

Thick Concepts, Analysis, and Reductionism

Edward Harcourt
Faculty of Philosophy, Oxford University, and Keble College, Oxford

Alan Thomas
Tilburg University

published in Simon Kirchin (ed.), *Thick Concepts* (Oxford: Oxford University Press, 2013)

“Thick concepts” has emerged as a topic of interest because it sits on a fault-line in metaethics. Some utterances which deploy “thick” ethical vocabulary – terms such as “obsequious”, “small-minded”, “priggish”, “resilient”, “tolerant”, each with different and comparatively specific satisfaction-conditions – have been claimed to express ethical *concepts* thanks to which the thoughts and judgments expressed by these utterances on the one hand are candidates for truth and falsity (“world-guided”) and, on the other, supply defeasible reasons for action (“action-guiding”). (See for example Williams 1985, p. 141). But it is a widely-held doctrine in metaethics that no single thing can play both roles. So there is considerable pressure to find an alternative account of the same phenomenon.

The alternatives take two basic forms. One is ~~to offer~~ an account of the thoughts and judgments expressed by “thick” ethical vocabulary according to which their truth-evaluability and their relation to action are carried by two independently specifiable components, typically – though only typically - something with a truth-value plus an expressive or otherwise non-cognitive extra. s–The other affirms s the simplicity of these thoughts and judgments while denying the link to action. Thise latter option is externalism in the theory of motivation. As we will assume for present purposes that externalism is unattractive, it looks as if either the former alternative

Formatted: Indent: First line: 1.27 cm

must be made to work, or else utterances deploying “thick” ethical vocabulary really do – contrary to the doctrine – express ethical concepts of the kind we have outlined.

Our aim in this paper is to defend the latter view. However, some [preliminary](#) commentary is required ~~de first of all~~ on the way the dispute about “thick concepts” is often set up.

1 How *Not* to Identify the Issue

Let us assume for the time being that the ethical concepts whose existence we wish to defend are usefully labelled “thick ethical concepts”, or “thick concepts” for short. (We will return to this assumption later.) Thick concepts are commonly defined as concepts that “have both a descriptive and an evaluative component”, perhaps with the qualification that these components are ~~an~~ “compound rather than a mixture” (Blackburn 1992) or an “indissoluble amalgam” (Payne 2006). ~~W~~However, we object to this kind of definition, with or without the qualification, ~~in either strength~~, because the word “descriptive” has acquired a technical usage in philosophy according to which it ~~has come simply to means~~ “non-evaluative”. Defining thick concepts as concepts that have both ~~an~~ (in this sense) descriptive and an evaluative component thus implies that, even if utterances deploying thick vocabulary *do* express judgments deploying thick concepts, the fact that judgments deploying these concepts express evaluations can be no thanks to whatever it is that makes them truth-evaluable. So our position, namely, that what it ~~i-~~s for judgments to deploy ethical concepts is for them to be apt to express evaluations precisely thanks to what makes them truth-evaluable,ⁱ goes out of the window with the mere definition of the topic.

Formatted: Font: Italic

A related and equally objectionable idea is that the most visible version of the view we are opposing here – the claim that what is expressed by utterances deploying “thick” vocabulary resolves into a truth-evaluable component plus a non-cognitive extra – is properly characterized as the claim that thick *concepts* are capable of a two-component reductive analysis. When “bachelor” is analyzed into “unmarried person” and “man”, both components of the analysis are themselves concepts. But when it comes to what is expressed by thick ethical vocabulary, on the view now in focus this is not so. For, assuming a connection between conceptually structured judgements and truth, the second component of the analysis is precisely *non-conceptual*, since by the non-cognitivist’s lights whatever connects the judgment to action cannot be truth-relevant. But now ~~if~~ ^{if} the two supposed analysantia of what thick vocabulary expresses are a conceptual and a non-conceptual item, what it expresses is not, properly speaking, a thick *concept* after all. Those who deny that utterances involving thick ethical vocabulary express thoughts and judgments with the two features we claim for them ought to deny that there are any thick ethical concepts, not concede their existence and then offer a reductive analysis of them.ⁱⁱ

~~This complaint connects with (Blackburn 1992 is on the ball here — his project is showing the eliminability of (supposed?) thick concepts.)~~

~~This point connects with~~ Elstein and Hurka’s helpful point useful recent thought (2009) that two- component approaches (or indeed multiple-component ones, but we suppress this qualification henceforth for the sake of brevity) to the thick are not the exclusive preserve of non-cognitivist reductionism (Elstein and Hurka 2009). They observe that though non-cognitivism *requires* a two-component approach, it-it

