p150).

The argument that Miller is trying to find a way past goes as follows (Nelson 2006):

- moral properties are only real if they figure in the best explanation of our experience;
- moral properties only figure in the best explanation of our experience if they are causally relevant;
- moral properties are only causally relevant if they are causally efficacious;
- moral properties are not causally efficacious;
- therefore moral properties are not real.

Why are moral properties not causally efficacious? Well, to claim that they are efficacious would require some bizarre account of how the wrongness of an action, say, somehow *causes* my judgement that it is wrong. 'Wrongness' does not feature as a property in our fundamental physical theories, so it is hard to see how wrongness could be cited as a cause. Since Jackson and Pettit's programme explanation offers causal relevance without causal efficacy, it breaks the chain of this argument against moral realism.

Miller points out that the program explanation of the container cracking - that the water boiled - adds two pieces of information to the process explanation - that certain molecules collided with the container. One piece of information is the fact that the water was boiling; for all we otherwise know, the molecule momentum might have been to do with something else. The second piece of information is that, in other possible worlds, molecules other than A, B and C might have done the cracking. However, putting it like this, what the programme explanation improves is *knowledge* of the cause and event, and *knowledge* of other possible worlds. Miller therefore claims that none of this helps the moral realist. An epistemic improvement does not and cannot support any ontological claim that moral properties are real. We can see this, Miller says, if we posit an epistemically unlimited being (God) who has access to everything there is to be known. To this being, the programme explanation just does not add anything, it is superfluous.

In response to this result, Nelson (2006) argues that there is another piece of information to be gained from the programme explanation. The programme explanation tells us that, had molecules A, B and C not caused the cracking of the container, other molecules *would* have caused the same effect. An omniscient being could know all there is to know about the actual sequence of events up to the cracking event, and could know all about the other possible worlds. However, absent the F type explanation, the fact that the water is boiling, even an omniscient being cannot have access to this third piece of information. Consequently, thinks Miller, programme explanation does help the moral realist.

5.3 Programme explanation, thick concepts and normative reasons

How does the notion of programme explanation apply to the question of whether thick evaluative concepts can or do provide normative reasons for action? The metaphor in use in programme explanation is that the existence of a higher-level causally relevant property *programmes for* the existence of some lower level more granular causally efficacious property, such that exactly which lower level property is instantiated is subject to multiple realizability — there are a number of possibilities in any given case. Rather than ‘programmes for’ we could say that one property warrants or guarantees the instantiation of another. Jackson and Pettit are not saying that the two properties are alternatives, in the sense that they are both sufficient but neither is necessary providing that the other exists. Both properties co-exist. The claim is that the causally relevant property is still doing some useful work in the causal explanation, even though it is not causally efficacious.

Miller and Nelson’s discussions of moral properties are still concerned with the notion of causal relevance. In the case of thick concepts and reasons for action, we are not concerned with causation. No one at all claims that a fact causes a reason. I also do not propose to ‘read across’ relevance and efficacy from causation to reasons, and argue that there are such notions as ‘reasons relevance’ and ‘reasons efficacy’ or ‘normative relevance’ and ‘normative efficacy’. ‘Efficacy’ suggests that the instantiation of one feature will be sufficient to instantiate the higher-level relevant evaluative fact. We saw in chapter three that evaluative concepts ‘outrun’ the possibility of fixing their extensions, and a feature that instantiates the evaluation in one case may not do so in another. For what grounds evaluative facts then, we have to say that whether or not a feature is efficacious depends on all the other features. Rather, I propose just to embrace aspects of the analogy and consider the possibility that there are facts that stand in reason relations to agents and actions, and those facts are at different ‘levels’ of generalisation or specificity.

In the case of normative reasons for action, the objection or claim under consideration is that thick concept facts are superfluous as reasons for action because they are, at best, merely a description of what some of the other considerations or facts are, and it is those facts that are the reasons for action. (Those other facts may also ground why in the instance in question the thick concept fact applies. The fact that the cat will be hurt if I kick her is a reason not to kick the cat, and is also part of the explanation of why kicking the cat is CRUEL. I will talk more about this below.) The suggestion in response then, following programme explanation, is that the thick-concept fact is in an analogous role to the causally relevant property. The thick-concept fact is still relevant as a reason for action, and is doing some useful work *as* a reason for action, even though the thick-concept proposition is true at the same time as other

propositions are true, and other relevant considerations are available.

One of Pettit and Jackson's concerns was to keep open the possibility of causal explanations at the level of ordinary day-to-day language. That worry does not seem to be alive in the case of reasons. Reasons are facts and facts are true propositions, and, as I discussed in chapter two, all propositions are 'thinkables', the contents or potential contents of thoughts, following McDowell and Hornsby. I will not have a difficulty that the proposition 'kicking the cat hurts her' is less accessible than the proposition 'kicking the cat is cruel'. I will be able to believe (or take some other propositional attitude to) either thinkable.

