

**ON THE NATURE AND SIGNIFICANCE OF THE DISTINCTION
BETWEEN THICK AND THIN ETHICAL CONCEPTS**

Michael Smith

Bernard Williams famously distinguishes "thin" from "thick" ethical concepts (Williams 1985, pp.140-142, 150-152). Examples of the former include the concepts of *goodness*, *badness*, *rightness*, *wrongness*, and *obligatoriness*. Examples of the latter include not just nice concepts like *justice*, *fairness*, and *kindness*, but also nasty concepts like *chastity*, *slutiness*, and *poofiness*. The nature and significance of this distinction has long been a matter of dispute. Here I focus on two issues concerning the nature of the distinction, issues that bear directly on the distinction's significance.

Imagine someone who says to someone who is evidently heterosexual, 'You're a poof!', or someone who says of those who take no account of the fact that others may be especially sensitive, and so suffer hurt feelings when they engage them in conversation, 'They are being unfair!'. Absent any further details, these claims should be dismissed as misapplications of the concepts in question. Whatever else we might think of those who call people 'poofs', someone's being evidently heterosexual makes it impossible for him to be correctly described as a 'poof'. And while insensitivity to the feelings of others may indicate that someone is unkind or crass, it isn't something in virtue of which he could correctly be called 'unfair'. Such uses of these concepts may of course be justified on other grounds. Perhaps they are being used as metaphors, or they are an attempt to subvert some social practice, or perhaps some other feature of the context makes their use felicitous. On their face, however, such uses are misapplications plain and simple.

Cases like these underscore a deep truth about the application of the thick ethical concepts. Their application requires that those to whom the concepts are applied have certain non-ethical features rather than others. Their application is thus "world-guided", as Williams puts it. The first issue to be addressed is whether this is a *difference* between thick ethical concepts and thin ones. Are thick ethical concepts world-guided in some way in which the thin are not? Answering this question requires us to do two distinct things. We must first spell out what it is about thick ethical concepts that explains their world-guidedness, and we must then ask whether thin ethical concepts are similar to the thick in this respect. To anticipate, it turns out that the very feature of thick ethical concepts that explains the non-ethical constraint on their application is a feature of thin ethical concepts as well. The upshot is that all ethical concepts, even the maximally thin ones, turn out to be a little bit thick. They are all world-guided.

The second issue concerns the order of definition or analysis. Should we analyze the thick ethical concepts in terms of the thin, or the thin in terms of the thick, or should we simultaneously define thick and thin ethical concepts in terms of each other, perhaps along with their relations to other concepts? Perhaps the best-known answer to this question is R.M. Hare's (1952, 1963). Hare thinks that the thick ethical concepts can all be given a two-part definition. Whatever non-ethical features are such that people's being unkind is a matter of their having those non-ethical features—call these features X, Y, Z—Hare thinks that the claim that someone is unkind means that they have X, Y, Z and their having those features is bad. Others, like Williams, seem to think that the order of definition goes in the other direction (Williams 1985, p.130). Someone's being bad (say) means nothing over and above his instantiating certain negative thick ethical concepts: his being unkind, or unfair, or cowardly, or whatever. Again, to anticipate, the correct answer to this question turns out to

be closer to Hare's than to Williams'. Once we have analyses of the maximally thin thick ethical concepts, we can use them to analyse the rest of the thick ethical concepts.

Once we appreciate that and why these two issues have to be resolved in these ways, it will become clear that the significance of the distinction between thick and thin ethical concepts is not so easy to assess. On the one hand, since all ethical concepts are a little bit thick, it follows that there isn't any distinction to be drawn between them, and hence nothing whose significance we can assess. There is only a distinction between the maximally-thin thick ethical concepts and the not-maximally-thin thick ethical concepts. On the other hand, since we can define the rest of the thick ethical concepts in terms of the maximally-thin ones, it also follows that the maximally-thin thick ethical concepts have a significance that is over and above that of the rest of the thick ethical concepts, their significance lying in their definitional priority.

Before confronting these two issues, however, it will be helpful if we step back and settle the answer to a question that lurks in the background of all discussions of thick and thin ethical concepts. What makes a concept an ethical concept, as opposed to a non-ethical concept? Equipped with an answer to this question, it will be much easier to frame the discussion of the two issues we have identified, and to see why proper responses are the ones suggested.

1. *What makes a concept an ethical concept?*

Though no account of what makes a concept an ethical concept will be completely uncontroversial, I suggest that we work with an account that builds on a distinction made by Thomas M. Scanlon. According to Scanlon, there is a distinctive class of mental states, the judgement-sensitive attitudes, which comprises all and only those attitudes "for which reasons can sensibly be asked or offered" (Scanlon 1998, p.20). In terms of this distinction, my suggestion is that a concept is an ethical concept if and only if, if someone believes that that concept is instantiated, then that person is committed to believing that there is a reason for a certain sort of the judgement-sensitive attitude. Let me explain this suggestion in greater detail.

According to Scanlon, judgement-sensitive attitudes are those

...that an ideally rational person would come to have whenever that person judged there to be sufficient reasons for them, and that would, in an ideally rational person, 'extinguish' when that person judged them not to be supported by reasons of the appropriate kind. (Scanlon 1998, p.20)

Reasons "of the appropriate kind" are reasons in the "standard normative sense", which Scanlon explain by way of a paradigm example: the considerations that support the truth of our beliefs. These are the paradigmatic example of reasons in the standard normative sense because, insofar as we are rational, we acquire and maintain our beliefs when we judge there to be considerations that support the truth of what we believe, and we lose beliefs when we judge there to be no such considerations. Beliefs are thus judgement-sensitive attitudes *par excellence*. But belief is not the only judgement-sensitive attitude, or so Scanlon tells us. Intention, desire, hope, fear, admiration, respect, contempt, indignation, and a range of other attitudes as well, are all judgement-sensitive attitudes, as he sees things (Scanlon 1998, pp.20-21). These mental states are all judgement-sensitive attitudes because they too are sensitive to our appreciation of reasons in the standard normative sense for acquiring them, maintaining them, and getting rid of them.

The relevant contrast is with mental states that do not come and go according to our

assessment of the reasons for and against them. Dizziness is an obvious such example, as we don't have reasons in the standard normative sense at all for being dizzy (compare Thomson 2008, p.208). But another less obvious example might be perceptual appearances like its seeming to us that the lines in the Müller-Lyer illusion are different lengths. Such perceptual appearances, like all perceptual appearances, can be more or less faithful to reality—that is, they can get things wrong—but it doesn't follow from this that they are sensitive to the considerations that bear on whether or not they get things wrong. Even when we know that the lines in the Müller-Lyer illusion are the same length, the appearance of their being the different lengths remains stubbornly in place.