Formatted: Indent: First line: 1.27 cm

Formatted: Font: Italic

[is's](#) also open to “Moorean” cognitivists, who hold that the word “good” expresses a concept apt to figure in truth-evaluable judgments but that is also simple and unanalyzable, to argue that the concepts expressed by thick ethical vocabulary have a reductive analysis (into “good” plus some non-evaluative component). That seems exactly right about “Mooreans”, but Elstein and Hurka do [notn't](#) quite go far enough. If non-cognitivists ought properly speaking to deny that there are any thick concepts, reductive analyses of thick ethical concepts are in fact on offer *only* from “Mooreans”, because only they combine a two-component approach to what the thick vocabulary expresses (and so have an analysis to offer) with the claim that both its evaluative and its non-evaluative functions are carried by semantic features of the utterance (and so can properly claim that there are some thick concepts, rather than merely some uses of ‘thick’ vocabulary, to analyze). We reject two-component approaches whether in their “Moorean” (properly speaking reductionist) or in their more visible, non-cognitivist (properly speaking eliminativist) version, though for ease of comprehension we use ‘reductive’ and ‘two-component’ interchangeably as there is no contradictory of ‘two-component’ ready to hand.

~~A (The situation gets worse, not better, if ‘non-evaluative’ is substituted for ‘descriptive’ in its technical sense: in reply, we would need to substitute ‘implies’ for ‘strongly suggests’.)~~ A distinct reason to protest at the way the topic is usually presented relates not to the thick but to the thin, and this [point](#) applies [as much](#) to Williams’s definition ~~as much~~ as it does to Blackburn’s or Payne’s. “Action-guiding and world-guided” as a characterization exclusive to the thick strongly suggests that ~~the~~ thin [ethical](#) concepts are [notn't](#) world-guided since, on pain of abolishing the thick/thin distinction, there must be something that thick concepts have as a class which thin ones do not. So, for example, an utterance deploying only the

paradigmatically thin term “good” will (just as non-cognitivists say) not express anything that’s a candidate for truth.ⁱⁱⁱ That is far too substantial a conclusion to get out of a mere definition. In order that the discussion not be skewed in favour of the non-cognitivist from the very start, therefore, this familiar definition of the thick/thin contrast is not tenable.

But why persist in the assumption we have been making so far, that the thick/thin distinction is needed at all? One bad reason for distinguishing the thin from the thick would be to appeal to the idea that thin evaluative concepts distinctively express verdictive ethical considerations (Foot, 2002). If that were so, it would imply that judgments deploying thin concepts are, distinctively, not *both* truth-evaluable *and* expressive of reasons for action, on the grounds that they are not expressive of reasons for action. A compressed argument to this conclusion is that verdictive judgments expressed by the word ‘right’, for example, express a verdict on one’s reasons and cannot in themselves add a further reason on pain of launching a regress of reasons. (Stratton-Lake, 2000, pp. 14-15; Dancy, 2004, pp. 16-17) However, while the evidential/verdictive distinction is a good one, we see it as marking off two distinct functional roles for a reason involving *any* kind of term (thick or thin), and thus as orthogonal to the thick/thin distinction. It cannot, therefore, be used to differentiate thin and thick judgements as a class.

In fact we agree with Scheffler that the distinction between thick and thin ethical concepts is best seen as a continuum. (Scheffler, 1987) If this is right, for a concept to be (relatively) thick is just for it to be more specific and to have narrower satisfaction conditions than one that is (relatively) thin, so there is no need to make differentiating claims about the marks of thick and thin concepts as classes. Moreover our non-reductive cognitivism about judgments deploying thick vocabulary applies

Formatted: Justified, Indent: First line: 1.27 cm

just as much to judgments expressed by paradigmatically thin vocabulary such as “good”. (This simply mirrors something the non-cognitivist can also say, namely that even “good” places *some* restrictions on the range of things that can satisfy it (as it might be, persons or states of affairs), but in so far as it does, even here there is room to unbundle what the word does into “description” – which earns the utterance its truth-conditions – and evaluation.) But it does not follow from any of this that there is

∴

~~To raise one final worry about how to set the issue up, let me return to thin concepts: is there anything the anti-reductionist (non-Moorean) cognitivist – that is, we – should say about thick concepts that we should not also say about thin ones? A number of commentators, following Scheffler (1989?), have commented that ‘thick’ and ‘thin’ is better seen as a continuum than as a distinction. Thick concepts, on this view, are just more specific or have narrower satisfaction conditions than thin ones. This is not to say that there is~~ nothing interesting going on towards one end of the continuum that ~~i~~’s not also going on towards the other. ~~For~~ – for instance, the thicker the concept, the more likely it is to be merely local, and the less likely it is to be portable outside a larger web of sensibilities, understandings and, ways of life whose characterization will require further use of thick concepts. (Thomas, 2006, pp. 156-160) So we feel no sense of awkwardness in continuing to claim that the ethical concepts expressed by (relatively) thick vocabulary are indeed (relatively) thick ethical concepts, rather than just ethical concepts without qualification. (It is also consistent with this to maintain that the concepts expressed by *thin* ethical vocabulary are sometimes (relatively) thick – but that is another story.)