Perhaps though, that objection is the result of the simplicity of the example. In everyday life, we may act for good reason on a relevant fact, without being able to fully elaborate why the fact is a fact, why the proposition is true. For example, I buy a particular product in a shop rather than an equivalent product, because the packaging of the one product is 'environmentally sustainable' and the packaging of the other product is not. The container may be biodegradable or recyclable, and exactly how that is the case could be further elaborated, perhaps it is made of potato starch. Now, grant that all this is true - it is a fact that the container is environmentally sustainable, that it is biodegradable, and so on. It is straightforward to say that the fact that is my motivating reason, the reason in light of which I acted, is that the container is environmentally sustainable. However, even though this is a fact, the normative reason that favours my action is some more detailed fact that makes the more generalised fact true. The thing is I may have no understanding, or inadequate understanding, of these more detailed facts. So, on the superfluity picture, I do not in this instance act on the normative reason, and indeed may be incapable of acting for normative reasons on many occasions. This result speaks against our intuition that it must be possible to act for the normative reasons that there are.

Now, of course, 'environmentally sustainable', is not a thick concept that I have seen anyone else suggest. I would argue myself, certainly on the 'liberal' view of thick concepts that Kirchin gives that it is one. For the current point at issue, I don't need absolutely to establish that ENVIRONMENTALLY SUSTAINABLE is thick. I just need to suggest that there can be propositions that are true when other propositions are true, and I may be able to take a propositional attitude to one proposition and not the other.

Here is another example. Person A works for Trumpton city council and needs to purchase a piece of equipment. A could buy this piece of equipment from a number of suppliers. The supplier that A chooses to buy from is not the supplier offering the lowest price for the equipment. A normative reason for the choice of supplier - what makes buying from the supplier right or good to some degree - is that the supplier is in Trumpton. We can also say

that a reason for the choice of supplier is that the supplier is *local*. Not only is this important to A, but the procurement rules of her employer require her to do this within certain margins of price difference.³³

We could tell the same story about Person B, who works for Chigley city council and chooses one supplier over another because the chosen supplier is a Chigley business. On one view, A and B have acted for different reasons; one reason is the fact that the supplier is in Trumpton and the other is the fact that the supplier is in Chigley. On another view, they have acted for the same reason and that is the fact that the supplier is local.

Someone who does not want to accept that thick concepts provide reasons could say that this example is a clear illustration of their view. “The supplier is local” does not add to the reasons in either case. What it does is to re-describe the reason, in a way that makes the similarity of the reasons across the two cases apparent, and also helps to make the reason why the agent or agents so acted intelligible. However, perhaps I could push back and say that neither that the supplier is in Trumpton in the one case or that the supplier is in Chigley in the other are normative reasons but are facts that ground the normative reason which is the fact that the supplier is LOCAL. Less strongly, I could argue that the priority of the fact that the supplier is in Trumpton over the fact that the supplier is LOCAL is not established. This, I think is a matter of intuition. Suppose a parent sees one of her children take a toy from another child and the second child wails. The parent says to the first child, ‘You must give the toy back because to take it from her is MEAN.’ Although the parent could say more about why the action is MEAN, the fact that to take the toy is MEAN is the reason not to do it.

Can the kind of causal explanations that Jackson and Pettit are concerned with be called ‘generalist’? Crisp (2000) points out that such explanations are stated in ‘generalist’ terms. Here ‘generalist’ is in contrast to ‘particularist’.

Consider Newtonian mechanics. This theory can provide an explanation of any event in terms of the principle $F = ma$. When some mass accelerates, there *must* be some balance of forces that produces it. This principle is in a sense an ultimate explanatory reason, analogous to the ultimate grounding reasons of which we seek an understanding in moral philosophy. (Crisp 2000 p35)

Crisp’s point is that in the physical explanation case it would be irrational not to accept that there are some principles that can be amplified by further enquiry to explain a particular case. If we don’t accept this, then our position must be that when on different occasions a billiard

³³ This is a real phenomenon. See, for example, Manchester’s work on progressive procurement as reported by the Centre for Local Economic Strategies cles.org.uk

ball collides with another, it is just a 'brute' fact that sometimes the second ball moves and sometime it doesn't.

Crisp's view of a grounding reason is essentially what I have been calling an externalist view of a normative reason.

Behind the notion of justifying reasons and beliefs lies the way things are. Here we find the kind of reasons I shall primarily be discussing, and which I think are of most interest and importance in ethics. I shall call them grounding reasons. In my example, the grounding reason for my giving her the pill, when she is in pain, is that doing so will decrease (her) suffering. It is not a belief I have about my action, but a property of that action. Our question concerns how such properties behave as reasons. (Crisp 2000 pp33-34)

This brings me on to my next topic.