What exactly it means for non-belief attitudes like intention, desire, hope, fear, admiration, and the rest to be supported by reasons in the standard normative sense, a sense explained by having us focus on the case of the considerations that support the truth of our beliefs, is a matter of some dispute. Scanlon think that it requires us to posit a *sui generis* *favouring* relation in terms of which we can explain both reasons for beliefs and reasons for these other judgement-sensitive attitudes: reasons are those considerations, whatever they are, that favour our having judgement-sensitive attitudes (Scanlon 1998, p.17). Others think that we can explain reasons for these non-belief states in terms of reasons for belief in particular (Thomson 2008, p.130). More specifically, they think that reasons for non-belief mental states are considerations that support the truth of the belief that the correctness condition of the relevant non-belief mental state obtains. Reasons for desiring are considerations that support the truth of the belief that the thing desired is desirable; reasons for admiring are considerations that support the truth of the belief that the thing admired is admirable; and so on. For present purposes, we do not need to decide between these ways of spelling out the idea of a reason in the standard normative sense. It suffices that some such story can be told.

Moreover, what exactly it means for these any of these judgement-sensitive attitudes to be supported by reasons in the standard normative sense, as distinct from their being supported by reasons in the sense of its being good to acquire them, is also a matter of some dispute. Many think that the best way to explain this difference is to appeal to whatever the correct account of reasons in the standard normative sense is (Parfit 2011, Vol.1, Appendix A). As they see things, we can explain what its being good to have some judgement-sensitive attitude amounts to in terms of the existence of reasons in the standard normative sense for wanting to have that judgement-sensitive attitude. The distinction, in other words, is that between there being reasons in the standard normative sense for having some judgement sensitive attitude and there being reasons in the standard normative sense for wanting to have that judgement-sensitive attitude. For present purposes, I propose that we simply assume that this way of making the distinction is correct. Nothing will turn on this in what follows, however.

With the class of judgment-sensitive attitudes clearly in view, my suggestion is that a concept is an ethical concept if and only if, if someone believes that that concept is instantiated, then that person believes that there is a reason for him to have some *non-belief* judgement-sensitive attitude: a reason for intending, or desiring, or hoping, or fearing, and so on. Given that these non-belief attitudes all seem either to be or to entail a desire of some sort—an intention seems to be, *inter alia*, the desire that the desirer himself performs some action available to him; hope seems to be, *inter alia*, the desire that things turn out a certain way; fear seems to be, *inter alia*, the desire that one or one's associates be out of harm's way; and so on—the suggestion can be simplified. A concept is an ethical concept, I suggest, if and only if, if someone believes that that concept is instantiated, then that person believes that there is a reason for him to desire that the world be a certain way. As the examples help make plain, there is no assumption the world could be that way without standing in a certain

relation to the desirer himself—in other words, the desire may have *de se* content, as well as *de dicto* content—and nor is there any assumption that the option of making it that way is available.

This suggestion should sound at least somewhat familiar. The signature doctrine of moral rationalism is, after all, that facts about what agents morally ought to do are just facts about what they have reasons to do, and hence that beliefs about moral obligations are beliefs about what agents have reasons to do. If, as Scanlon suggests, facts about what agents have reason to do are just facts about what they have reasons in the standard normative sense to intend to do, and if intention is just a kind of desire, then we can generalize the rationalist's signature doctrine as follows (Scanlon 1998, p.21). An ethical fact, whether thick or thin, is just a fact about what there is reason to desire: facts about what's morally obligatory are facts about what there is reason to desire to do of the options available; facts about what's good are facts about what there is reason to desire be the case; facts about what's bad are facts about what there is reason to be averse to, or to desire not be the case; facts about justice and fairness are a subspecies of facts about what there is reason to desire be the case; and so on.

But though I have no objection to this generalization of the signature doctrine of moral rationalism to cover all ethical facts, I don't myself think that it constitutes, all by itself, a story of what makes a concept an ethical concept. Consider, for example, the most straightforward way of turning it into a story of what makes a concept an ethical concept. A concept is an ethical concept, the idea would be, if and only if facts about that concept's instantiation entail the existence of reasons for corresponding desires. The difficulty with this idea is that it doesn't tell us why those ethical concepts that are *uninstantiated* are ethical concepts. Think once again about the concept of *poofiness*. The concept of *poofiness* is an ethical concept, indeed a thick ethical concept, but it isn't the case that facts about this concept's instantiation are facts about the reasons that there are to (say) be averse to people on account of their being homosexual. This isn't the case because there are no such facts. A story about what makes a concept an ethical concept must therefore allow for the possibility that such concepts are not instantiated. Ethical concepts that are not instantiated include all the nasty thick ethical concepts: *poofiness*, *slutiness*, *chastity*, and so on. These concepts are not instantiated for the simple reason that their instantiation would require there to be reasons for desiring that there aren't. The proper theory of these is an error theory. My suggestion, by contrast, is tailor-made to avoid this difficulty. Since what makes a concept an ethical concept is the fact that someone who believes that that concept is instantiated believes that there is a corresponding reason for desiring, the concept of *poofiness* counts as an ethical concept simply in virtue of the fact that someone who believes of certain people that they are poofs thereby believes that there are reasons to be averse to them on account of their being homosexual. The fact that their belief makes a false presupposition about the reasons that there are is neither here nor there.

Note that this suggestion is also similar to, though importantly different from, a similar proposal of Bernard Williams's. Williams tells us that the distinctive feature of thick ethical concepts, and hence by implication thin ethical concepts too, is that they have a general connection with 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. (Williams 1985, p.140)

Williams's proposal should sound familiar, as it is very like that of the moral rationalist's. This should, however, give us pause for thought, as throughout his career Williams was famously opposed to moral rationalism.

As Williams sees things, all reasons for action are what he calls "internal reasons," where an agent has an internal reason to act in some way just in case she would be motivated to act in that way if she were to deliberate correctly (Williams 1981). Correct deliberation, for Williams, is a matter of making sure that the elements in one's "subjective motivational set" (or "S", as he terms it)—we can think of all these as one's desired ends—are determinate; getting these desired ends to cohere with each other in a manner suited to their satisfaction over time; using one's imagination to make sure that desire satisfaction really would be the result of the attempt to satisfy them; and, finally, taking the means to satisfy these desired ends in the light of knowledge of what those means are. What an agent would be motivated to do after deliberating correctly, as Williams conceives of correct deliberation, and hence what she has reason to do, is therefore relative to her S: that is, relative to what her antecedently desired ends are. The contrast with what she morally ought to do thus could not be starker. When we say to a man who is cruel to his wife that he morally ought to treat her better, he does not refute us simply by pointing out something that may already be evident, namely, that he has no desire that would be satisfied by treating her better. But that is precisely how he refutes the claim that he has a reason to treat her better, according to Williams.