What of the converse worry, that there's nothing to distinguish ethical concepts (thick or thin) from concepts generally? Some [very general reflections on concept use, derived from the later Wittgenstein](#), can make it seem as if [the claims that non-reductive cognitivists make about thick ethical concepts stem from a generalized operationalism that ties all concept use to practicality](#), and so rope in non-ethical concepts as well as ethical ones. [Operationalism is at home in the philosophy of science as a distinctive kind of explanation of concept use: grasp of a concept is identical to grasp of "a certain set of operations"](#). (Bridgman, 1927) [But this view is open to decisive objections even in its proper domain.](#) (Chihara and Fodor, 1965) So we need to show that the practicality of ethical concepts rests on something different, and proper to those concepts in particular.

The answer is that [both thick and thin evaluative concepts presuppose that which David Wiggins and John McDowell have called a prior "evaluative interest" in a way that other concepts do not.](#) [Adrian Moore explains clearly why there have to be two arguments here, not one:](#)

[Practical reasoning, on this reconstruction, includes a pure element: keeping faith with concepts. Theoretical reasoning also involves keeping faith with concepts. What makes it possible for keeping faith with concepts to have a practical dimension as well as its more familiar theoretical dimension is, ultimately, the fact that some concepts –... ethical concepts – equip those who possess them with certain reasons for doing things.](#) (Moore, 2005a)

[Because of the role of an evaluative interest in explaining ethical concepts, the interpretation of both thin and thick concepts places special demands on the interpretation of social groups, demands that do not extend to the explanation of all concepts. Given the nature of the ethical, we can expect both classes of ethical concepts to be highly perspectival: they are tied to our human perspective and its](#)

distinctive peculiarities. Once again, however, it is a matter of degree that thick concepts place more demands on social explanations than do the thin. That is because they give you more specific information about the particular way in which, to use Bernard Williams's phrase, users of thick concepts "find their way around" a particular social world. This fact is simply a consequence of their comparative specificity compared to thin concepts. They capture that which is to be explained in a more fine-grained way than descriptions that deploy thin concepts.

_____ In this paper we cannot hope to say all that needs to be said in defence of non-reductive cognitivism about what thick ethical vocabulary expresses. In particular we do not address what, to some, fundamentally underpins opposition to it, namely the doctrine that nothing can be both truth-evaluable and bear the relation to action which we claim for ethical concepts. (If correct, this would of course rule out reductive ("Moorean") cognitivism about the thick as much as it would our own view.) We focus instead on considerations as it were upstream of that doctrine designed to make two-component accounts seem inevitable, arguing that our non-reductive view is at least as well-motivated as its rivals.

2 Evaluative Properties and the Determinate/Determinable Relation

One of the main arguments in favour of a two-component approach to judgments deploying thick vocabulary is that a two-component approach *follows from* the idea that the thick/thin contrast is a contrast between more and less 'descriptive content'. As long as 'more descriptive content' means 'greater specificity' or 'narrower satisfaction-conditions', this interpretation of the contrast is unimpeachably neutral, so we now want to show why it is no argument for a two-component approach.

Consider a thick ethical concept such as BRUTAL and an occasion of its use, such as “the prisoners were treated brutally”. We can form a corresponding judgement using one of its thin counterparts such as BAD, and substitute the judgement “the prisoners were treated badly” as a less specific description of the same situation. In the former case you know much more specifically the particular way in which the prisoners were treated badly and one can establish a defeasible tie to a reason for action (perhaps preventing the brutal treatment if it is in your power to do so) in virtue of the correct application of a more fine-grained description. Now it is tempting to imagine that what one is doing here is subtracting the extra descriptive content contributed by ‘brutal’ to isolate the element – whether it’s regarded as purely evaluative or not is another matter, though it usually is – expressed by ‘bad’. But this explanation is far from mandatory.

Consider an analogy with colours. ‘Red’ has more ‘descriptive content’ than ‘coloured’, in that whatever is red is coloured, but we don’t learn *which* colour something is – as it might be, red – just by being told it’s coloured. But there is no property distinct from redness itself such that something is red iff it has that property and is coloured. There is nothing one can conjoin with colouredness to yield redness; or to put the point the other way about, colouredness is not a *component* of redness. (Williamson, 2000, p. x; Jackson, 2002, p. 518ff). But since it is not, it is a mistake to infer from the mere fact that ‘brutal’ has more descriptive content than ‘bad’ that brutality can be analyzed into two components, badness plus some further independently specifiable condition. You could say, trivially, that *x* is brutal iff it is bad and brutal, but that is no comfort to a two-component approach.