5.4 The evaluation of actions

Here are two passages from Bernard Williams' essay "Acting as the Virtuous Person Acts" (Williams 2006). Williams is discussing Aristotle, and I will come back to more of his discussion in the next section.

'One important sort of consideration involved is that which serves to dismiss alternative courses of action as variously unworthy: thus, a modest person might dismiss a course of action as vainglorious.' (Williams 2006 p193)

'Some of the reasons X will be simply "factual" statements ("it will embarrass her") which might be seen to be true by someone not disposed to count them as reasons in a V-linked way. But it is easy to move out of this area (consider "she needs help"), and many of the most significant reasons X - notably among the negative considerations mentioned before - will be "thick" ethical concepts ("inconsiderate", "disloyal", "shabby"), where the questions of truth and relevance can be separated only in relation to particular cases.' (Williams 2006 p196)

What these passages point out is that, in an instance of practical deliberation about how to act, and having assembled the options for actions that are means to achieve a goal, and the reasons for and against those actions, then some further considerations about so-acting may present themselves. Williams' observation is that these further thoughts about the action often involve thick ethical concepts, and particularly, he says, the considerations that deter acting in a certain way. Suppose I want to go out this evening, and believe that the pub will be a pleasant place to go to for a drink. It's likely that some friends will be there, but if

not I can read a book or watch some TV sport while having a drink. If I tell my housemate that I am going, then she will come with me, and that will make for a less pleasant evening for me, because she will talk a lot about things that don't interest me. Unlike me, she doesn't know people in the neighbourhood and never goes out without company. Practical reasoning tells me that the best available way to achieve my goal of a relaxing evening is to go to the pub without telling my housemate what I am doing. However, such an action is SHABBY and UNGENEROUS. This further consideration may deter me from going to the pub without telling her.

It seems to me that, whether or not the action is SHABBY is a normative reason, the judgement is an evaluation of the action. There is a proposition, and in the proposition the thick concept SHABBY is predicated of the action (where action is the type of action that is being considered). This does not have to be the case for all true propositions that are relevant reasons; "she needs help" is Williams' example above, or "the weather is fine". Whatever a writer's (or reader's) position on thick concepts and evaluation, it is hard, I suggest, not to hear a judgement about an *action* as good in some way or to some degree or bad in some way or to some degree. The action is CONVENIENT, the action is TACTLESS, the action is SELFISH. How could a proposition about an action (remembering that an action is an intentional doing or thing done, not an event or an accidental happening) not be a consideration in favour or against the doing of the action, albeit it may be an un-compelling consideration, and easily defeated?

There could be counterexamples. A sceptic about thick concept could say that the following are not evaluative; 'the action is swift'; 'the action is uncharacteristic'; 'the action is ELEGANT'. I argue that these counterexamples are about the *manner* of performing the action in its physical manifestation, or they are about the agent, not the action. Or, in the case of 'the action is ELEGANT', this could be an occasion when a concept usually regarded as aesthetic is relevant to an action; 'the resolution of the problem is elegant and has saved a world of grief'.

If I am right, then the *intuitive* appeal of the claim that thick concepts are action-guiding can be accounted for by examples in which propositions predicate thick concepts of actions. Evaluations are reasons, where the evaluation is of an action, and it remains to be seen whether the reason is explanatory or normative.

5.5 Virtue concepts and the virtuous person

Virtue theory is a theory of ethics where the 'test' or standard of right action is not

adherence to some rules or principles, as with deontological theories, or the success or otherwise of the outcomes of the action, as with consequentialist theories. In virtue theory, right action is doing what a *virtuous person* would do. It has its origin in Greek philosophy and G.E.M. Anscombe is sometimes credited with raising a renewed interest in the subject with her essay³⁴ 'Modern Moral Philosophy' (Anscombe 1958) In chapter two (section 2.3) I referred to John McDowell's³⁵ discussion of the virtuous person. 'The right thing to do' in a given situation is established *necessarily* by reference to what a 'virtuous' person does or would do in that situation. A *virtuous* person grasps the relevant reasons and this grasp issues in her doing the right thing. (McDowell 1979). Here is David Wiggins, also discussing Aristotle:

The person of real practical wisdom is the one who brings to bear upon a situation the greatest number of genuinely pertinent concerns and genuinely relevant considerations commensurate with the importance of the deliberative context. (Wiggins 1976 reprinted in Wiggins 1998 p233)

If there is no real prospect of an ordinary scientific or simply empirical theory of all action and deliberation as such, then the thing we should look for may be precisely what Aristotle provides — namely, a conceptual framework which we can apply to particular cases, which articulates the reciprocal relations of an agent's concerns and his perception of how things objectively are in the world; and a schema of description *which relates the complex ideal the agent tries in the process of living his life to make real to the form that the world impresses, both by way of opportunity and by way of limitation, upon that ideal.* (Ibid p237)

In his essay 'Acting as the Virtuous Person Acts' (Williams 2006) discusses how this is supposed to work, in terms of specifying what the virtuous person's reasons are. Firstly, we need to establish that the case is an even minimally virtuous action and to rule out that an agent, any agent, does what the virtuous person would do, but for *different* reasons than the virtuous person. Williams' example is of a person sending a donation to a hospital, not because he is LIBERAL, but in order to gain a reputation for LIBERALITY. Hence:

'A (A) A (fully) V [Virtuous] act is what a V person would do, but only if it is done as the V person does such a thing.' (Williams 2006 p 189).