Williams's idea that facts about reasons for action are relative to antecedently desired ends also sets claims about reasons apart from claims couched in thick and thin ethical terms. He tells us, for example, that it would be apt to make all of the following claims about a man who is cruel to his wife, some thick and some thin:

...that he is ungrateful, inconsiderate, hard, sexist, nasty, selfish, brutal, and many other disadvantageous things. I shall presumably say, whatever else I say, that it would be better if he were nicer to her. (Williams 1995, p.39)

But as before, if his desired ends are sufficiently horrible and ingrained, Williams admits that, as he sees things, the husband may have no reason to be grateful, or to be considerate, or to treat her better. As he puts it,

I say, 'You have a reason to be nicer to her'. He says, 'What reason?' I say, 'Because she is your wife.' He says—and he is a very hard case—'I don't care. Don't you understand? I really do not care.' I try various things on him, and try to involve him in this business; and I find that he really is a hard case: there is nothing in his motivational set that gives him a reason to be nicer to his wife as things are. (Williams 1995, p.39)

The hard case he imagines is appropriately describable in thick and thin ethical terms, but Williams thinks that he does not have corresponding reasons for action because he would not have corresponding desires if he were to deliberate correctly. Whatever he has in mind when he says that there is a general connection between the application of thick and thin ethical concepts and reasons for action, Williams therefore most certainly isn't suggesting that someone who is appropriately describable in thick or thin ethical terms has corresponding reasons. So what does he have in mind?

Here is the passage in which Williams hints at how he would develop his view if he were to spell it out fully:

It may well be that 'thick' ethical concepts are, to an adequate degree, both 'world-guided' and 'action-guiding'. People who use a given concept of this sort will find their application of it guided by their experience, and also accept that it gives them reasons for or against various kinds of action. Then this disposition will figure in their S, in rather the same way as the disposition to avoid the poisonous or the disgusting may figure in it.

But this does not mean that a speaker who does use a given concept of this kind (*chastity* is an example that focuses the mind) can truly say that another agent who does not use the concept has a reason to avoid or pursue certain courses of action in virtue of that concept's application. To show this, the speaker would need to show that the agent has reason to use that concept, to structure his or her experience in those terms. That is a different and larger matter; all the work remains to be done. (Williams 1995, pp.37-38)

Though Williams does not elaborate further, his idea here seems to be that the role of ethical concepts as such is to make salient to users of those concepts features that bear in some very direct way on what they have reason to do. Just as different people have reasons to act in different ways, depending on their antecedently desired ends, so, according to Williams, different people have reasons to use different ethical concepts, depending on their antecedently desired ends, because different ethical concepts make salient different features of their circumstances. His strategy is to turn this feature of ethical concepts into an account of what makes a concept an ethical concept because of its connection to reasons for action.

Imagine someone who conceptualizes his experience in terms of the concept of *chastity*. Williams's idea seems to be that, for such a person, whether or not women are restraining their sexual behaviour will be a very salient feature of their circumstances. Who would have a reason to make that feature of their circumstances salient? His answer is: those who desire that women restrain their sexual behaviour. Such people have a reason to make *use* the concept of *chastity*, and they have corresponding reasons for action, but all this is consistent with others who have no such desire having no such reasons. The connection that Williams sees between ethical concepts and reasons for action is thus mediated by such reasons as they have to use ethical concepts. To put Williams's idea in our terms, a concept is an ethical concept if and only if those who have reasons to use that concept have corresponding reasons for action, and hence, given the general connection between reasons for action and reasons for intention, if and only if they have corresponding reasons for intention. To return to the example of chastity, what makes this an ethical concept isn't anything about our use of the concept, and it isn't anything about the reasons had by those to whom the concept is applied, but is rather the fact that those who have reasons to use the concept of chastity have corresponding reasons for intention.

Williams's proposal differs from ours in several respects, two of which are worth emphasizing. First of all, unlike ours, Williams's proposal tells us that what makes the concept of chastity an ethical concept is the fact that there are certain people, namely those who have reasons to use that concept, who in fact have reasons for corresponding intentions: reasons to intend to (say) restrain women's sexual activity. Our suggestion has no such upshot. What makes the concept of chastity an ethical concept, according to our suggestion, is the fact that those who believe that certain women are chaste *believe* that they have reasons to have such intentions, not that they really do have such reasons. The second difference is that Williams's proposal tells us that there is a connection between ethical concepts and corresponding reasons for intentions in particular. Our suggestion, by contrast, allows that though certain ethical concepts have a connection with reasons for intention, other ethical

concepts have no such connection, though of course they do have a connection with reasons for desiring that the world be a certain way.

In order to see how this second difference makes a difference, consider some examples. Consider the thin ethical claim that it would be better if, in the natural order of things, not as a result of anyone's intervention, good people flourished and bad people suffered. Or consider the thick ethical claim that it would be fairer if, in the natural order of things, not as a result of anyone's intervention, natural disasters were the burden of all equally, not the special burden of only some. Not only would the truth of such claims not require anyone to have the option of doing something to make things other than the way they are, the truth of such claims rules out the possibility of there being such reasons for intention. This creates a difficulty for Williams's proposal, as it seems incapable of explaining what makes it the case that those who have a reason to use these concepts in this way—the concept of betterness and fairness without any tie to reasons for intentions—are using ethical concepts at all, given that they plainly are not being used to make salient a feature of an agent's circumstances that is relevant to what he has reason to intend to do.

Our suggestion, by contrast, faces no such difficulty. For our suggestion is simply that, if someone believes that it would be better if, in the natural order of things, not as a result of anyone's intervention, good people flourished and bad people suffered, or if he believes that it would be fairer if, in the natural order of things, not as a result of anyone's intervention, natural disasters were the burden of all equally, not the special burden of only some, then he believes that he has reasons to prefer that the world be these ways. The fact that he has no reasons to intend anything is irrelevant.

2. Ethical concepts, moral concepts, and world-guidedness

Before using this account of what makes a concept an ethical concept to give an account of what makes thick ethical concepts world-guided, it will be helpful to remind ourselves of what Hare has to say about the thin ethical concepts. Once we have Hare's account of the thin ethical concepts before us, it will be clear that the thin ethical concepts are world-guided in exactly the same way in which the thick ethical concepts are world-guided. In what follows we will focus on what Hare has to say about the thin ethical concept of an obligation, but there is nothing special about this thin ethical concept. We could have chosen any of the thin ethical concepts to make the point.

Suppose that A says to B, 'You ought to ϕ .' Hare thinks that in saying this, A may or may not be expressing an ethical belief. His basic idea, one that permeates all of his work in ethical theory, is that A expresses an ethical belief only if A's grounds for saying that B ought to ϕ is his commitment to some universal principle (Hare 1952, Hare 1963, Hare 1981). A expresses a non-ethical belief, by contrast, if his grounds for saying that B ought to ϕ has nothing to do with a universal principle, but instead turns essentially on facts about some particular. Imagine, for example, that A says B ought to give him some money. If A's grounds for saying that B ought to give him some money is that his doing so would make A richer—perhaps A is a robber—then A's grounds essentially involve a particular, namely, the effect of B's giving money to A on satisfying A's desire for money. But if A's grounds for saying that B ought to give him some money is that his doing so would replay a debt, then A's grounds do not essentially involve any particular. His claim is rather grounded in the universal principle that people ought to repay their debts.

Adapting this to our account of what makes a concept an ethical concept, we can put the point schematically as follows. If Hare is right, then when A says to B, 'You ought to ϕ ', A expresses a belief about the instantiation of an ethical concept only if:

(i^{universal}) A believes that there is some sub-set, C, of the universal features of the circumstances in which B finds himself; and

(ii^{universal}) A believes that being in C is a reason for anyone who finds himself in those circumstances to desire that he ϕ s.