Does this show that thick concepts are unanalysable? No: at this stage all we have shown is that a two-component analysis of what thick ethical vocabulary

expresses is not inevitable, [because the data is compatible both with that proposal and with an alternative](#), namely that ‘bad’ stands to ‘brutal’ as ‘coloured’ stands to ‘red’, at least in so far as the former is an *abstraction from* the latter (or the latter a *specification of* the former). However, we mention the alternative not just because its mere possibility shows no two-component approach can be inferred straight from the agreed facts about more and less ‘descriptive content’, but because we think this view is correct. To elaborate: [following Christine Tappolet, we argue that a *pro tanto* judgment using the word “good” stands to a counterpart judgment using a thick ethical concept](#) in a way analogous to that in which [the concept RED stands to the concept COLOURED, namely, the relation of a determinate to its covering determinable](#) at the level of the corresponding properties ([Johnson, 1921; Prior, 1949; Funkhouser, 2006](#)).

Why only “analogous to”? Because evaluative properties lack the kind of necessary inter-property relations within the entire field of properties that allowed Wittgenstein to speak of a “geometry of colour” in the case of colour concepts. (Wittgenstein, 1977) (There is no analogy, in the case of evaluative properties, to a colour wheel exhibiting these necessary relations within a holistic structure.) Nevertheless, the [relations between properties such as](#) the property of badness [and](#) the property of being brutal [preserve the truisms that Funkhouser lists as minimally required for the presence of this kind of determination relation between properties.](#) ([Funkhouser, ibid, pp. 548-549](#))

[First, it is a relation that holds relative to pairs of properties. So, on this account, the property of brutality is determinable relative to thuggish, but determinate relative to bad. Secondly, any determinate property is an instantiation of a determinable in some specific way, where specificity is distinct from, inter alia, the](#)

genus-species relation, so brutality is a determinate way of being bad. Thirdly, we agree with Judith Jarvis Thomson when she claims that:

It seems very plausible to think that a thing's being good must consist in its being good in some way (...) if that is the case then there is no metaphysically mysterious property goodness. (Thomson, 1996, pp. 128-129)

Compare: no object is just coloured without being some particular colour. This preserves the truism that an object which instantiates a determinate instantiates all determinables that it falls under, so a thuggish action is also brutal and bad. (Funkhouser, ibid.) The way in which the relation between properties picked out by thick and thin predications matches this profile seems to us a strong indication that our account is correct.

We would like to add a further, epistemic gloss to Thomson's ontological claim that to be good is always to be good in a particular way. It is an advantage of our account that it promises to explain, in a way that is favourable to cognitivism, a line of thought which has been made to bear too much weight in meta-ethics. Consider a uniformly red surface filling one's field of vision. In a certain obvious sense there is no further feature of the surface one needs to see beyond its redness in order to judge that it is coloured. If one were sufficiently impressed by this fact, therefore, it might be tempting to say that the colouredness of the surface is a property of it that cannot be seen. One would then need a theory of what sort of property colour is, and how it relates to the visible. But there's a simpler explanation of the facts stated. The reason there is no further feature of the surface one needs to see beyond its redness in order to judge that it is coloured is just the relationship between the two concepts: in someone who possesses both concepts, to see that it is red is to see that it is coloured. So there's no warrant for saying of colouredness, any more than there is of redness, that it is not a property that can be discerned by looking.

This allows one to see G. E. Moore as having taken a wrong turning precisely where the determinate/determinable analogy shows it to be avoidable. Moore's "intuition" that "good" is simple, indefinable, but non-natural (i.e. not discernible by the senses or by introspection) seems almost exactly to parallel the misguided construction of a theory of colouredness on the assumption that colour is not a visible feature. Goodness is a real property of things and, like "yellow", it is (let's grant) simple; but unlike "yellow" – so Moore reasoned - it is unobservable. So we must come to know it by "intuition". We do not want to argue that the property of goodness can literally be seen, but Moore overlooked the possibility that, like the exercise of "coloured", judgments of goodness might be the bringing to bear on the observable of a complex conceptual capacity, concluding instead that they must involve the exercise of a conceptually unstructured quasi-perceptual faculty which delivers knowledge of what cannot be perceived. That is one path towards an intuitionist moral epistemology that seems to us very unappealing. The positive model we have put forward explains how it can be avoided. If, as on our view, to be good is always to be good in some particular way, in coming to know (by whatever epistemic route) that something is that particular way one also comes to know (by that route) that it is good.

We now propose to strengthen our view by defending it against two recent rivals.

3 Tappolet's Reductive Account of Thick Concepts

Our non-reductive cognitivist treatment of what's expressed by thick ethical vocabulary claims that thick concepts are unanalysable. The first challenge to this view comes from Christine Tappolet. (Tappolet, 2004) Despite holding that the thin/thick relation is analogous to that of the determinable/determinate relation at the

level of properties and the specific/abstract relation at the level of descriptions, Tappolet also maintains that as we have an independent theoretical grasp on a theory of warranted attitude expression, we must be able to develop a two-component analysis of thick concepts that is still at the service of cognitivism.