The conditions for this to be met, Williams thinks are the following:

- '(i) The agent knows
- (ii) He does it choosing and choosing for its own sake.
- (iii) He is in a steady, unchangeable state.' (Ibid)

Williams' interest is in the second condition. What does it mean to choose to do

³⁴ See Crisp and Slote (1997) *Virtue Ethics*

³⁵ JM would not have it that reasons are specifiable as Williams has it

something virtuous for its own sake (*di auto* in the Greek)? Does premise (ii) entail that the following is the case:

‘(B) A V person chooses V acts qua V acts.’

(A) and (B) together seem to make the ideas of a V act and a V person depend on each other. There is no problem just in this, but there will be problem if it leads to vacuity. This will be so, if we cannot distinguish one virtue from another, and or the V agent is left with no determinate content to his thoughts. Both of these threaten if (B) and no more is adopted. But (B) is in general false, in a *de dicto*³⁶ sense: courageous people rarely choose acts as courageous, and modest people never choose modest behaviour as modest. Justice is about the only case where it clearly holds.’ (Ibid p 191)

So, Williams asks, what could stand in place of premise (B)?

‘The right answer starts from the idea that *di’ auto* should be read negatively: the agent does the V thing not for the wrong reasons. This is correct, but it needs positive reinforcement, in the form of an account of the kinds of reasons that appropriately go with various Vs...(B) is very roughly right, but only if “qua V” is read *de re*. We say that the agent did the generous (e.g.) thing because it was the generous thing to do, and we understand what this means because we understand what it is about the situation and the action that makes this action in this situation something that would seem to a generous person the appropriate thing to do.’ ...

This does imply that there is a way in which the action seems to the agent the appropriate thing to do: a way in which such an action commends itself in such a situation to a generous person. What this means is that typically there is a kind of reason or consideration present to the agent’s thought that goes with the act’s being of this particular V kind. One important sort of consideration involved is that which serves to dismiss alternative courses of action as variously unworthy: thus, a modest person might dismiss a course of action as vainglorious. A person with a particular V disposition will have a specific repertoire of considerations that operate for or against courses of action. He may include both kinds, positive and negative, under the label “acting for X reasons.” (Ibid p 192 - 193)

With the introduction of the notion ‘acting for X reasons’ we can now have, non-vacuously:

- ‘- A V act is an act done for X reasons
- A V person is disposed to do V acts.’ (Ibid p193)

Here is where the thick concepts come in. They may be some of the most important reasons X that link with particular virtues.

‘One important sort of consideration involved is that which serves to dismiss alternative courses

³⁶ *De dicto* modal statements answer to definitions and *de re* modal statements are concerned with whether specific things have certain properties essentially – Papineau (2012)

of action as variously unworthy: thus, a modest person might dismiss a course of action as vainglorious.’ (Williams 2006 p193)

‘Some of the reasons X will be simply “factual” statements (“it will embarrass her”) which might be seen to be true by someone not disposed to count them as reasons in a V-linked way. But it is easy to move out of this area (consider “she needs help”), and many of the most significant reasons X - notably among the negative considerations mentioned before - will be “thick” ethical concepts (“inconsiderate”, “disloyal”, “shabby”), where the questions of truth and relevance can be separated only in relation to particular cases.’ (Williams 2006 p196)

Williams acknowledges that his account does not apply to the executive virtues of COURAGE and SELF-CONTROL because these are never done for their own sake.

Here we have the idea that we had before, that reasons stack or underpin each other.

5.5 Particularism, generalism and virtue concepts

For another view of the role of virtue concepts in reasons, there is Roger Crisp’s discussion of Jonathan Dancy’s position (Crisp 2000). Dancy has defended ‘particularism’ about reasons and Crisp quotes a passage from Dancy to indicate what this position is:

The leading thought behind particularism is...that the behaviour of a reason...in a new case cannot be predicted from its behaviour elsewhere. The way in which the consideration functions *here* either will or at least may be affected by other considerations here present. So there is no ground for the hope that we can find out here how that consideration functions *in general*...nor for the hope that we can move in any smooth way to how it will function in a different case. (Dancy 1993 p60 quoted Crisp 2000 p32-33).

