


This article was downloaded by: 10.3.97.143

On: 28 Mar 2023

Access details: *subscription number*

Publisher: *Routledge*

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: 5 Howick Place, London SW1P 1WG, UK



## The Routledge Handbook of Metaethics

Tristram McPherson, David Plunkett

### Thick Concepts

Publication details

<https://www.routledgehandbooks.com/doi/10.4324/9781315213217.ch13>

Debbie Roberts

**Published online on: 28 Aug 2017**

**How to cite :-** Debbie Roberts. 28 Aug 2017, *Thick Concepts from: The Routledge Handbook of Metaethics* Routledge

Accessed on: 28 Mar 2023

<https://www.routledgehandbooks.com/doi/10.4324/9781315213217.ch13>

**PLEASE SCROLL DOWN FOR DOCUMENT**

Full terms and conditions of use: <https://www.routledgehandbooks.com/legal-notices/terms>

This Document PDF may be used for research, teaching and private study purposes. Any substantial or systematic reproductions, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The publisher shall not be liable for an loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

## 13

## Thick Concepts

Debbie Roberts

There has been a renewed flurry of interest in thick concepts in the last ten years or so. Despite this, they remain a fairly niche topic in metaethics, and in metanormative theorising more generally. In this chapter, I aim to show that thick concepts deserve more attention in general than they currently enjoy.

Metaethics, and metanormative theorising more generally, concerns giving an account of all normative thought and talk. *Prima facie*, thick terms and concepts are part of the relevant range of terms and concepts, so an account of them is part of the task of metanormative theorising. In addition, as I will argue below, thick terms and concepts are intrinsically puzzling because they seem to combine evaluation and non-evaluative description. While it would, of course, be intrinsically interesting to figure out how they do this, solving this puzzle turns out to have interesting and important implications for metanormative theory more generally.

‘What Are Thick Concepts?’ is a brief characterisation of what thick concepts are. I then outline the main views in the debate in the section ‘The Main Views’. The section three that follows this is the meat of this chapter. There I lay out the current state of play, pointing out the main points of contention between the different views. I end that section by examining which concepts end up counting as thick according to the different positions. The final section, ‘Why Are Thick Concepts Important?’, considers the significance of thick concepts.

## WHAT ARE THICK CONCEPTS?

What is a theory of thick concepts a theory of? There are at least two routes to answering that question. The first provides some paradigm cases. The second provides an answer in theoretical terms. Here, I take the first route.

A standard way of introducing thick concepts is to point to an intuitive contrast between concepts like GENEROUS, BRUTAL, CRUDE and TACTLESS, and concepts like GOOD,

BAD, RIGHT and WRONG (though the difference may be one of degree rather than one of kind [Scheffler 1987]). Bernard Williams (1985: 129) introduces to ethics the terminology of ‘thick’ and ‘thin’ to mark this contrast and gives the following examples of thick concepts: PROMISE, TREACHERY, BRUTALITY, COURAGE, COWARD and LIE.

As well as examples drawn from the *ethical*, and the *practical* realm more generally, like those above, paradigmatic cases in the literature are also drawn from other normative domains. In *aesthetics*, DAINTY, DUMPY, DELICATE and DASHING are among these. And specifically *political* concepts seem to belong on the list too, for example, DEMOCRACY, LEGITIMACY and DISTRIBUTIVE JUSTICE. More controversially, some have claimed that there are *legal* thick concepts, for example, NEGLIGENCE and MURDER, as well as LAW and LEGAL themselves. And others have claimed that there are *epistemic* thick concepts, for example, GULLIBLE, CONSCIENTIOUS and OPEN-MINDED (compare *Philosophical Papers* 2008 [37]).

There are other examples of terms and concepts that, like the legal and epistemic examples, don’t typically feature as paradigmatic cases but seem intuitively like they might either turn out to be thick or be interestingly similar-but-different to the thick. These include *slurs and pejoratives*, like ‘slut’, ‘kraut’, ‘girly’ and ‘poof’; *affective notions* like ‘admirable’, ‘desirable’ and ‘contemptible’; certain *complex predicates* that explicitly concern ways of being good or bad, like ‘good torturer’ and ‘bad dancer’; and, lastly, *prima facie uncontroversially nonevaluative terms* that are often used to convey evaluations pragmatically, e.g., ‘painful’, ‘pleasurable’, ‘fat’ and ‘athletic’. Whether any of these kinds of terms end up counting as thick will depend on the details of the correct account. On the approach to answering the question that I started this section with, working out which account of thick concepts one ought to adopt will in part involve working out which generates the most plausible-looking list.

I’ve moved from talking about thick concepts to talking about thick terms and concepts. Roughly put, the idea that most in this literature are working with is the simple one that terms express concepts; the semantic content of a term is, or overlaps exactly with, the concept.

So how are we to characterise the intuitive contrast between thick and thin? BRUTAL and SPITEFUL are thick and BAD thin, seemingly in virtue of BRUTAL and SPITEFUL, each containing or conveying more detailed information than BAD. However, more precise accounts that nonetheless remain neutral between particular views on the nature of thick concepts are non-existent. (Compare Väyrynen 2013: 4–7; Eklund 2011.) This is because part of what is at issue in this debate is just what makes thick concepts thick.

For the moment, since we must start somewhere, let’s take as a working attempt to characterise the distinction the claim that thick concepts and the terms that express them somehow combine evaluation and non-evaluative description whereas thin concepts are (or are more purely) evaluative or normative. Even this is not a neutral distinction, as should become clear below. But this seems at least an initially plausible explanation of how it is that thick terms and concepts seemingly convey or contain more detailed information than that conveyed or contained by the thin.