Throughout her paper she distinguishes between a determinate/determinable model of the relation between the thick and the thin and a species/genus model. The hallmark of the species/genus model is that species must always have an independently specifiable differentia, as in “man is a rational animal” (contrast: “man is a human animal”), while on the determinable/determinate model this is not the case. Tappolet then argues that thick ethical concepts cannot stand to “good *pro tanto*” as species to genus because there is no such differentia. We agree. But she goes on to argue that the proper way to view thick ethical concepts is as conjunctions of what she calls an “affective concept” – defined in terms of reasons for certain feelings, such as ADMIRABLE and DESIRABLE - and some further condition, the non-evaluative grounds in virtue of which one judges whatever it is to be admirable, desirable, and so on. So her putative analyses run as follows:

BRAVERY \equiv ADMIRABILITY + withstanding danger (or whatever other descriptive candidate would be adequate),

KINDNESS \equiv ADMIRABILITY + giving to others, etc.

Thus there's some relation between thick concepts and the determinates of *pro tanto* goodness, as Tappolet conceives of them, but not a straightforward one.

Tappolet's account seems internally inconsistent. If there is a reason to reject the species/genus model of the thick concept/*pro tanto* goodness relation, then that very same reason will force rejection of Tappolet's account itself. For what non-evaluative conditions are we to conjoin with admirability to give necessary and

sufficient conditions for bravery, generosity and so on? At this point Tappolet faces a dilemma: if such conditions are stateable, they could presumably also be pressed into service as differentiae on the species/genus model, which would leave Tappolet's account unmotivated vis-à-vis its closest competitor. Or she could argue that though there are such conditions, they are not stateable by those competent with the concept. That is, following the example of a content-externalist treatment of natural kind terms, it might be argued that descriptive necessary and sufficient conditions for bravery are not part of the meaning of the term or of what it is to grasp the concept. So the fact that one could be competent in the use of the concept BRAVE and not be able to state any such conditions is compatible with the truth of the account. But if Tappolet can draw upon such an argument, so can the defender of the species/genus model. Furthermore, this is no longer an analysis: analyses have to be made up of other concepts and cannot simply list what falls under the original concept.

At this point we appeal to a more general interpretationism about the mental: our aim is to characterize a thinker's grasp of concepts such as CRUEL, or BRUTAL, or BRAVE. Our methodological principle is that while there is no general interpretationist scruple about recognising externally individuated contents, there must be some evidence, from within our ordinary grasp of the concepts that make up these contents, that we implicitly defer to a set of conditions beyond the concept user's grasp. (The alternative view simply rewards theft over honest toil.) It must be part of the concept user's self-understanding not simply that the concept would tolerate such an expansion, but also that it could be represented as an enrichment of what we took ourselves to be doing all along. Tappolet offers no grounds to substantiate this claim and is, in fact, commendably hesitant to endorse the externalist line:

One might wonder whether it is plausible to claim that thick concepts and natural

kind concepts are similarly out of the ordinary speaker's reach. The suggestion that the necessary descriptive underpinnings of a thick concept might be unknown to ordinary speakers is bound to appear implausible. (Tappolet, p. 215)

She goes on to canvas the suggestion that determining the underlying nature of (say) cruelty is a (contestable) task for the community of concept users, but that suggestion is unconvincing. If the external determination of the set of descriptive necessary and sufficient conditions fails because there are no moral experts, then an appeal to community is not going to be an adequate substitute for that role. Deferring to community is simply a part of non-expert, quotidian, conceptual competence and does no work in substantiating Putnam-style externalism. We are, after all, perfectly good at applying a thick ethical concept without any theory of its underlying "nature" that seeks to determine its extension. How can it be reasonable to set us the task of extension determination in this way, even distributed across the linguistic community?

We conclude that the unstateability by competent speakers of descriptive necessary and sufficient conditions for thick terms works – unless we have missed something important – as much against Tappolet's own view as it does against the species/genus model. We suggest, then, that for the reasons we have given she stick with the separate proposal to understand the thin/thick relation in terms of the relation of determinable to determinate at the level of properties and abstract/specific at the level of concepts and descriptions.

4 Elstein and Hurka's Sophisticated Two-Component Account

Daniel Elstein and Thomas Hurka have jointly proposed a different kind of two-component approach to thick ethical judgments which differs both from Tappolet's and from the simple two-component views we have referred to so far. (Elstein and

[Hurka, 2009](#)) They point out [that the](#) latter [views](#) all take a ‘descriptively determinate two-part form’ (517), where the non-evaluative (‘descriptive’) component on its own fully determines the extension of the thick concept. As such, all are indeed vulnerable [to a](#) well-known argument of John [McDowell](#)’s. But, they argue, it is a mistake to think that reductive analyses must be committed to ‘descriptive determinacy’: as long as the non-evaluative component determines the concept’s extension only partially, the analysis can escape the force of McDowell’s argument while nonetheless meeting the standards for a reduction.^{iv}

The argument of McDowell’s they have in mind is his ["disentangling" argument](#) (McDowell, 1981). McDowell invites us to consider the things to which it’s agreed that some thick ethical term applies, and then asks what all these things have non-evaluatively in common. The answer, he says, is ‘nothing’: that is what’s meant by the claim that the extension of a putative thick ethical concept is non-evaluatively ‘shapeless’. But because it is thus shapeless, there is no non-evaluative term mastery of which would have enabled its user to gather all and only the things to which it’s agreed the thick term applies, or would enable the user to go on to new cases in the same way as one who has mastered the thick term itself. So, McDowell reasons, competence with a thick term can’t be broken down into competence with a non-evaluative predicate which determines the thick term’s extension, plus a non-cognitive extra. (McDowell does not consider ‘Moorean’ reductionism, but for present purposes that is unimportant.)