If these conditions are not met, then A doesn't express an ethical belief at all. Moreover, since Hare thinks that a similar universality condition applies to the other thin ethical concepts as well, it follows that, if he is right, then what makes beliefs about their instantiation ethical beliefs is that versions of (i^{universal}) and (ii^{universal}) are true of them too.

Hare himself thinks that there are further conditions on the application of the thin ethical concepts. For example, he thinks that for a claim about what ought to be done to be an ethical judgement, that judgement must entail commitment to an imperative, and the agent's commitment to that imperative must be overriding (Hare 1981). Whether or not we would take issue with these further conditions on the application of an ethical concept, it is virtually universally agreed that he is right that for a judgement to be an ethical judgement, that judgement must be universalizable. Nor is this surprising. The concept of an ethical claim just seems to be the concept of one that gives no privilege to any particular over any another, and the deeper explanation of this, if our suggestion is along the right lines, is that *reasons* must be universal features. The difference between the robber and someone who thinks B ought to repay his debts, in other words, is that though they both desire B to give them money, only B has that desire *for a reason*.

Assuming that thin ethical concepts really are subject to the universalizability condition, two important and related conclusions follow. The first is that there is a sense in which even the thin ethical concept of an *obligation* is world-guided. The worldly condition takes the form of a constraint on the sorts of considerations that people can believe provide reasons for desires to do things when they think that there are ethical obligations, a constraint that we can specify in non-ethical terms. These considerations, it turns out, must all be *universal* features of the circumstances in which those obligations are had. The concept of universality is not itself an ethical concept, however. Non-universal features of circumstances—features that concern the particulars involved—cannot be believed to be the features that provide reasons for desires. But though there is this constraint on the sorts of considerations that can provide reasons for desires, the constraint doesn't wholly determine what those considerations are. The considerations must be a sub-set of the universal features, but it is left open which sub-set. Moreover, since as I said it is somewhat arbitrary that we have chosen to focus on the thin ethical concept of *obligation*, we should suppose that the rest of the thin ethical concepts are subject to a similar universalizability condition. There is therefore a sense in which all the thin ethical concepts are world-guided.

The second conclusion is a corollary of this. When A says 'B ought to ϕ ' and D denies this, we must sharply distinguish between two very different possibilities. One is that A and D have a substantive ethical disagreement about what B ought to do. Their disagreement is substantive when they both meet conditions (i^{universal}) and (ii^{universal}), but they believe that a different sub-set of the universal features of B's circumstances provide reasons for B's desires about what to do. The mere fact that A and D share the thin ethical concept of *obligation*, and have ethical beliefs about what B's obligations are, thus doesn't guarantee that they will have the same conception of what B's obligations are. A's and D's conceptions of B's obligations turn instead on which sub-set of universal features they believe provide B with reasons to desire.

The other possibility, however, is that A and D are not having a substantive disagreement at all, but are rather having a merely verbal disagreement. Their disagreement is merely verbal when one of them is making an ethical judgement and the other isn't. This would be the case if (say) A meets conditions (i^{universal}) and (ii^{universal}), but B doesn't. Imagine, again, that when A says 'B ought to ϕ ', he is simply applying the universal principle that people ought to repay their debts, whereas when D denies this, he is a robber who thinks that B should give all his money to him. A's use of the word 'ought' is the application of an ethical concept, whereas D's isn't: D's claim simply reports what B needs to do in order to satisfy D's desire for money. What makes this the case is the fact that their judgements are grounded in worldly features that we can discriminate between in non-ethical terms. The features that ground A's judgement are a sub-set of the set of universal features, whereas the features that ground D's judgement concern a particular, namely, D himself. A's and D's judgments could therefore both be true.

So far we have asked what's required for a concept to be an *ethical* concept. Let's now turn to ask what might turn out to be a slightly different question. What makes a concept a *moral* concept? For example, what has to be true for someone's claims about another's obligations to be claims about their *moral obligations*? Hare's own view is that the three conditions already described suffice: judgements of moral obligation are universalizable, prescriptive, and overriding. He therefore sees no difference between ethics and morality. But whatever Philippa Foot thinks is required for ethical obligation, she plainly thinks that Hare's account of is inadequate as an account of moral obligation.

[S]ome philosophers [argue] that a man can choose for himself, so long as he meets formal requirements of generality and consistency, what his ultimate moral principles are to be; while others insist that certain criteria of good and evil belong to the concept of morality itself. The first, or formalist, position seems to me indefensible, implying as it does that we might recognize as a moral system some entirely pointless set of prohibitions or taboos, on activities such as clapping one's hands, not even thought of as harmful, aggressive, treacherous, cowardly by the community in which the prohibitions exist. A moral system seems necessarily to be one aimed at removing particular dangers and securing certain benefits... (Foot 2002, pp.6-7)

Foot's objection here is that the principle that people ought to clap their hands in circumstances where their doing so will do nothing to remove harms or secure benefits is perfectly universalizable, but that that plainly doesn't suffice for it to be moral principle. Nor is the problem that we have taken no account of prescriptivity or overridingness. The problem is rather that, for a principle to be a moral principle and hence to give rise to moral obligations, the principle must concern itself with acts that have something to do with removing harms or securing benefits.

Is Foot right? Remember that earlier it seemed incoherent to suppose that someone was making an *ethical* judgement if their judgement wasn't based on a universal principle. For Foot to be right, it would have to be similarly incoherent to suppose that someone has a belief about what people *morally* ought to do if he supposes that that has nothing whatsoever to do with removing harms and securing benefits. For reasons I will explain presently, I am not sure that this is incoherent. But whether or not Foot is right isn't really relevant to the present concern. What's relevant is rather what follows on the supposition that she is right. So, suppose now that Foot is right. What minimal revision to conditions (i^{universal}) and (ii^{universal}) would we be forced to make in order to turn our earlier Harean account of what makes a concept an ethical concept into an account of what, as she sees things, makes a concept a moral concept?

When A says to B, 'You ought to ϕ ', we would have to suppose that A expresses a belief about the instantiation of the concept of *moral obligation* only if:

- (i^{universal/harms/benefits}) A believes that there is some sub-set of the universal features of the circumstances in which B finds himself, features having something to do with removing certain sorts of harms and securing certain sorts of benefits: call these circumstances 'C^{universal/harms/benefits}'; and
- (ii^{universal/harms/benefits}) A believes that being in circumstances C^{universal/harms/benefits} is a reason for anyone who finds himself in those circumstances to desire that he ϕ s.

If these conditions are not met, we would have to suppose, then A isn't expressing a belief about moral obligation at all. Moreover, we would also be forced to suppose that similar conditions apply to the application of the other thin moral concepts, concepts like *moral rightness*, *moral wrongness*, *moral goodness*, and *moral badness*. The thin moral concepts as such would have to be such that believing that they are instantiated in certain circumstances is a matter of believing that some more restricted subset of the universal features of the circumstances, a sub-set having to do with the harms and benefits that arise, provides a reason for having some appropriate desire about what is to be the case in those circumstances.