## THE MAIN VIEWS

Broadly speaking, there are three different types of view one might have on the nature of thick terms and concepts (Roberts 2013b). The first divergence arises depending on

how the evaluative aspect of thick terms and concepts is taken to be communicated. Some hold that thick terms and concepts are themselves evaluative as a matter of content (“inherently evaluative,” to use Väyrynen’s terminology) so the evaluation is, or is part of, the *semantic content* of the relevant terms. Others hold that the evaluations most closely associated with thick terms and concepts are communicated by some kind of *pragmatic* mechanism, and thus that thick terms and concepts are not inherently evaluative, that is, not evaluative as a matter of content.<sup>1</sup>

Content views can be further divided into those who think that the content of thick terms and concepts can in principle be separated into non-evaluative descriptive and thin evaluative elements of content and those who think this cannot be done. I call these *reductive* and *non-reductive* content views, respectively.

The *non-reductive content* view holds that thick terms and concepts are inherently evaluative and irreducibly thick. It is on this view that thick terms and concepts are held to have distinctive significance (e.g., as an argument against non-cognitivist views and ‘the fact-value distinction’). The other kinds of views represent two different ways one might deflate the significance of the thick and avoid the problems for various theories that thick concepts are supposed to cause.

One view of the nature of thick terms and concepts, the *descriptive equivalence* view, has been widely discredited in this literature, though it is still worth mentioning here. The descriptive equivalence view claims that for every thick term (concept) there is, or could be, a purely non-evaluative descriptive term (concept) with the same extension (compare Mackie 1977: 16–17; Hare 1981: 72). Against this it is widely held that for thick terms and concepts ‘evaluation drives extension’, that is, in order to be able to apply the term to new cases, one has to at least imaginatively adopt the evaluative perspective within which the relevant concept has its point. If a non-evaluative descriptive term with the same extension were possible, an outsider would be able to see how to go on applying the term to new cases without ever even just imaginatively adopting the relevant evaluative perspective (see McDowell 1981: 201–202).

It seems very plausible that the descriptive equivalence view is false for at least *paradigmatic* examples of thick terms and concepts. (It is a widely held view in the slurs literature, but as I said at the start, it is controversial whether or not slurs are thick concepts.) When one starts to try to come up with extensional equivalents for ‘cruel’ or ‘elegant’, it’s very difficult not to reach for other seemingly *evaluative* terms (‘suffering’, ‘graceful’) and very easy to end up using ‘cruel’ and ‘elegant’ themselves. This looks to undermine not only the claim that we have in our current vocabulary the non-evaluative term that is extensionally equivalent to cruel, but also the claim that we could in principle invent one.

One might think that the failure of the descriptive equivalence view shows that the evaluative and non-evaluative content of thick concepts cannot be separated out or ‘disentangled.’<sup>2</sup> And, indeed, some do appear to have concluded this (see, e.g., Williams 1985: 141; Putnam 2002: 37–38). That the evaluative and non-evaluative content of thick concepts cannot be separated out, even in principle, is a crucial element of the non-reductive content view. So if we could conclude ‘no separation’ from the failure of descriptive equivalence, this would be good news for the non-reductivists.

However, unhappily for the non-reductive view, reductive content and pragmatic views can happily accommodate that the descriptive equivalence view is false, though they do so in very different ways. Most importantly, the reductive content view accepts

the claim that evaluation drives extension for thick terms and concepts, but the pragmatic view does not.

Let's take *reductive content* views first. Prominent defenders include Gibbard (1992) and Elstein and Hurka (2009). Elstein and Hurka, for example, put forward two reduction plans for thick terms, neither of which generate descriptive equivalents, so they are both compatible with the claim that evaluation drives extension for the thick. The first is a two-component analysis where the non-evaluative descriptive component does not fully specify what something has to be like to merit the application of the thick term. 'Cruel', for example, might be handled in this way:

'x is cruel' means 'x is bad and there are properties P, Q & R (unspecified) of non-evaluative sort c (specified) such that x has properties P, Q & R and P, Q & R make anything that has them bad'.

(2009: 522)

Elstein and Hurka's second reductive plan introduces an *embedded* evaluation, distinct from the *global* evaluation governing the whole use of the term. When it comes to thick terms and concepts, an evaluation is *global* if that evaluation applies to all the features that distinguish the things falling under that term or concept. We can contrast this with *embedded* evaluations, which are evaluations required to specify the very thing over which the global evaluation will take scope. Take 'distributively just' as an example, and assume that 'x is distributively just' means something like 'x has features X, Y and Z as a distribution and is good for having those features'. The 'good' that occurs in the analysis is a global evaluation. An embedded evaluation would be present if specifying the type of thing to which the global evaluation applies required evaluative information. In this case, if one or more of X, Y or Z were an evaluative feature (e.g., perhaps X is 'is the result of a fair procedure') then 'distributively just' would contain an embedded evaluation.

And now for the second reductive plan:

'x is courageous' means 'x is good [global evaluation], and x involves an agent accepting harm or the risk of harm for himself for the sake of *goods greater than the evil* [embedded evaluation] of that harm, where this property makes any act that has it good'.

(2009: 526).