In this passage, the suggestion that a ‘reason’ is affected by other considerations here present appears to suggest an ‘all things considered’ reason, or the most salient reason after all the considerations have been taken into account. This is in contrast to a defeasible reason or consideration, which is the notion of reason that I am using in trying to defend the claim that thick concepts provide reasons for action. Crisp goes on:

Behind the notion of justifying reasons and beliefs lies the way things are... I shall call them grounding reasons... It (a grounding reason) is not a belief I have about my action, but a property of that action. Our question concerns how such properties behave as reasons. (Crisp 2000 pp33-34)

Dancy’s support for holism and particularism, says Crisp, comes from examples. These examples illustrate that a fact can be a reason for an action or not, depending on what other facts are also true of the situation. I will quote just two of the Dancy examples that Crisp

highlights, and then go on to the point that Crisp makes.

The Book Case 1: My having borrowed a book from you is a reason for my returning it to you.
Case 2: My having borrowed a book from you is not a reason for my returning it to you, since it turns out that you have stolen it from the library.

Hunting Case 1: That an action is fun is a reason for doing it.
Case 2: The fun experienced by hunters is a reason against hunting.

All of these examples have a common factor. To demonstrate it, I must draw a distinction between ultimate and non-ultimate reasons. Take first *The Book*. The fact that my having borrowed a book from you can both sometimes count and sometimes not count in favour of my returning the book to you suggests that we are not dealing here with an ultimate reason, that is, a reason that we can rest satisfied with as grounding the actions in question. To rest satisfied here would be like accepting it as a brute fact that sometimes one billiard ball moves another, and sometimes it does not. Dancy's example is reminiscent of that used by Socrates in the first book of the Republic to show that justice cannot be identified with returning what one has borrowed. If one has borrowed a machete from someone who subsequently goes off her head, it would not be just to return it. One of Socrates's points here is that an action's being just is an ultimate reason in favour of it. If I can demonstrate the justice of my action, you cannot go on to ask whether I have any ground for doing it, for justice is such a ground. Thus, in Dancy's example, I have a reason to return the book to you in the ordinary case because it would be just; but it would be unjust to return a stolen book to the person who stole it. Justice, then, is an ultimate reason. The fact that I have borrowed the book from you, in an ordinary case, can be described as a reason for returning it to you; but this is only because that is what constitutes justice in this particular case. (Crisp 2000 p 37)

Virtue concepts yield ultimate reasons which are not outweighed. In the example of the book, it would be just to return the book in case 1, and unjust to return it in case 2. The application of virtue concepts involves getting things right, and they will not be reversed by other considerations. In the hunting example, it may in normal cases be PRUDENT to do what is fun or pleasurable, but PRUDENCE does not extend to doing things that are CRUEL.

If we take virtue concepts to be some of the most uncontroversial examples of thick concepts (as writers do, albeit that writers go on to disagree about what thick concepts are or whether there are any such concepts), then Crisp's position could be said to be that this group of thick concepts *always* give normative reasons, because they are invariable grounding reasons. However, this result follows because they are correctly applied only in response to all the considerations and features that there are in a particular case. If a virtue concept is correctly applied to an action, then the resulting fact is a not a defeasible reason. To return to my example of whether it is CRUEL to kick the cat when the cat has hold of a baby rabbit, if there is some justification for kicking the cat such that it outweighs other considerations — perhaps it is the only way of saving the rabbit, and the pain caused to the cat will be less than

the suffering of the rabbit — then to kick the cat just is not CRUEL in this instance.

However, the upshot of this position is that the challenge to thick-concept facts as normative reasons as I have presented it goes through. The challenge is that if a thick-concept fact applies there will always be other considerations that are the case and the thick-concept fact in and of itself adds nothing to the reasons. We have to be careful about terminology, as Crisp uses the term 'grounding reason', but his 'ultimate grounding reason' is an all-things-considered reason. His point is that the application of virtue concepts *can only be* all-things-considered reasons.

At this point I think I have to say that virtue concepts are a special class of thick concepts, that don't fit the notion of a normative reason as I am trying to restrict it i.e. to potentially defeasible reasons.

5.6 Conclusion

At the beginning of this chapter, I remained with my intuition that some examples of thick-concept facts supposedly providing reasons for action are plausible. In order to make good on my intuition, at least in some cases, I thought there were two promising lines of discussion. The first was to maintain that thick concept facts are normative reasons for action, where there is an action in view to which they are relevant. To maintain this claim, I needed to deal with the objection that thick-concept facts are grounded by other features of the situation and it is these features are the real reasons. I attempted to meet this objection by suggesting some counter examples, and also by applying the notion of programme explanation to reasons relevance.

The second line of discussion was the suggestion that part of what is going on, in the most intuitive examples of when thick-concept facts provide reasons for action is that the action or prospective action itself is being evaluated. This might connect to the thought that thick-concept facts are best understood as reasons *why*, in that they highlight the feature of the situation that has particular appeal or salience to the agent. I connected this discussion to some thoughts about virtue concepts. It turns out that it can be claimed that virtue concepts, a distinctive sub-set of thick concepts, always provide reasons for action, in that if they apply it can be claimed they are not defeasible. On the other hand they appear to allow the objection I had been attempting to resist.