In reply, [Elstein](#) and [Hurka](#) simply agree, but without abandoning reductionism about the thick, on the grounds that reductionists need not assume that the extension of a thick concept must be determined by a non-evaluative component all on its own. Thus on the simpler of their two types of analysis – they also advocate

a more complex pattern of analysis for some virtue terms, of which more shortly -

they propose

(To) replace(s) the fully determinate two part analysis ... with an only partly determinate one, which narrows the possible extension of the target concept to some degree, but does not fix it completely, so the remaining work is done by the evaluative part. (Elstein and Hurka, 2009, p. REF)

Thus, “x is selfish” is to be analyzed as

something like ‘x is wrong, and there are properties X, Y, and Z (not specified) that acts have in virtue of somehow bringing about the agent’s happiness rather than other people’s, such that x has X, Y, and Z, and X, Y, and Z make any acts that have them wrong.

‘Brings about the agent’s happiness rather than other people’s’ *partly* determines the extension of ‘selfish’, explaining for example why it would be a mistake to apply the term to distributions of goods between people. Thus a ‘strict impartialist’ (who thinks self-preference *per se* is never good) and a believer in ‘agent-centred prerogatives’ (who thinks it sometimes is) agree that satisfying ‘brings about... (etc.)’ is a necessary condition for satisfying ‘selfish’. But it doesn’t determine the extension of ‘selfish’ all on its own, because the moralists in dispute are plausibly to be pictured not as employing different concepts but as disagreeing over what falls under a given concept (E&H, 521-2). But because the non-evaluative component doesn’t on its own determine any extension for ‘selfish’ itself, and so *a fortiori* not a ‘shapeless’ one, the proposed analysis is immune to the ‘disentangling’ objection. For all that, however, the analysis is reductive because it

uses only thin evaluative terms (here, ‘wrong’) plus non-evaluative ones (here ‘brings about the agent’s happiness rather than other people’s’): any other concepts are represented in the analysis only by means of quantifiers.

Before we raise a problem for Elstein and Hurka’s two-part analysis, something needs to be said about the more complex three-part analysis they offer of ‘some virtue-concepts’. Here their idea is that thin evaluative concepts figure in the analysis of the relevant virtue-concepts twice over: once as on the two-part analysis (the thin evaluation of some unspecified further properties) and once more, independently, in the specification of those further properties themselves. Thus their analysis of ‘x is courageous’ is (roughly)

x is good, and x involves an agent’s accepting harm or the risk of harm for himself for the sake of goods greater than the evil of that harm, where this property makes any act that has it good. (Elstein & Hurka, 2009, p. 527)

As they say, the second ‘good’ here is the extra, ‘embedded’ evaluation. The three-part style of analysis is motivated by the quite proper observation that an act can only count as (say) courageous if it is done in pursuit of a good (or a good enough) end: drinking a pint of vodka in one go in order to show off is dangerous, but not courageous.

Two points about the three-part style of analysis deserve comment. The first is that whereas Elstein and Hurka’s two-part analyses studiously avoid specifying *any* property that fully determines the extension of the analysandum (so the analysis of ‘selfish’ tells us merely that there are *some* (unspecified) wrong-making properties that an act has in virtue of satisfying the non-evaluative necessary but not sufficient condition for selfishness), the three-part analyses *do* specify an extension-determining property for the thick term – that of ‘involving an agent’s accepting harm or the risk

of harm for himself for the sake of goods greater than the evil of that harm'. To be sure, the specification of that property embeds a thin evaluation, so the analysis does not violate the requirement that no extension-determining *non-evaluative* property should feature in it. But granted that, this property *does* determine the extension of 'courageous', whereas the counterpart specification in the two-part analyses merely 'narrows the extension of the target concept'. One might therefore have expected the evaluation-embedding property in the three-part analysis (i.e. 'involving an agent's accepting harm or the risk of harm' etc etc) also to do no more than that, leaving the final determination of the target concept's extension to some further set of properties of an act, unspecified in the analysis. But if there is a problem here, we suspect it can be fixed by a more complex formula, so we will not dwell on it.