Both reductive plans are compatible with the claim that evaluation drives extension, and the falsity of the descriptive equivalence view, for in both cases, in order to determine the extension of the term, one has to make some evaluations. In the first case, evaluation is required to settle precisely which non-evaluative features of type c are the bad-making, and thus cruel, ones. In the second, evaluation is required to settle when accepting harm or the risk of harm is worthwhile and thus courageous (it's worth noting that Elstein and Hurka are assuming the 'harm' is non-evaluative).

On the *pragmatic view*, thick terms and concepts are not inherently evaluative. Prominent defenders include Hare (1952; 1963; 1981), Blackburn (1992) and Väyrynen (2013). On Väyrynen's view, the fact that the descriptive equivalence view is false, and that one can't go on to apply 'cruel' correctly in new cases without drawing on one's

understanding of ‘cruel’, doesn’t show anything more than would be shown by observing that one can’t go on to apply ‘red’ or ‘chair’ correctly in new cases without drawing on one’s understanding of ‘red’ or ‘chair’. If this is correct, the ‘evaluation drives extension’ point that non-reductivists appear to have made so much of may really be no more than a misidentification of the unexciting point that thick terms and concepts lack a classical structure. Moreover, Väyrynen argues, thick terms should be understood as context-sensitive gradable terms (like ‘tall’). And for such terms, there is really no such thing as ‘the’ extension (Väyrynen 2013: chapters 7 and 8).

To sum up, there are three options currently on the table: The *non-reductive* content view, the *reductive* content view and the *pragmatic* view. The main point historically taken to favour the non-reductive view, that the descriptive equivalence view is false, can be accommodated by all three. How, then, are we to decide between them? In the next section, I set out the main points of contention between these views.

### THE CURRENT STATE OF PLAY

The main issues in the current debate concern objectionable thick terms and concepts, what makes a concept evaluative, the question of whether evaluation drives extension for thick terms and concepts, the distinction between global and embedded evaluation, the non-evaluative shapelessness of thick concepts and, lastly, the question of which terms and concepts get to count as thick.

#### Objectionable thick terms and concepts

In classifying the available positions on the nature of thick concepts, I said above that the first divergence occurs depending on whether the (global) evaluation most closely associated with a thick term is held to be a part of the semantic concept of the term or not, that is, depending on whether thick terms and concepts are held to be *inherently* evaluative. Content views claim that they are, pragmatic views that they are not. Central to the issue here is linguistic data concerning the use of *objectionable* thick terms and concepts.

This issue is best introduced by considering Väyrynen’s argument for the pragmatic view. Väyrynen takes evaluation to be “information to the effect that something has a positive or negative standing—merit or demerit, worth or unworthy—relative to a certain kind of standard ... the kind that is capable of grounding claims of merit or worth” (Väyrynen 2013: 29). And a term or concept T is inherently evaluative on this view if “all literal uses of sentences of the form x is T in normal contexts entail, as a conceptual matter or by virtue of a semantic rule, that x is good or bad, depending on T in a certain way” (2013: 34). This last suggests a way to test whether T-evaluations are a part of the meaning of thick terms and concepts.

If T-evaluations were a part of the meaning of thick terms and concepts, then we would expect them to be semantic entailments, semantic presuppositions or conventional implicatures. Standard methods in linguistics give us ways to test for these things. If they turn out to be none of these things, then they are related to thick terms and concepts pragmatically.

Here is where objectionable thick terms and concepts come into the argument for the pragmatic view. A thick term or concept is objectionable if the global evaluation most closely associated with it doesn't fit the things which the term or concept is true of. For example, we might think that 'lewd' is objectionable because overt displays of sexuality are not bad for being such displays. The linchpin of the argument is a hypothesis that linguistic data concerning objectionable thick terms and concepts shows that T-evaluations *project*—they survive embedding in various contexts that cancel semantic entailments—and are *defeasible* in certain ways that semantic entailments aren't. If someone says "Miley's show is lewd", we'd be inclined to think that they thought the show was bad in a certain way. But is that it is bad in a certain way *entailed* by 'lewd'? Väyrynen argues that it is not, on the grounds that because the global evaluation projects in all of the following cases: "Is Miley's show lewd?," "Miley's show might be lewd" and "If Miley's show is lewd, the *Daily Mail* will be all over it." And one can perfectly well say things like "Miley's show is lewd, but not bad in any way."<sup>3</sup>

Projection and defeasibility are also key to Väyrynen's argument that T-evaluations are not conventional implicatures or semantic presuppositions. He argues that the best explanation of the behaviour of T-evaluations is that they are related to thick terms and concepts pragmatically. Specifically, his view is that global T-evaluations are implications of utterances containing thick terms which are part of background, 'not-at-issue' content in normal contexts and which arise conversationally.

Broadly speaking, there are two different strategies for replying to this argument. The first can be used by defenders of either the reductive or the non-reductive content views, and this is to directly deny that the data concerning objectionable thick terms and concepts supports the pragmatic view by challenging either the supposed data itself, or that the data shows that the relevant evaluations are not part of the semantic content of the relevant terms. Brent Kyle is an example of someone who has pursued this strategy, and he introduces further linguistic data which he argues supports what he calls the semantic view (Kyle 2013).

The second strategy for replying to Väyrynen's argument can be revealed by considering the question of whether *thin* terms and concepts can be found objectionable, that is, as embodying evaluations that ought not to be endorsed. Take 'morally good' for example. Perhaps a Marxist or a Nietzschean can think that there is such a property but offer an account in which moral goodness is something to be condemned. Any difficulty we have in imagining that moral goodness could be found objectionable could be accounted for by the imaginative difficulty in occupying evaluative perspectives very remote from our own.