Chapter 6: Conclusion

The notion of a thick concept is a concept that is held to be *evaluative* but is also more *empirically informative* and specific than ‘thin’ evaluative concepts such as good, bad, right or wrong. Thick concepts are usually introduced by way of examples, and this is partly because there is no generally accepted account of thick concepts, and this has been their main interest in the literature. The term ‘thick concepts’ was introduced by Bernard Williams. Williams claimed that thick concepts, as well as being evaluative, typically provide reasons for action. While Williams’ assertion that ‘they usually...provide reasons for action’ has been quoted in the literature, and the basic claim acknowledged, particularly by writers who agree that thick concepts are evaluative the mechanics of how this is supposed to work have not been extensively examined, apart from research on Williams’ own philosophical views. This thesis has been an attempt at that exploration.

In the second chapter I considered the appeal of expressivist theories of language, particularly the account of Simon Blackburn. Blackburn’s ‘projectivism’ denies that there can be any genuinely thick concepts, and claims to better explain the behaviour of thick terms than a cognitivist account does. Opposing this John McDowell argued that the distinction between fact and value which is essential to Blackburn’s account, and to non-cognitivism more generally, is not available. I find McDowell’s arguments both plausible and persuasive, so I reject Blackburn’s claim that there are no genuinely thick concepts, where evaluation is part of the content of the concept.

Finding the arguments for projectivism wanting was important for the project of this thesis, as it is essential to my project that the application of a thick concept can be truth evaluable, can yield a fact, contra expressivism. This is because the notion of reasons for action that I am using is that reasons are facts. If propositions involving thick concepts cannot be true, cannot be facts, then the claim that thick concepts provide reasons for action falls at the first hurdle.

In the following chapter, I moved on to some more recent literature, which has thick concepts as its subject. One issue taken up in this writing is that, while McDowell’s arguments are effective against the idea that thick concepts are systematically analysable or reducible, this does not secure that they are *evaluative*. This result might also have undermined my project: if ‘thick concepts’ are non-evaluative, what argument could there be that they are reason-giving where (other) non-evaluative concepts are not. I looked at Elstein and Hurka’s proposal for a reductive analysis of thick concepts and thought there were arguments to resist it. Värynen on the other hand, acknowledges that paradigmatic thick concepts may be non-

analysable, but that is because, in most cases, their content is not evaluative at all on his view. However, I didn't find Värnyen's positive account of evaluative language convincing either. Debbie Roberts and Simon Kirchin respond to Värnyen's arguments about the over reliance of non-reductivist accounts on the disentangling arguments of John McDowell. They attempt to give alternative positive accounts of evaluative thickness. I concluded that truth-evaluable thick evaluative propositions are still an available option.

In the fourth chapter, I explored various ways that a claim that thick concepts provide reasons for action may be understood. To do this I put in place my understanding of what an action is. I outlined how the question of what reasons for action there are can be a question about what reasons justify or favour the action, reasons for, or it can be a question about what motivates the agent or explains her action, reasons why.

I suggested that a claim that thick concepts are apt to provide reasons for action could be understood as a claim that thick concepts are special because where propositions involving thick evaluative concepts are relevant to an action there will always be a (defeasible) reason for action. I called this the Always thesis. Alternatively, a claim that thick concepts are apt to provide reasons for action could be understood as a claim that thick concepts are special because only propositions involving thick evaluative concepts can provide reasons for action and propositions involving non-evaluative concepts do not provide reasons.

The obvious place to look for the connection between thick concepts and reasons for action is in the work of Bernard Williams, who introduced the claim. The literature on the topic of thick concepts and reasons for action is mainly about Williams' writings. However, writers who have studied Williams' views have shown how his two positions have inconsistencies, perhaps fatal ones and that one or other of his positions would have to be adapted.

I went through the possibilities of understanding what reasons for action are as I categorised them in section 4.3, and in each case I considered both readings of the claim about thick concepts — the Always thesis and the Only thesis. There is very little to be said in favour of the Only thesis. This is not because thin evaluative concepts can also be action-guiding, but because it is straightforward to see that non-evaluative facts can also provide reasons.

I remained with my intuition that some examples of thick-concept facts supposedly providing reasons for action are plausible. In order to make good on my intuition, at least in some cases, I thought there were two promising lines of discussion. The first was to maintain that thick concept facts are normative reasons for action, where there is an action in view to which they are relevant. To maintain this claim, I needed to deal with the objection that thick-

concept facts are grounded by other features of the situation and it is these features are the real reasons. I attempted to meet this objection by suggesting some counter examples, and also by applying the notion of programme explanation to reasons relevance.