The second point is that we do not see the three-part style of analysis as supplying any further reasons, beyond those putatively supplied by the two-part style, for the priority of the thin over the thick in ethical judgments. Let us suppose for argument's sake that the embedded thin judgment (about the goodness of the end) is verdictive or all-in. It is thus determined by all the evaluative features bearing on judgement, both thick and thin, and thus is no argument for the priority of the thin. Now suppose that the judgment is not verdictive: in fact there's just one thing that makes it the case that the agent is acting courageously, namely that their end (in the relevant circumstances, in the presence of the appropriate emotional dispositions etc.) is honesty. Of course it would proper to report this by saying that their end is good. But again, there is no evidence here for the priority of the thin: the data are compatible with the claim that 'good' here is an abstraction from the prior (evidential or *pro tanto*) thick judgment of honesty. Either way, what we have in judgments of courage here is not evidence for the priority of the thin, but one point of entry into an

inescapable holism of the thick. It is not clear in any case what marks off the cases that call for a three-part analysis from those where a two-part analysis is appropriate. ('Just', like 'courageous', is a virtue-term and lewdness – another of Elstein and Hurka's examples – is a vice.) So perhaps the three-part treatment applies only to virtues like courage and integrity where there's no particular end or type of end for the sake of which one needs to act in order to display the virtue (all the more reason, we would have thought, for some quantifiers in the three-part analyses). Consequently we focus henceforth on Elstein and Hurka's proposed two-part style of analysis, to which our main objection to their view pertains.

To develop this objection let us return to the trivial kind of 'analysis' that we hinted at earlier, BRUTAL = BAD + BRUTAL or, as it might be, RED = COLOURED + RED. Here is a more complex (but no less trivial) version of it: RED = (is true of surfaces) + COLOURED + (there is some further property of surfaces such that any surface satisfies COLOURED in virtue of having this property). This cumbersome formulation mimics Elstein and Hurka's two-part analyses, because its first clause simply 'narrows the possible extension of the target concept'. However, it differs from our earlier 'analysis' of RED in that it does not specify *which* property needs to be added to COLOURED to yield RED: it just says there is one. But it will be obvious that the only property that can be a value of the property-variable in our new quantified 'analysis' is, in fact, RED itself. The example thus dramatizes the fact that a would-be analysis is not improved just by quantifying in to the good- (or bad-) making property contexts if the only possible values of the good-making property variables are values that yield a merely trivial 'analysis'.

Now according to us, thin concepts are abstractions from (stand as determinables to) thick evaluations. So, according to us, the only value of the good-

making property variable in Elstein and Hurka's bipartite analysis that will make the analysis true is the thick evaluative property which is in fact the analysandum - as it might be, 'distributively just'. So cannot we just accept Elstein and Hurka's two-part style of analysis, on the grounds that its availability does not vindicate their reductionism over our anti-reductionism, as it is compatible with the only analysis available being the trivial one we have given?

That would in fact be hasty, for of course the unspecified good-making properties Elstein and Hurka have in mind are not just any old properties, but non-evaluative ones, while the good-making property which we acknowledge as satisfying the analysis is a thick evaluation itself. So the question is why we should accept that (once the range of application has been narrowed down) there is a further set of *non-evaluative* properties such that (whatever the example) *x* is wrong in virtue of having them, and anything that has them is thereby made wrong?

It might look as if our disagreement with Elstein and Hurka here concerns universalizability, and of course it is controversial whether if some properties make an item distributively just on one occasion, they make anything else that has them distributively just (e.g. the property 'having an equal number of sweets'). However, Elstein and Hurka don't argue for universalizability: they merely say, reasonably enough, that since the disentangling argument has been claimed to rule out reductionism even on the assumption of universalizability, if they can show, on the same assumption, that the argument fails to do this, that is a worthwhile result, notwithstanding the fact that some anti-reductionists reject the assumption (520-1). So if our objection to their two-part style of analysis depended only on rejecting universalizability, that would limit its scope and thereby its interest.

Now as it happens we do not think that if something makes an item (let's say)

distributively just, it must make everything that possesses it distributively just. But that is not where the dispute between ourselves and Elstein and Hurka lies. Universalizability says that if something (as it might be within a given range) is made good by its being F, then anything else within that range that is F is also thereby made good. But it doesn't say that anything within that range that is good is made good by being F. That is, universalizability is a constraint on what counts as a good-making property on an occasion: if that property would not make everything else that has it good, then it's not good-making here and now. But it does not commit one to the claim that the extension of a thick evaluative term has non-evaluative 'shape'.

The focal point of our dispute with Elstein and Hurka is, rather, their claim that full 'descriptive determinacy' of a thick term is 'provided by evaluation', given that what they mean by 'evaluation' seems to be the mapping of some small set of non-evaluative properties onto goodness. The novel feature of their analysis is that it eliminates extension-determining non-evaluative terms in favour of quantifiers, so nothing non-evaluative that's specified in the analysis is itself extension-determining for the thick term in question.