We could use the case of a Nietzschean finding 'morally good' objectionable to argue that the global evaluation associated with 'morally good' projects is defeasible in the same way as the global evaluation associated with 'lewd'. Someone who found 'morally good' objectionable would be unlikely to say any of the following: "Is helpfulness morally good?," "Helpfulness might be morally good" or "If helpfulness is morally good, Bryan will help Patrick." That is, they would be unlikely to use the term in a question, embedded in a possibility modal or epistemic modal, or in the antecedent of a conditional because the positive evaluation projects. And it seems possible to deny that an action is good for being morally good.

The crucial point is that this would not necessarily show that 'morally good' is not an evaluative concept or that it fails to pick out an evaluative property. To reject a concept, as Dancy puts it, one has to see its point and reject it for that reason (Dancy 1995: 269). Given that we are talking about evaluative concepts, another way to put this would be

that one has to make the evaluation and then evaluate that evaluation as objectionable. A rough-and-ready characterisation of what is going on here could be that there are two evaluations: one conceptual and one substantive. If we accept Väyrynen's account of such data, what this would show is that any substantive evaluation of 'morally good' is plausibly not itself a part of the content of 'morally good'.

Might not the same be true in the thick case? If global evaluations are no part of the semantic content of thick terms, this might be a reflection of the fact that substantive evaluations of thick properties are no part of the concept that picks out those properties. We could take as evidence for this the case familiar in the literature of the bully who is quite competent with 'cruel' but evaluates cruelty positively.

In fact, this is part of the pragmatic view as Väyrynen defends it: that the substantive evaluations that groups of people make are no part of the content of the relevant term or concept but instead ought to be viewed as communicated via pragmatic mechanisms. However, importantly for the defenders of content views, this by itself does not answer the question of whether the thick concepts or properties *themselves* are evaluative, any more than a substantive evaluation of 'morally good' being no part of its content should lead us to conclude that 'morally good' fails to be an evaluative concept or pick out an evaluative property.

### What makes a concept evaluative?

The previous section brought to light an important fundamental issue raised by the thick concepts debate, that is, the question of what makes a concept evaluative in the first place. This is a tricky issue, and again one where it looks to be difficult to find an account that is appropriately metaethically neutral, and neutral between the different views in this debate. This is particularly difficult as at least some non-reductive content views hold, or at least suggest, that there is no sharp distinction to be made between evaluation and non-evaluative description (see, e.g., Putman 2002: chapter 2; Dancy 1995; Roberts 2013a).

Because, historically, much of the thick concepts debate merely assumed that thick concepts were evaluative concepts, this issue is relatively under-explored. The second reply to Väyrynen's argument in the previous section relied on the intuitive, pre-theoretical view that the concept MORALLY GOOD is an evaluative concept. But MORALLY GOOD is thicker (richer, more informative) than GOOD. Perhaps then it is open to defenders of the pragmatic view to argue that MORALLY GOOD is itself a non-evaluative descriptive concept, merely one that we tend to think is evaluative because it is so closely associated with the substantive evaluation that it is good.

For the purposes of this chapter, it is important to note that, given the nature of the debate, we cannot rely on our pre-theoretical intuitions about which concepts are evaluative, any more than we can rely solely on our intuitions concerning which concepts are thick. Part of any account of the thick, then, should be an account of what makes evaluative concepts evaluative in the first place (see Eklund 2013; Roberts 2013a; Väyrynen 2013: chapter 2).

### Does evaluation drive extension?

As noted above, that thick concepts are evaluative is assumed without argument in most of the literature on the thick. Those wishing to defend this assumption, and thus some version of the content view, might look to the considerations cited in favour of



the common claim in the literature that evaluation drives extension for thick terms and concepts. Borrowing Väyrynen's terminology, I consider two here: *parochiality*, or the claim that the point of these concepts is only apparent from within a particular evaluative point of view, and *underdetermination*, the way in which their non-evaluative content underdetermines their extensions.

Evidence for parochiality is taken to come from the plausible thought that an outsider would not be able to see how to 'go on' in applying the concept to new cases without at least imaginatively adopting the *evaluative* perspective within which the concept has its point. This is best illustrated by considering examples of thick concepts historically and/or culturally remote from our own, like perhaps the Quaker value of simplicity or the Southern African value of *ubuntu*.

Parochiality is closely related to the claim that the non-evaluative components of thick terms and concepts underdetermine their extensions. Here, the thought is that even the most detailed non-evaluative description which a given thick term may plausibly be thought to entail will not be sufficient to determine the extension of the term. Consider 'distributively just', an example that Elstein and Hurka discuss. If a distribution is just, that seems to entail nothing more non-evaluatively speaking than that it is a distribution. Or take 'dainty', which appears to entail nothing more than 'small,' non-evaluatively speaking. One plausible explanation for this is that it is evaluation as part of the semantic content of the relevant terms that is at least partially determining extension.

Gibbard illustrates this thought nicely in his discussion of *GOPA*, his fictional example of a thick concept used by the (fictional) Kumi tribe for certain kinds of killings:

I allow there might be descriptive constraints on a thick concept. Calling the eating of vegetables *gopa*, say, might violate a rule of language: it may just not be the sort of thing that could count as *gopa*. These constraints, though, will in many cases be too sparse to give the whole meaning—even when combined with an evaluative component. A good treatment of thick terms may combine clearly descriptive and clearly evaluative elements—indeed, my own analyses will. My claim has been that a statement like 'this killing was *gopa*' has no descriptive meaning as a whole—or not enough of one that, in some combination with its evaluative meaning, it can yield the full meaning of the statement. There is no statement 'This act was descriptively *gopa*' which somehow combines with an evaluative meaning to give the full meaning of 'This act was *gopa*.' In thick concepts, descriptive and evaluative components intermesh more tightly than that.