If this is right, then note that two conclusions follow, conclusions that are very similar to those that followed from the constraint of universalizability in the ethical case. Focus again on the thin concept of moral obligation. The first is that there is also a sense in which the concept of moral obligation is world-guided, indeed *more* world-guided than the ethical concept of what ought to be done. Imagine, for example, that A says 'B morally ought to clap his hands' and that D denies this. If A believes that B's clapping his hands in his circumstances has nothing whatsoever to do with removing harms and securing benefits, then what he says is false for conceptual reasons. This is because the circumstances that A thinks provide B with a reason to desire to clap his hands have a feature characterizable in non-ethical and non-moral terms—the feature of being in circumstances in which B's clapping his hands would have nothing whatsoever to do with removing harms or securing benefits—that precludes them from being the sorts of circumstances which, even if they did provide B with reasons for desiring to clap his hands, are such that their providing B with reasons to desire to clap his hands makes it the case that B morally ought to clap his hands.

Of course, if these circumstances are still characterizable wholly in universal terms, then it may yet be true that B *ethically* ought to ϕ . Though it might initially sound fanciful to suppose that this is so, it isn't hard to imagine situations in which it at least some people might well think that it is so. Imagine, for example, that B's clapping his hands will in some miraculous way preserve a species from extinction, but neither the act of species preservation, nor the allowing of species extinction, has any effect on harms or benefits. Perhaps the species would have died out harmlessly, but, if it did, its ecological niche would have been occupied by a new species that would have brought about harms and benefits that are the same as those that would have obtained if the species had been preserved. The only difference made by B's clapping his hands, in other words, is the existence, rather than the non-existence, of the species itself. In these circumstances, it is easy to imagine that a radical environmentalist might think that B has a reason to desire to clap his hands, and hence believe that he ought to do so. But given that we have stipulated that this feature of B's circumstances has nothing whatsoever to do with removing harms or securing benefits, the claim that B *morally* ought to clap his hands in these circumstances is false on conceptual grounds. We may, however, still suppose that the radical environmentalist believes that B *ethically* ought to clap his hands.

(It should now be clear why I said that I am not sure that Foot is right. For the example of the radical environmentalist puts pressure on the idea that it is incoherent to suppose that someone has a belief about what people *morally* ought to do if he supposes that what they morally ought to do has nothing whatsoever to do with removing harms and securing benefits. The idea that radical environmentalism of the kind described is an ethical view, but not a moral view, seems more like a stipulation about how to use the words 'ethical' and 'moral', not an *a priori* truth about either ethics or morality. At this stage I am therefore inclined to side with Hare rather than with Foot. The ethical concept of obligation *is* the moral concept of obligation, and the only conceptual constraint on the concept of obligation is the constraint of universalizability.)

The second conclusion is a corollary of this. When A says 'B ought to clap his hands' and D denies this, we must sharply distinguish between two very different possibilities. One is that A and D have a substantive moral disagreement about what B ought to do. Their disagreement is a substantive moral disagreement when they both meet conditions (i^{universal/harms/benefits}) and (ii^{universal/harms/benefits}), but they believe that a different sub-set of the universal features of B's circumstances having something to do with removing harms and securing benefits provide B with reasons to desire to clap his hands. Perhaps A is a utilitarian and D is a deontologist. The mere fact that A and D share the thin concept of *moral obligation*, that they use the word 'ought' to express this concept, and that they both have beliefs about what B's moral obligations are, thus doesn't guarantee that they will have the same conception of what B's moral obligations are. A's and D's conceptions of B's moral obligations turn instead on which sub-set of universal features having to do with removing harms and securing benefits they believe provide B with reasons to desire.

The other possibility, however, is that A and D are not having a substantive disagreement at all, but are rather having a merely verbal disagreement. Their disagreement is merely verbal when one of them is making a moral judgement and the other isn't. This would be the case if (say) A meets conditions (i^{universal/harms/benefits}) and (ii^{universal/harms/benefits}), but B's merely meets (i^{universal}) and (ii^{universal}). Think again about the possibility that A is a radical environmentalist in the circumstances we imagined above, and suppose that D is a utilitarian in those same circumstances. A's use of the word 'ought' is the application of the concept of ethical obligation, not the concept of moral obligation, whereas D's use of 'ought' is the application of the concept of moral obligation. What makes their disagreement merely verbal is thus the fact that their judgements are grounded in differently worldly features, features that we can distinguish between in non-ethical terms. The features that ground A's judgement are a sub-set of the universal features, but features that have nothing to do with removing harms or securing benefits, whereas the universal features that ground D's judgement do have something to do with removing harms and securing benefits. This explains why their judgements could both be true together.

Let me sum up the discussion thus far. We've spent some time discussing Hare's suggestion that thin ethical concepts are universalizable, and Foot's suggestion that thin moral concepts aren't just universalizable, but must have something to do with removing harms and securing benefits, and we have seen how we can build these insights into our account of what makes a concept an ethical concept. In doing so we have learned that, if Foot is right, then we may have to distinguish sharply between the *thin ethical* concepts, on the one hand, and the *thin moral* concepts on the other. We could do this by acknowledging that there is a sense in which, though both the thin ethical concepts and the thin moral concepts are world-guided, the thin moral concepts are more world-guided than the thin ethical concepts. In other words, a more demanding worldly condition must be satisfied for a thin moral concept to be applicable. Finally, we have seen that satisfying these worldly conditions, worldly

conditions that are themselves characterizable in non-ethical and non-moral terms, doesn't wholly determine which features someone who applies these concepts believes to provide reasons for desiring. The worldly condition narrows the set of features that can be believed to provide reasons, but doesn't wholly determine it.

With this discussion of thin ethical and thin moral concepts under our belt, let's now turn to consider the thick ethical concepts. What is it about the thick ethical concepts that makes them world-guided? The best place to start in providing an answer to this question, I think, is with the exemplary work done by Daniel Elstein and Thomas Hurka (Elstein and Hurka 2009). Elstein and Hurka's aim is to show that thick ethical concepts can all be given a reductive two-part analysis in the spirit of Hare's two-part analysis. Of course they don't analyze every thick ethical concept. Instead they illustrate their general strategy of analysis by working through an example, namely, how to analyze the concept of distributive justice. Their claim, which they leave as a homework task for those who doubt them, is that similar analyses can be given of all the other thick ethical concepts.

Here is the passage in which Hurka and Elstein explain the difficulty to be faced in providing such a reductive analysis.