The second line of discussion was the suggestion that part of what is going on, in the most intuitive examples of when thick-concept facts provide reasons for action is that the action or prospective action itself is being evaluated. This might connect to the thought that thick-concept facts are best understood as reasons *why*, in that they highlight the feature of the situation that has particular appeal or salience to the agent. I connected this discussion to some thoughts about virtue concepts. It turns out that it can be claimed that virtue concepts, a distinctive sub-set of thick concepts, always provide reasons for action, in that if they apply it can be claimed they are not defeasible. On the other hand they appear to allow the objection I had been attempting to resist.

References

- Alvarez, M. (2017) 'Reasons for Action: Justification, Motivation, Explanation' in the Stanford Encyclopaedia of Philosophy
- Anscombe, G.E.M. (1958) 'Modern Moral Philosophy' reprinted in Crisp and Slote (1997) *Virtue Ethics*
- Anscombe, G.E.M. (1979), 'Under a Description' NOUS 13
- Beaver, D. and Geurts, B. (2014) "Presupposition", The Stanford Encyclopaedia of Philosophy (Winter 2014 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2014/entries/presupposition/>>.
- Blackburn, S. (1973) 'Moral Realism' reprinted in Blackburn, S (1993) *Essays in Quasi-Realism*, Oxford University Press
- Blackburn, S. (1980) 'Truth, Realism, and the Regulation of Theory' reprinted in Blackburn, S (1993) *Essays in Quasi-Realism*, Oxford University Press
- Blackburn, S. (1981) 'Reply: Rule-Following and Moral Realism' in Holtzman, S and Leich, C Eds. (1981) Wittgenstein: To Follow a Rule
- Blackburn, S. (1984) *Spreading the Word*, Oxford University Press
- Blackburn, S. (1985) 'Supervenience Revisited' reprinted in Blackburn, S (1993) *Essays in Quasi-Realism*, Oxford University Press
- Blackburn, S. (1988) 'Attitudes and Contents' *Ethics* 98 (3) pp 501-517
- Blackburn, S. (1992) 'Through Thick and Thin' *Proceedings of the Aristotelian Society Supplementary Volume* 66 pp 285-299 reprinted in Blackburn, S (2010) *Practical Tortoise Raising: and other philosophical essays*
- Blackburn, S. (2013) 'Disentangling disentangling' in Kirchin, S., (ed.) *Thick concepts*: Oxford University Press.
- Crisp, R. (2000) 'Particularising Particularism' in Hooker, B. and Little, M. (Eds) *Moral Particularism*
- Dancy, J. (1995). 'In Defense of Thick Concepts'. *Midwest Studies In Philosophy* 20(1) pp 263-279.
- Dancy, J. (2000) *Practical Reality*, Oxford University Press

Davidson, D. (1980) *Essays on Actions and Events*

Davis, W. (2014) "Implicature", The Stanford Encyclopaedia of Philosophy (Fall 2014 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2014/entries/implicature/>>.

Eklund, M. (2011). 'What are Thick Concepts?' *Canadian Journal of Philosophy* 41(1) pp 25-49.

Elstein, D. Y. and Hurka, T (2009) 'From Thick to Thin: Two Moral Reduction Plans' *Canadian Journal of Philosophy* 39 (4) pp 515-535

Everson, S. (2009) 'What are reasons for action?' in Sandis, C. (ed.) *New Essays on the Explanation of Action* Palgrave-McMillan

Finlay, S. (2009) 'The Obscurity of Internal Reasons' *Philosophers Imprint* (9,7)

Finlay, S. and Schroeder, M. (2017) 'Reasons for Action: Internal vs. External' in the Stanford Encyclopaedia of Philosophy

Foot, P (1958) 'Moral Beliefs' *Proceedings of the Aristotelian Society* Vol 59 pp 83-104

Gallie, W.B. (1956) 'Essentially Contested Concepts' *Proceedings of the Aristotelian Society*

Gibbard, A. (1990) *Wise Choices, Apt Feelings*, Oxford University Press

Harcourt, E and Thomas, A (2013) 'Thick concepts, analysis and reductionism' in Kirchin, S., ed. *Thick Concepts*: Oxford University Press.

Heuer, U. (2004) 'Reasons for Actions and Desires' *Philosophical Studies* 121 pp 43-63

Heuer, U. (2012) 'Thick Concepts and Internal Reasons' in Heuer, U and Lang, G Eds (2012) *Luck, Value and Commitment: Themes from the Ethics of Bernard Williams*: Oxford University Press

Hornsby, J. (1980) *Actions*. Henley-on-Thames: Routledge and Kegan Paul

Hornsby, J. (1997) 'Truth: the identity theory' *Proceedings of the Aristotelian Society* 97 pp 1-24

Hurley, S.L. (1989) *Natural Reasons* Oxford University Press

Jackson, F. and Pettit, P (1990) 'Program Explanation: A General Perspective' *Analysis* Vol 50 pp 107 - 117

Kirchin, S. (2013a) 'Thick concepts and thick descriptions' in Kirchin, S., ed. *Thick Concepts*:

Oxford University Press.