(1992: 277–278)

Now it may seem, on the face of it, that both parochiality and underdetermination tell strongly in favour of holding that evaluation plays an extension-determining role for thick terms and concepts by being part of their content. The pragmatic view is, however, not without resources here. Väyrynen, for example, has argued that the pragmatic view, together with general features of context-sensitive gradable terms (and he argues that thick terms are plausibly context-sensitive gradable terms) can also explain these considerations. He then argues on the grounds of Grice's razor—that other things being equal, we should prefer the view that postulates fewer semantic properties—that the pragmatic explanation is to be preferred.

Defenders of reductive and non-reductive content views may reply to this in a number of ways. First, this account of parochiality and underdetermination depends on the success of the argument for the pragmatic view discussed in the section ‘Objectionable thick terms and concepts’ above. If the direct response to that argument (that is, denying that the data concerning objectionable thick terms and concepts supports the pragmatic view) succeeds, then this account of parochiality and underdetermination is unavailable. Second, it’s worth noting that a term or concept’s being evaluative is compatible with its being a context-sensitive gradable term—thin evaluative concepts fall precisely into this category. If, as I discussed in ‘Objectionable thick terms and concepts,’ it is substantive evaluations that are communicated pragmatically, the pragmatic view combined with general features of context-sensitive gradable terms could nonetheless be combined with the view that thick concepts and properties are evaluative.

This is perhaps just a different way of making the same point that Väyrynen makes, which is that the issue of whether thick concepts are *themselves* evaluative is really to be decided by the employment of a generalisation of Grice’s razor; *other things being equal* we should prefer the pragmatic view. But are other things equal? In the rest of this section, I aim to show that there is as yet no reason to be confident that they are, so both reductive and non-reductive content views remain as live options.

### Evaluation: Global and embedded?

Whether thick terms and concepts contain embedded evaluation, regardless of whether they also contain global evaluation, is an important issue. Broadly speaking, it’s important for the pragmatic view that if there are any thick terms that contain embedded evaluation, that this isn’t a characteristic of paradigmatic thick terms in general. And it is important for both the pragmatic and the reductive content views that this embedded evaluation is thin and that it can in principle be “unembedded,” i.e., separated out from the non-evaluative content of the thick term. I explain both these points in more detail below.

The main argument for the pragmatic view, discussed above, targets only global evaluation. Väyrynen allows that paradigmatic thick terms and concepts may contain embedded evaluation but holds that thick terms and concepts won’t be inherently evaluative, or have deep or distinctive significance, if their meanings only contain embedded, and not global, evaluations. This is for two reasons, he claims. First, he suggests, only some thick terms have embedded evaluation, so embedded evaluation cannot be what is in general distinctive of the thick. Second, he follows Elstein and Hurka (defenders of a reductive content view) in claiming that because embedded evaluations are independent of whether they figure in the meaning of thick terms and concepts, if embedded evaluations had deep or distinctive significance, then they’d have this independently of their relation to the thick terms and concepts that embed them.

Both these reasons can be questioned, with potentially problematic results for both pragmatic and reductive content views. First, the non-reductivist can attempt to make a case for the claim that all (or just most) paradigmatic cases of thick concepts contain embedded evaluation. Second, regardless of whether this case is successful, the claim that embedded evaluations couldn’t give thick terms and concepts deep or distinctive significance because they are independent of thick terms and concepts that embed them is

really only plausible if (i) all embedded evaluations are *thin*, and (ii) if they are genuinely embedded, that is, could at least in principle be unembedded. And the non-reductivist can attempt to deny these things.

Defenders of non-reductive content views might well ask why we should assume either (i) or (ii). On (i), if some embedded evaluations are *themselves* thick, why assume that there will be some chain of embeddings which bottoms out in an analysis that includes only thin evaluation and non-evaluative description?<sup>4</sup> Moreover, why assume that there are neat conceptual divisions such that each embedded evaluation can be regarded as independent? Our evaluative concepts may form a network with no such sharp divisions.

For example, grasping the concept 'courage' may involve grasping when it is *worthwhile* to pursue *goods* regardless of the badness of the *harm* or risk of *harm* one is accepting. Grasping 'harm' may require grasping what *welfare* and *flourishing* and the *worthwhile* are. And grasping 'welfare', 'flourishing' and 'the worthwhile' may require grasping what *courage* is.

On (ii), if so-called embedded evaluations cannot be separated from that in which they are embedded, it may not make sense to think of them as being independent or, indeed, as embedded.

At this point, it is worth noting that it may be that what is referred to as embedded evaluation, and what is dismissed as thus being unable to ground claims concerning the distinctive significance of thick concepts, may well have been what, historically, non-reductivists about the thick were concerned to focus on. That is, the issue was never really primarily whether the relevant global (thin) evaluation could be detached (be separated out by analysis on the reductive content view, or shown not to be part of the semantic content on the pragmatic view). The issue was whether the very object of the global evaluation was *itself* evaluative. In my view, the way to put this point is that the real issue in this debate is whether the object of the global evaluation is an evaluative *property*.