Consider the term 'just' in the specific sense used in discussions of distributive justice. And imagine two people who disagree about the ground of distributive justice, one saying that just distributions of, say, happiness are equal distributions, while another says they are distributions proportioned to people's merits, which may be unequal. If the reductive view had to use a fully determinate two-part analysis of 'distributively just,' it would have to say these two people use different concepts of justice. For the egalitarian, 'x is just' means 'x is an equal distribution and is good for being so,' while for the desert-theorist it means 'x is a distribution proportioned to merit and is good for being so.' It would not follow that the two cannot disagree about distributive justice, since the claims their concepts imply about the ground of goodness in distribution contradict each other. But in debating the justice of a particular distribution they could not be disagreeing about whether a single concept applies to it, since they use different concepts. But does it not look as if they are disagreement about the application of a single concept? And is an analysis not preferable that allows this? (Elstein and Hurka 2009, pp.521-522)

The problem they describe was the main problem faced by the two-part analyses Hare himself offered. Hare thought that all thick ethical concepts could be analysed into a non-ethical component and a thin ethical component. Moreover, he thought that the non-ethical component had to be a specification of the non-ethical features in virtue of which the thin ethical concept that figures in the ethical component of the analysis is instantiated. Thus, for example, if Hare were to provide an analysis of the concept of distributive justice, then, given that he is a utilitarian, he would have said that 'x is just' means 'x is a distribution in which individuals are satisfied according to the principle of utility and is good for being so.' What's crucial here is that the first component of the analysis tells us everything about the distribution in virtue of which it is good: it is, in Elstein and Hurka's terms, "fully determinate". And the trouble they see with this is that it makes it impossible for an egalitarian or a desert-theorist's use of the expression 'x is just' to have the same meaning. For they don't believe that a distribution is good for being one in which individuals are satisfied according to the principle of utility. An egalitarian and a desert-theorist cannot express their disagreement with each other by using a single concept of justice.

Elstein and Hurka go on to explain how we can provide a reductive two-part analysis of thick ethical concepts that doesn't face this problem.

We can construct such an analysis if we...give 'distributively just' an only partly determinate descriptive component. Then 'x is distributively just' will mean something like 'There are some properties X, Y, and Z (not specified) that distributions have as distributions, or in virtue of its distributive shape, such that x has X, Y, and Z, and X, Y, and Z make any distribution that has them good.' This analysis certainly places significant restrictions on the extension of 'distributively just.' If someone says an act of generously helping a stranger is distributively just, she is misusing the concept, because only distributions can be just in this sense. And if she calls a distribution just because it was brought about by generous actions, she is likewise misusing the concept, because only a distribution's properties as a distribution can bear on its justice. But these restrictions do not completely determine the concept's extension, and in particular do not determine whether it contains equal distributions or ones based on desert. That depends on which properties of distributions are in fact good-making... To know the intended extension of 'distributively just' as used by some person or community, it is not enough to know the descriptive part of that term's meaning; we must also know what evaluations they make, that is, which properties they take to make extensions good. And to know the term's actual extension, we must know which properties in fact make distributions good. (Elstein and Hurka 2009, p522)

The crucial feature of Elstein and Hurka's analysis of justice, unlike the earlier two-part analysis, is that the non-ethical component is not fully determinate. It constrains, but doesn't wholly determine, the features in virtue of which the thin ethical concept that figures in the ethical component is instantiated. The attraction of this less determinate analysis is that it is therefore an analysis of a concept that could be shared and used by an egalitarian, a desert-theorist, and a utilitarian alike, a concept that they could use to express their disagreements with each other about what makes a distribution just. For though they disagree about the grounds of goodness in a distribution, they can also agree that there is no conceptual barrier to supposing that something is just on different grounds, so long as those grounds have something to do with the distributive shape of the thing that is being judged to be just.

Elstein and Hurka's analysis of justice can easily be reformulated in our terms. So reformulated, and taking into account what we learned from our discussion of Hare, the idea would be that when A says to B, 'x is just', A expresses a belief about the instantiation of the thick ethical concept of distributive justice only if:

(i^{universal/distributive shape}) A believes that x has, among its universal features, a distributive shape, and he further believes that that distributive shape has a further feature F; and

(ii^{universal/distributive shape}) A believes that the fact that x's distributive shape has feature F is a reason to desire that x has that distributive shape.

It thus isn't a requirement on anyone's having a belief about the justice of some state of affairs that the feature of the distributive shape that they believe provides reasons to desire that that state of affairs obtains is the same as the one everyone else believes provides such reasons. It is sufficient that they believe that the reason-providing feature is some feature or other of the distributive shape of the state of affairs. A and B can therefore disagree with each other about whether x is just, precisely because they disagree with each other about which specific feature of the distributive shape provides reasons for desiring that that state of affairs obtains.

On the assumption that the concept of justice is a *moral* concept, then, if what Foot says is right about the analysis of moral concepts, our analysis of justice requires further

amendment. When A says to B, 'x is just', A expresses a belief about the instantiation of the thick *moral* concept of justice only if:

(i^{universal/harms/benefits/distributive shape}) A believes that x has, among its universal features that have something to do with harms and benefits, a distributive shape, and he further believes that that distributive shape has feature F; and

(ii^{universal/harms/benefits/distributive shape}) A believes that the fact that x's distributive shape has feature F is a reason to desire that x has that distributive shape.

In other words, if Foot is right, then one has a fully-fledged moral belief about the justice of some state of affairs only if one believes that the feature of its distributive shape that provides one with a reason has something to do with harms and benefits. But as is I hope clear, it is far from obvious that this is correct. Those who reject all forms of welfarism—for example, those like Rawls who think that justice concerns the distribution of primary goods (primary goods, remember, are all purpose means like opportunities, the social bases of self-respect, and the like), rather than the distribution of welfare—deny that just distributions need have anything to do with harms and benefits (Rawls 1971). Since it doesn't seem incoherent to suppose that these theorists are right, it seems to me that we should therefore stick with our original reformulation of Hurka and Elstein's analysis, the one that insists merely that the feature of distributions has to be a universal feature. In other words, we should suppose that the ethical concept of justice *is* the moral concept of justice.

Note that our original reformulation of Hurka and Elstein's analysis explains why the thick ethical concept of justice is world-guided. The concept is world-guided because a worldly condition must be satisfied by the considerations that someone who employs the concept thinks provide reasons for desire, a worldly condition that we can describe in non-ethical terms. These considerations must be *universal*, and they must have something to do with a *distributive shape*. The upshot is that if someone applies the concept of justice to something in virtue of something other than its distributive shape, then that person is making a straightforward conceptual error. Their use of the concept doesn't satisfy the worldly condition. For example, imagine a disagreement between an environmentalist and an egalitarian about whether a certain distribution is just or unjust, but suppose further that the radical environmentalist's grounds for desiring that that distribution obtains is the fact that it is a distribution in which a certain species is preserved. This disagreement is purely verbal, as something can only be just if it the feature that provides a reason for desiring is a universal feature having something to do with a distributive shape. But if an egalitarian says 'x is just' and a desert-theorist says 'It is not the case that x is just', then this is a substantive disagreement. It is a substantive disagreement because they are both believe that the considerations that provide reasons to have desires concerning x are universal features having something to do with x's distributive shape, it is just that they disagree about which such features provide such reasons.