Kirchin, S. (2017) *Thick Evaluation*

McDowell, J. (1978) 'Are Moral Requirements Hypothetical Imperatives?' reprinted in McDowell, J. (1998) *Mind, Value and Reality*: Harvard University Press

McDowell, J. (1979) 'Virtue and Reason' reprinted in McDowell, J. (1998) *Mind, Value and Reality*: Harvard University Press

McDowell, J. (1981) 'Non-Cognitivism and Rule Following' reprinted in McDowell, J. (1998) *Mind, Value and Reality*: Harvard University Press

McDowell, J. (1985) 'Values and Secondary Qualities' reprinted in McDowell, J. (1998) *Mind, Value and Reality*: Harvard University Press

McDowell, J. (1987) 'Projection and Truth in Ethics' reprinted in McDowell, J. (1998) *Mind, Value and Reality*: Harvard University Press

McDowell, J. (1987) 'Might there be external reasons' reprinted in McDowell, J. (1998) *Mind, Value and Reality*: Harvard University Press

Miller, A. (2003) *An introduction to contemporary metaethics* Polity Press

Nelson, M.T. (2006) Moral realism and program explanation, *Australasian Journal of Philosophy*, 84:3, 417-428

Papineau, D. (2012) *Philosophical Devices* Oxford University Press

Price, H. (2013) *Expressivism, Pragmatism and Representationalism*, Cambridge University Press

Putnam, H. (2002) *The Collapse of the Fact Value Dichotomy*: Harvard University Press

Roberts, D. (2011) 'Shapelessness and the Thick' *Ethics* 121 (3) pp 489-520

Roberts, D. (2013b) 'Thick Concepts' in *Philosophy Compass* 8 (8) pp 677-688

Roberts, D. (2013a) 'It's evaluation, only thicker' in Kirchin, S., ed. *Thick concepts*: Oxford University Press.

Roberts, D. (2015) 'Värynen, Pekka. The Lewd, the Rude and The Nasty: A Study of Thick Concepts in Ethics' (Review) *Ethics* April 2015 pp 910-915

Roberts, D. (2018) 'Why believe in normative supervenience?' in Shafer-Landau, R. (ed) *Studies in Metaethics* Oxford: Oxford University Press

- Schroeder, M (2007) *Slaves of the Passions* Oxford: Oxford University Press
- Strawson, P. F. (2000) 'Concepts and Properties or Predication and Copulation' reprinted in Strawson, P F (2000) *Entity and Identity: And other essays* Oxford: Oxford University Press
- Trueman, R. (2021) *Properties and Propositions: The Metaphysics of Higher-Order Logic*: Cambridge University Press
- Van der Linden, B. and Freeman, R. E. 'Profit and Other Values: Thick Evaluation in Decision Making' *Business Ethics Quarterly* 27:3 (July 2017) pp 353-379
- Väyrynen, P. (2013). *The Lewd, the Rude and the Nasty: A Study of Thick Concepts in Ethics*: Oxford University Press.
- Wiggins, D. (1976) 'Deliberation and practical reason' reprinted in Wiggins, D. (1998). *Needs, Values, Truth: Essays in the Philosophy of Value* 3rd edition: Oxford University Press.
- Wiggins, D. (1984). 'The sense and reference of predicates: a running repair to Frege's formula and a plea for the copula' *The Philosophical Quarterly* 34(136) pp 311-328.
- Wiggins, D. (1987) 'A sensible subjectivism?' reprinted in Wiggins, D. (1998). *Needs, Values, Truth: Essays in the Philosophy of Value* 3rd edition: Oxford University Press.
- Wiggins, D. (1976) 'Truth, Invention and the Meaning of Life' in Wiggins, D. (1998). *Needs, Values, Truth: Essays in the Philosophy of Value* 3rd edition: Oxford University Press.
- Williams, B. A. O. (1979), "Internal and External Reasons", reprinted in *Moral Luck*. Cambridge: Cambridge University Press, 1981: 101–13.
- Williams, B. A. O. (1985). *Ethics and the Limits of Philosophy*. London: Fontana
republished (2011) London: Routledge
- Williams, B. A. O. (2006) 'Acting as the Virtuous Person Acts' in Burnyeat, M (ed) *The Sense of the Past* Princeton University Press 2006
- Wright C. (2003) 'Realism, Anti-Realism, Irrealism, Quasi-Realism' in Wright C (2003) *Saving the differences: essays on themes from Truth and Objectivity*: Harvard University Press
- <https://www.oed.com/view/Entry/55781?redirectedFrom=distribution#eid>

Acknowledgements

I would like to thank everyone at the University of York who has helped my research, in particular my supervisors and thesis advisory panel.