This gets to the heart of what at least some defenders of irreducible thickness were concerned about, that is, showing that kindness, cruelty, courage and the like are genuine evaluative features of the world and thus a problem for anti-realist views (see, e.g., McDowell 1981; Dancy 1995; Putnam 2002). The thought might be that no one, not even the metaethical anti-realist, wants to claim that kindness, generosity, cruelty, harm, suffering, brutality and courage are not genuine features of the world (using the relevant robust standard of genuineness), though they may be happy to do so in the case of thin properties. So thick terms and concepts represent the possibility of a distinctive challenge to the sharp ontological, and perhaps also semantic, fact-value distinctions made by these sorts of views, if it can be shown that the object of the global evaluation, for any thick concept, is *itself evaluative* and evaluative in such a way that cannot be reductively analyzed into some combination of thin evaluation and non-evaluative description.

### Shapelessness

Another common claim, typically made by non-reductivists, is that thick terms and concepts are *non-evaluatively shapeless*. Assuming that thick concepts are themselves evaluative, this is thought by some to tell in favour of the non-reductive view.

What is non-evaluative shapelessness? The 'shape' of a concept, we can say, is what all the things falling under the concept have in common: the unifying feature or the real

resemblance (Roberts 2011: 505). If evaluative concepts are non-evaluatively shapeless then that unifying feature or real resemblance is *not* non-evaluative: it is not simply that we lack a term for the non-evaluative feature of the world that we are nonetheless sensitive to in applying our evaluative concept, *it is not there*.<sup>5</sup>

Many have found the shapelessness thesis plausible (see, e.g., McDowell 1981; Dancy 1993, 1995; Wiggins 1993; Putnam 2002; Majors 2003, 2009; Kirchin 2010; Roberts 2011, 2013a). If shapelessness is the case, does this tell against either the reductive content view or the pragmatic view? The answer here is bound up with the issue of embedded evaluation discussed in the previous section. If we assume that Elstein and Hurka and Väyrynen are correct about embedded evaluation, then I think it is safe to say that both the pragmatic view and the reductive content view can accommodate shapelessness. But if they are not, then I think that shapelessness will tell in favour of the non-reductive view.

(vi) *What makes concepts thick? The list revisited*

On the pragmatic view, thick terms and concepts are not inherently evaluative. One might then ask why (paradigmatic) thick concepts have been so universally thought to be evaluative. Väyrynen's answer is that thick terms and concepts concern aspects of human life that matter: they "are widely invested with evaluative significance or systematically evoke various affective responses that are connected to evaluation" (Väyrynen 2013: 133). On this view, there is no difference in kind between 'courageous', 'cruel' and 'loyal,' on the one hand, and 'fun', 'athletic' and 'chocolate,' on the other, where these are non-evaluative terms commonly used to imply positive evaluation. Väyrynen holds that we can explain the appearance of a difference in kind by noting a cluster of contingent differences of degree. Chief among these are the differences in the degree to which the relevant evaluations are generalised: the greater the degree of generalisation, the easier it is for a term or concept to appear inherently evaluative.

A response open to reductive and non-reductive content views here is that it leaves it very mysterious that 'pleasure' and 'pain' aren't paradigmatic thick terms. These terms would seem to satisfy all the relevant criteria for being a paradigmatic thick term on Väyrynen's account. For one, they are typically used against a very broad and very strong background of agreement regarding their evaluative significance.

Väyrynen offers two speculations as to why 'pleasure' and 'pain' aren't regarded as thick: that the properties ascribed by thick terms are less easy to characterise in general terms, and that many of us are confident that it is no part of what it is to be pain that it be bad and no part of what it is to be pleasant that it should be good and thus see no essential need to use evaluative notions in describing pleasure and pain. The worry remains, though, for another way to put this might be that while we are confident that pain and pleasure are non-evaluative, the same cannot be said for the paradigmatically thick.

A plausible explanation here may have to do, again, with so-called embedded evaluation. Perhaps paradigmatic thick terms and concepts are those that involve this kind of evaluation. Once again, there seems to be a reason for further exploration of this issue. Moreover, appealing to embedded evaluation may be what explains the pre-theoretical intuition that complex predicates like 'bad dancer' and 'good assassin' are not thick, and these also come out as thick on Väyrynen's view.

Let us pause for a moment to take stock. The data concerning objectionable thick concepts does not conclusively favour any particular view of the nature of the thick. Moreover, all three positions are able to explain parochiality, underdetermination and

shapelessness. Other things being equal, employing Grice's razor at this point means the pragmatic view comes out on top. However, there are further considerations. There is the issue of which concepts get to count as thick, as just discussed above. The pragmatic view is possibly vulnerable on this point for including too many concepts we don't ordinarily think of as evaluative or thick, and for not having a convincing rationale for excluding some that are not—i.e. pleasure and pain. At this point, plausibly, content views are out ahead. The issue of which concepts count as thick is connected to the interesting question of what makes concepts evaluative in the first place. In my view, this question is unlikely to yield non-question-begging answers.

The last remaining consideration is that of so-called "embedded evaluation". Many, if not all, of the examples of thick terms and concepts which are given as paradigmatic examples are plausibly understood as requiring evaluative information to characterise the very thing that is then subject to global evaluation. For these paradigmatic examples, it is very difficult not to reach for further thick terms when characterising that thing, as Elstein and Hurka inadvertently reveal in using the thick term 'harm' in their supposed reductive analysis of courage. Exactly what the upshot of this is, is an issue that requires more careful consideration than I have space for here. Suffice it to say that this issue concerning the nature of embedded evaluation (including whether it really makes sense to call it that) should now be centre stage in the debate about the nature of the thick.