What's much more striking about this analysis of the thick ethical concept of justice, however, is the fact that it has exactly the same form as the earlier analysis the thin ethical concepts inspired by Hare. Thick and thin ethical concepts both turn out to be world-guided, and the explanation of their being world-guided is exactly the same in each case. The world-guidedness of an ethical concept, whether thick or thin, is a matter of there being a constraint on the considerations that those who employ the concept believe to provide reasons for desiring, a constraint that, so far at any rate, we have been able to describe in non-ethical terms. In the case of the thin ethical concepts, the constraint is that those features must be universal, and in the case of the thick ethical concept of justice, the constraint is that those features must be universal and have something to do with distributive shape. In order to

complete our homework task, all we need to do is to work through the remaining thick ethical concepts and identify the constraints on the reasons for desires that they are believed to provide.

Moreover, just to keep the record straight, note that if we were to do this then we would discover that such constraints *are not* always describable in non-ethical terms. For example, Elstein and Hurka ask us to

...consider the virtue of integrity. It is a morally good trait that involves, roughly, sticking to one's ideals and projects despite temptations or distractions. But not any fidelity to a project counts as integrity. Someone who persists in building his beer-mat collection despite the rise of Nazism around him and the temptation to fight against it would hardly be described as acting with integrity. The reason is that only fidelity to good or important goals counts as integrity. An initial analysis of 'x is an act of integrity' therefore runs something like 'x is good, and x involves an agent's sticking to a significantly good goal despite distractions and temptations, where this property makes any act that has it good,' and where the second 'good' indicates an embedded evaluation. Given this analysis, we can only know what counts as integrity if we know which goals are independently good, and there can be disputes about this. Consider a pope who remains faithful to the views on homosexuality and the ordination of women that he and his church have long held, despite calls from modernizers within and outside the church for change. Supporters of traditional Catholic teaching may say his resistance to reform shows integrity; those who reject that teaching will not. (Elstein and Hurka 2009, p.526)

Certain thick ethical concepts are thus such that some of the believed constraints on the features that provide reasons for desire are themselves only describable ethically. In our terms, Elstein and Hurka's suggestion can be recast as follows. When A says 'B has integrity', A has a belief about B's *integrity* if and only if:

- (i^{universal/temptation-imperviousness/ideals}) A believes that B has, among his universal features, some universal features having something to do with sticking to his ideals and imperviousness to temptations: call these features 'G'; and
- (ii^{universal/temptation-imperviousness/ideals}) A believes that B's ideals are themselves aimed at things that have some universal feature that provides reasons to desire that those things obtain; and
- (iii^{universal/temptation-imperviousness/ideals}) A believes that G plus the fact that B's ideals are themselves aimed at things that have some universal feature that provides reasons to desire those things, provides reasons to desire that B sticks to those ideals and remains impervious to temptations.

In order to fully understand the world-guidedness of the thick ethical concept of integrity, it is therefore crucial that we appreciate that ethical considerations can themselves serve as reasons for desires.

We began this paper with two questions. The first was whether thick ethical concepts are world-guided in some way in which the thin are not. We have now answered this question and we have done so by doing two distinct things. We first gave an analysis of the thin ethical concepts and saw that there is a sense in which they are world-guided. We then gave an analysis of the thick ethical concepts and saw that they too are world-guided, indeed, world-guided in the very same sense in which the thin ethical concepts are world-guided. The world-guidedness of an ethical concept, whether thick or thin, has turned out to be a matter of there being a constraint on the sorts of considerations that those who employ the concept believe to provide reasons for desiring. The upshot is thus the one anticipated at the

very beginning. There is no distinction between thick and thin ethical concepts in terms of world-guidedness. All ethical concepts, even the maximally thin ones, turn out to be a little bit thick.

3. *Thick ethical concepts, thin ethical concepts, and the issue of definitional priority*

Let's finally turn to a brief discussion of the second question with which we began. Should we analyze the thick ethical concepts in terms of the thin, or the thin in terms of the thick, or should we simultaneously define thick and thin ethical concepts in terms of each other, perhaps along with their relations to other concepts?

It might be thought that the answer to this question has already been given. For haven't we seen that the thin ethical concepts—*obligatoriness*, *rightness*, *wrongness*, *goodness*, and *badness*—are all defined in terms of the existence of some universal features that provide reasons for desiring, and haven't we seen that the thick ethical concepts are all simply definable in terms of some further restriction on these reasons for desiring? So haven't we already seen that the thick ethical concepts are indeed to be defined in terms of the thin ethical concepts? But tempting though this line of thought may be, we shouldn't go along with it straight away. The line of thought starts from the observation that the features that provide reasons for desires must themselves be universal. But since instantiating a thick ethical concept is itself a universal feature, it follows that the observation doesn't rule out the possibility that the universal features that provide reasons for desires are always further restricted in the way in which some thick ethical concept is restricted. The question remains whether this is so and, if it is, whether it follows that we can define the thin ethical concepts in terms of the thick (or, perhaps, simultaneously define them in terms of each other).

What exactly such a definition would look like isn't something that anyone has fully spelled out, at least to my knowledge. Susan Hurley has, however, done us the great service of describing the idea in broad outline.

A feature common to many philosophical accounts of ethical concepts is that the general concepts, *right* and *ought*, are taken to be logically prior to and independent of the specific concepts, such as *just* and *unkind*. According to such accounts, the general concepts carry a core meaning...that also provides the specific concepts with reason-giving status... I shall refer to accounts that take the general concepts in some category to be logically prior to and independent of the specific as *centralist*. (Hurley 1985, p.56)

The tempting line of thought just described is an instance of centralist thinking. As Hurley goes on to point out, however, this isn't the only possible view.

Non-centralism about reasons for action rejects the view that the general concepts *right* and *ought* are logically prior to and independent of specific reason-giving concepts such as *just* and *unkind*. Instead it may take the identification of discrete particular values as a starting point, subject to revision, and give an account of the interdependence between the general concepts and the specific concepts. Coherentist views provide an example of non-centralism. According to such a view, to say that a certain act ought to be done is to say that it is favoured by the best account of the relationships among the specific values that apply to the alternatives in question... (Hurley 1985, p.56)

A specification of how the reason-providing universal features associated with thin ethical concepts are further restricted may therefore itself derive from an account of the way in which different thick ethical features present in a situation relate to each other: that is, it may

be an account of the normative significance of the various thick ethical features vis a vis each other, something that is in turn fixed (say) by the best theory of their relations.

The concept of an act's being obligatory in certain circumstances, according to non-centralism, is thus the concept of an act's having the thick ethical feature, or combination of features, that has the greatest normative significance in those circumstances, and this in turn is a concept which someone believes to be instantiated when he believes that that act has whichever thick ethical feature, or combination of thick ethical features, has the greatest normative significance in those circumstances. Someone who believes that an act is obligatory will therefore believe that the universal feature that provides a reason for desiring in those circumstances is further restricted in whatever way the concepts of those thick ethical features are restricted. But note that non-centralism itself takes no stand on whether we can define the thick ethical wholly independently of the thin ethical concepts because, for all that it tells us, defining the normative significance of the thick ethical concepts in terms of each other may require further appeal to the thin ethical concepts. In this way non-centralism allows that the thin ethical concepts may be simultaneously defined in terms of each other (compare Hurley 1985 p.58).

Note that there are two readings of non-centralism. According to the *de re* reading, there is a distinctive set of thick ethical features that have a certain normative significance vis a vis each other in various circumstances—for short, let's just suppose that these are exhausted by justice and kindness—and we define the thin ethical concepts in terms of these features (and perhaps vice-versa). To be obligatory is by definition to be either just or kind. This in turn has implications for what it is to believe that an act is obligatory. According to the *de re* reading, there is a distinctive set of thick ethical features that have a certain normative significance vis a vis each other in various circumstances—justice and kindness—such that, when someone believes that an act performed in those circumstances is obligatory, he believes that that act has those thick ethical features.

According to the *de dicto* reading, the thin ethical concepts are to be defined in terms of some set of thick ethical features that have a certain normative significance vis a vis each other in various circumstances, but that set of thick ethical features is left unspecified. To be obligatory is by definition to have some thick ethical feature or other. When someone believes that an act performed in certain circumstances is obligatory, he therefore believes that there is some set of thick ethical features that have a certain normative significance in those circumstances, and that the act has those thick ethical features. The *de dicto* reading thus allows people to have beliefs about obligatoriness based on their *false* beliefs about the thick ethical features that there are. It allows, for example, that the homophobe can have beliefs about what's obligatory based on his beliefs about 'poofs', as he would put it. It simply requires that, to have a belief about what's obligatory, he has to believe that the act has *some* thick ethical features. The *de re* reading doesn't allow this.

When Hurley suggests that we can define the thin ethical concepts in terms of the thick (and perhaps vice-versa), she certainly seems to have the *de re* reading in mind. But given what this tells us about what it is to believe an act to be obligatory, we must surely conclude that non-centralism is false on the *de re* reading. The homophobe who says, 'Homosexuals ought to be shunned!' does not believe, of the thick ethical features of justice and kindness, that shunning homosexuals has *those* features. He believes that not to shun homosexuals would be 'poofy', as he would put it. This simplifies our discussion, for it leaves us with non-centralism on the *de dicto* reading. The question we must ask about the doctrine on this reading isn't just whether it is plausible, but whether it is an alternative definitional doctrine to centralism at all. Let's start with whether the doctrine is plausible.

On the *de dicto* reading, there is no constraint at all on the thick ethical features that non-centralism insists are definitive of obligation. What's true by definition is simply that an act's being obligatory is a matter of its having some thick ethical features or other. When someone believes an act to be obligatory, he therefore believes that this is so because of the presence of some thick ethical feature or other. Accordingly, it might be thought that non-centralism on this reading is manifestly implausible. For think again about the radical environmentalist who believes that someone ought to clap his hands because he will thereby preserve a species, and suppose further that the radical environmentalist believes this because of the loss of value that the extinction of a species would represent. He has a belief about the obligatoriness of acting to preserve the species, it might be thought, but he doesn't believe that any thick ethical feature is present. This seems to be a counterexample to non-centralism on the *de dicto* reading, or so it might be thought.

But I am not so sure about this, as on plausible assumptions, the radical environmentalist clearly does believe that a not-maximally-thin thick ethical feature is present. For we have seen that a not-maximally-thin thick ethical concept is just a concept whose application requires a further constraint on the sorts of considerations that provide reasons for desiring, over and above being universalizable, and the radical environmentalist does believe that such a concept is instantiated, as he believes that the universal feature of causing species-preservation provides a reason for desiring to clap hands. Of course, the radical environmentalist's language may not include a word whose dedicated role is to express this concept, and we may have no such word either. Perhaps the best either of us can do is to talk about how the following universal feature of state of affairs—the feature of that-state-of-affairs-having-the-universal-feature-of-being-one-in-which-a-species-is-preserved—provides a reason to desire that that state of affairs obtains. But this is neither here nor there. The important point is simply that this is what it is to believe that a thick ethical feature is present.

The real question, I think, is therefore not whether non-centralism so understood is plausible, but whether it is an alternative definitional doctrine to centralism at all. Given that the truth of non-centralism on the *de dicto* reading is entailed by the truth of centralism, myself I think it follows that it is not. Rather, as we have seen, what it means for acts to have some thin ethical feature is for them to have *some sub-set* of universal features that provide reasons for desiring, and on the *de dicto* reading of non-centralism, this is exactly what it means for some thick ethical feature to be present. Non-centralism on the *de dicto* reading thus isn't an alternative to centralism. The upshot, I think, is that we should go along with the tempting thought outlined at the very beginning of this section. The thick ethical concepts can indeed all be defined in terms of the thin because the thin ethical concepts are all defined in terms of the existence of some universal features that provide reasons for desiring, and the thick ethical concepts are all defined in terms of some further restriction on these reasons.

Conclusion

I said at the outset that the significance of the distinction between thick and thin ethical concepts is not easy to assess. I hope it is now clear why this is so. Since all ethical concepts are a little bit thick, it follows that there isn't really any distinction to be drawn between them. There is only a distinction between the maximally-thin thick ethical concepts and the not-maximally-thin thick ethical concepts. However since we have also seen that we can define the rest of the thick ethical concepts in terms of the maximally-thin ones, it follows that the maximally-thin thick ethical concepts have a significance that is over and above that of the not-maximally-thin thick ethical concepts. Their significance lies in their definitional priority.¹

REFERENCES

- Elstein, Daniel Y. and Thomas Hurka 2009: 'From Thick to Thin: Two Moral Reduction Plans' in *Canadian Journal of Philosophy* (39) pp. 515-536.
- Foot, Philippa 2002: 'Morality and Art' in her *Moral Dilemmas—and Other Topics in Moral Philosophy* (Oxford: Oxford University Press).
- Hare, R. M. 1952: *The Language of Morals* (Oxford: Oxford University Press).
- Hare, R. M. 1963: *Freedom and Reason* (Oxford: Oxford University Press).
- Hare, R.M. 1981: *Moral Thinking* (Oxford: Oxford University Press).
- Hurley, S.L. 1985: 'Objectivity and Disagreement' in Ted Honderich, ed, *Morality and Objectivity* (London: Routledge and Kegan Paul).
- Parfit, Derek 2011: *On What Matters: Volume One* (Oxford: Oxford University Press)
- Rawls, John A *Theory of Justice* (Cambridge MA: The Belknap Press of Harvard University Press).
- Scanlon, Thomas M. 1998: *What We Owe to Each Other* (Cambridge: Harvard University Press).
- Thomson, Judith Jarvis 2008: *Normativity* (Chicago: Open Court Publishing Company).
- Williams, Bernard 1981: 'Internal and External Reasons' in his *Moral Luck* (Cambridge: Cambridge University Press).
- Williams, Bernard 1995: 'Internal Reasons and the Obscurity of Blame' in his *Making Sense of Humanity* (Cambridge: Cambridge University Press).

¹ Earlier versions of this paper were presented at *Thick Concepts*, University of Kent, Canterbury, England, June 2009; the Australasian Association of Philosophy Annual Conference, Melbourne University, June 2009; and Rice University, January 2010. I am grateful for all of the helpful comments I received on these occasions, but especially for those received from Simon Kirchin and the skeptical audience at Rice.