### WHY ARE THICK CONCEPTS IMPORTANT?

Precisely because these terms and concepts combine evaluation and non-evaluative description, some have held that they have significance beyond their intrinsically interesting nature. It's held to be distinctive because it is supposedly significance that isn't shared by thin concepts. Arguments for the distinctive significance of thick concepts have focussed on thick ethical terms but if there turn out to be thick concepts in other normative domains, then presumably the same kinds of significance would follow.

Typically, those who think that thick concepts have distinctive significance hold the non-reductive content view. Defenders of the pragmatic view deflate the significance of the thick by claiming that thick concepts are not inherently evaluative. And those who defend reductive content views claim that the evaluative and non-evaluative component of thick concepts can, at least in principle, be separated out by analysis. Non-reductive content views of the thick are most at home within a non-reductivist robust realist view in metaethics. Defenders of other kinds of metaethical views are likely to be drawn to one or other of the reductive content or the pragmatic view, depending on how wedded they are to the idea that thick concepts are inherently evaluative.

What then are the distinctive significances that (irreducibly) thick concepts have been claimed to have?

#### Fact and value

There are at least two of what may be loosely termed "fact-value" distinctions that thick terms and concepts have been held to break down. If these terms and concepts are

successful in this regard, then they constitute an objection to any metaethical view that requires either or both of these distinctions.

First, we might make a *semantic* distinction between two different kinds of meaning, two different kinds of things we can do with language: we can evaluate and we can describe. One way to understand expressivism is as the view that evaluative utterances express non-cognitive attitudes, and that evaluation is thus not a species of description. On this understanding, expressivism requires a sharp distinction between evaluation and description. If thick terms and concepts are both evaluative and descriptive in such a way that they don't combine evaluative and non-evaluative descriptive elements that can in principle be separated from each other, then this is a problem for expressivism so understood. Many read Williams, McDowell, Putnam and Dancy, among others, as making this sort of thick concept argument against expressivism.

Secondly, we might make a related *ontological* distinction concerning what we admit into our ontology as genuine features of the world. The view that evaluative utterances are not truth-apt usually goes along with the view that there aren't any evaluative properties (in some appropriately robust sense) in the world. There are of course all sorts of other properties on this view, properties which, loosely speaking, are what the facts are about. Expressivism, as characterised above, thus also accepts a sharp ontological distinction between fact and value in this sense. If thick concepts were to undermine this ontological distinction, then this is a problem for expressivism. But it is a problem, too, for any anti-realist metaethical view that denies the existence of evaluative properties (compare McDowell 1981; Putnam 2002; Wiggins 1993; Dancy 1995). Error theories and fictionalist views would thus also be vulnerable to a thick concept argument.

In what way would all these anti-realisms be vulnerable? Recall from the discussion of so-called embedded evaluation above that the issue there was, for any thick concept, whether the very object of the global evaluation is *itself* evaluative. I said that the best way to put this is that the real issue in this debate is whether the object of the global evaluation is an evaluative *property*, for this gets to the heart of what at least some defenders of irreducible thickness were concerned about, that is, showing that kindness, cruelty, harm and the like are genuine evaluative features of the world. It is a significant cost to a view, one might think, if it has to claim that kindness, cruelty, harm and the like are not genuine features of the world (by the relevant robust standard of genuineness).

Interestingly, if thick concepts are a problem for anti-realist views because they're irreducibly thick and undermine an ontological distinction between fact and value, then they may also turn out to be a problem for *reductive realist* views in metaethics (Roberts 2011, 2013a). Moreover, thick concepts may even be recruited to undermine the idea that the evaluative, in general, supervenes on the non-evaluative (Dancy 1995; Roberts Forthcoming).

Of these two challenges to "fact-value distinctions," the second looks to be more fundamental and potentially more powerful, if it can be brought off, and there is still much interesting work to be done.

### **Evaluative practice, conceptual priority and normative ethics**

If thick terms and concepts are evaluative in their own right, and they do not combine non-evaluative and evaluative elements of content in a way that can in principle be

separated out, then this has been thought to have implications for evaluative thought and discourse, and normative ethical theories.

A large range of ethical theories focuses primarily, even exclusively, on thin terms and concepts. It is at least implicit in these views that all moral judgements can be expressed using thin terms and concepts, and that the thin is primary in some sense. It is also sometimes explicitly a part of these theories that thick moral concepts can be derived from the thin. If one is attracted to these kinds of views in normative ethics, one is likely to be attracted to reductive content or pragmatic views of the thick.

If thickness is irreducible, this may generate problems for the relevant theories in normative ethics. For example, we might worry, as some anti-theorists and virtue ethicists have, that ethical theories that focus only on the thin impose an oversimplified, distorting structure on ethical thought and the complexity of ethical life as we actually live it (compare Anscombe 1958; Williams 1985). In addition, we might think that the traditional approach to ethical theory is thin-centric, in that it assumes that the thin is somehow (conceptually, explanatorily) fundamental. If thick concepts are not derivable from the thin, we might question this thin-centralism and perhaps advocate either thick centralism or some kind of no-priority view in its place (compare Hurley 1989: 11).

### Ethics and beyond

Lastly, whatever account of thick concepts turns out to be correct is likely to have implications for accounts of particular virtues, vices and values in ethics—and in other domains if the distinction between thick and thin turns out to hold there too. So, there are potential implications for accounts of various concepts important in ethics, applied ethics, politics, law, aesthetics and epistemology. To mention just a few potential examples, some repeated from above: CONSENT, RESPECT, DIGNITY, HARM, MORAL STANDING, LEGITIMACY, NEGLIGENCE, LEGAL, INTELLECTUAL HUMILITY and TRUSTWORTHY. Arguably, the implications for accounts of such concepts will be most profound if thickness turns out to be irreducible.

### NOTES

1. Strictly speaking, the pragmatic view holds that it is the global evaluations most closely associated with thick terms and concepts that are not part of their content. See below.
2. This is a claim that is often made in this literature, and it is not always precisely clear what it means. There are at least two ways it has been interpreted: that evaluative and non-evaluative content/information cannot even in principle be separated out by conceptual/semantic analysis, or that the property picked out by the relevant term is itself an irreducibly evaluative property.
3. “Guy loves hummus” entails that Guy has very positive feelings towards hummus. Compare “Does Guy love hummus?” “Guy might love hummus” and “If Guy loves hummus, then he’ll enjoy this restaurant.” The entailment does not project. And nor is it plausibly defeasible: “Guy loves hummus but he doesn’t have any positive feelings toward it” is infelicitous.
4. Thanks to Tristram McPherson for this point.
5. This requires adopting a certain view of the nature of properties on which massive disjunctions, for example, are not properties.

### ACKNOWLEDGEMENTS

Thanks to the editors, James Fritz and Guy Fletcher for extremely helpful comments on an earlier draft.

## REFERENCES

- Anscombe, G. E. M. (1958) 'Modern Moral Philosophy', *Philosophy* 33 (124), 1–19.
- Blackburn, S. (1992) 'Through Thick and Thin', *Proceedings of the Aristotelian Society*, Supplementary Volume 66, 285–299.
- Dancy, J. (1993) *Moral Reasons* (Oxford: Blackwell).
- Dancy, J. (1995) 'In Defense of Thick Concepts', in P. A. French, T. E. Uehling Jr. & H. K. Wettstein (eds.) *Midwest Studies in Philosophy XX: Moral Concepts* (Notre Dame: University of Notre Dame Press), 263–279.
- Eklund, M. (2011) 'What Are Thick Concepts?', *Canadian Journal of Philosophy* 41, 25–49.
- Eklund, M. (2013) 'Evaluative Language and Evaluative Reality' in S. Kirchin (ed.) *Thick Concepts* (Oxford: Oxford University Press), 162–181.
- Elstein, D. & Hurka, T. (2009) 'From Thick to Thin: Two Moral Reduction Plans', *Canadian Journal of Philosophy* 39, 515–536.
- Gibbard, A. (1992) 'Thick Concepts and Warrant for Feelings', *Proceedings of the Aristotelian Society*, Supplementary Volume, 66, 267–289.
- Hare, R. M. (1952) *The Language of Morals* (Oxford: Clarendon Press).
- Hare, R. M. (1963) *Freedom and Reason* (Oxford: Clarendon Press).
- Hare, R. M. (1981) *Moral Thinking: Its Levels, Methods and Point* (Oxford: Oxford University Press).
- Hurley, S. L. (1989) *Natural Reasons: Personality and Polity* (New York: Oxford University Press).
- Kirchin, S. (2010) 'The Shapelessness Hypothesis', *Philosophers' Imprint* 10 (4), 1–28.
- Kyle, B. (2013) 'How Are Thick Terms Evaluative?', *Philosophers' Imprint* 13 (1), 1–20.
- McDowell, J. (1981) 'Non-Cognitivism and Rule Following' in his (1998) *Mind, Value, and Reality* (Cambridge, MA: Harvard University Press), 198–218.
- Mackie, J. L. (1977) *Ethics: Inventing Right and Wrong* (London: Pelican Books).
- Majors, B. (2003) 'Moral Explanation and the Special Sciences', *Philosophical Studies* 113 (2), 121–152
- Majors, B. (2009) 'The Natural and the Normative' in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics* (Oxford: Clarendon Press) 4, 29–52.
- Putnam, H. (2002) 'The Entanglement of Fact and Value' in *The Collapse of the Fact/Value Dichotomy and Other Essays* (Cambridge, MA: Harvard University Press), 28–45.
- Roberts, D. (2011) 'Shapelessness and the Thick', *Ethics* 121, 489–520.
- Roberts, D. (2013a) 'It's Evaluation, Only Thicker' in S. Kirchin (ed.) *Thick Concepts* (Oxford: Oxford University Press), 79–96.
- Roberts, D. (2013b) 'Thick Concepts', *Philosophy Compass* 8, 577–588.
- Roberts, D. (Provisionally Forthcoming) 'Why Believe in Normative Supervenience?' in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics* (Oxford: Clarendon Press), 13.
- Scheffler, S. (1987) 'Morality through Thick and Thin: A Critical Notice of *Ethics and the Limits of Philosophy*', *Philosophical Review* 96, 411–434.
- Väyrynen, P. (2013) *The Lewd, the Rude and the Nasty: A Study of Thick Concepts in Ethics* (Oxford: Oxford University Press).
- Wiggins, D. (1993) 'A Neglected Position' in J. Haldane & C. Wright (eds.) *Reality, Representation and Projection* (Oxford: Oxford University Press), 329–336.
- Williams, B. A. O. (1985) *Ethics and the Limits of Philosophy* (Cambridge MA: Harvard University Press).