But now consider what Elstein and Hurka's quarrelling egalitarian and desert-theorist each thinks. What each thinks is that x is distributively just (on any occasion) *if and only if* it falls within some non-evaluatively specified range and it is F, G and H (for some non-evaluative F, G and H), the difference between them being that they disagree over *which* property at least one of F, G and H actually is. The striking thing here is not the commitment to universalizability (expressed by the 'on any occasion'), but rather the idea (expressed by 'if and only if') that any particular user of a thick evaluative term must take it that there *is* a fully determinate non-evaluative specification of the extension of the term, even though different users of the term may

differ over what the non-evaluative specification is. Quantification into the relevant non-evaluative property contexts lifts the mention of any particular such properties out of the analysis, so mastery of the non-evaluative concepts that (according to each disagreeing theorist) unify the term's extension aren't criterial for mastery of the thick term itself, thus allowing the egalitarian and the desert-theorist to agree in the concept they use.

But now: it seems only a slight improvement on the 'descriptive determinacy' view criticized by McDowell to say that someone who possesses the concept 'distributively just', or 'selfish', or whatever it might be, must be able to state in non-evaluative terms what (in his opinion) unifies its extension - even if doing that does not amount to saying what the word means. If the disentangling argument is right, there will not be anything non-evaluative that unifies it (the extension will be, as Elstein and Hurka seem to agree, 'shapeless'). So any opinions that users of a thick term may have about what non-evaluatively unifies its extension will be false. So – and to say this is simply to play a variation on the point we made by means of our complex but trivial 'analysis' of RED - there will be no substitutions for X, Y and Z in the two-part analysis which make the analysans come out true.

What if we drop the requirement that users of the term can state (or rather think they can state) what unifies the extension, and understand only that x is good, x falls within some non-evaluatively specified range, and there *are* some further non-evaluative properties that make x distributively just (etc. etc.), though they can't say what they are? There are two ways to understand this weakened requirement, neither of which helps.

First, if the extension of the thick term is non-evaluatively shapeless, there will not be any tidy set of non-evaluative properties (known or unknown) that unify

its extension, at least if the disentangling argument is valid. If on the other hand all that is meant is that the members of the extension of 'distributively just' *have* some non-evaluative properties, then we can hardly disagree: there must after all be something in virtue of which an item picks up a thick evaluation. But could that understanding explain competence with the thick concept to be analyzed? Of course not. (Compare our objection to the 'externalist' amendment to Tappolet's theory.) Indeed accepting this weakened formulation is consistent with maintaining that it is the evaluative interest that guides application of 'distributively just' from case to case, and that the concept itself is unanalyzable. The fact that no extension-determining non-evaluative properties are actually mentioned in Elstein and Hurka's sophisticated two-part proposal does not, in the end, advance the cause of reductionism very far. Their putative refutation of the "shapelessness" argument rests, fundamentally, on an assumption that grasp of a thick concept involves the belief that the extension of a thick term is, in fact, shaped by a determinate set of non-evaluative properties. Why would a non-reductionist accept *that*?

Some independent argument is needed to the effect that the unanalysability of thick concepts is problematic in its own terms. It would, indeed, be unfortunate if the thesis of unanalysability was constrained to a mute conservatism of the local folkways, but we do not think that it is. Is our thesis of unanalysability all that can be said about thick concepts? Can we say nothing about how such concepts are learned, how they can be essentially contested, or must we simply say "this is how we go on" with such terms? This raises a whole host of issues that, for reasons of space, we cannot go into here. However, we do believe that given the connection between meaning and explanations of meaning, there is a great deal to be said about the use of thick concepts that does not constitute an analysis of them. Any competent use of a

concept grasps various truisms about that concept, derived from rules of thumb used to teach the concepts, or specifications of its Aristotelian “field of application”. (For the latter, see Dancy, 1996) [Gibbard](#) points out that a great deal of language [learning involves a lot of catching on in the absence of explicit “definitions”](#). But none of this folk wisdom constitutes an *analysis* of a thick concept and our practices with such concepts proceeds without such analyses.

We also deny that Elstein and Hurka’s analysis is needed in order to make sense of moral disagreement. [We would naturally explain it using the distinction between concept and conception](#), introduced by Herbert Hart and [made famous by Rawls](#) in the course of his own discussion of exactly the same example, differing conceptions of the concept of justice. (Hart, 1961/1994, pp. 155-159; Rawls, 1971/1999, p. 5; Ezcurdia, 1998) In drawing the concept versus conception distinction one must be careful, as Ezcurdia notes, to distinguish “having a conception of a concept from having a conception *associated with* the concept which one takes to be *analytic to or constitutive of* that concept”. (Ezcurdia, 1998, p. 187) Conceptions are further beliefs about the proper domain of application of a concept that serve to vindicate the public character of meaning and the normativity of concept use. The fact that two people can disagree about justice while sharing the same concept is easily explained by appeal to the concept/conception distinction and we can appeal to that distinction even in the case of unanalysable concepts.

Elstein and Hurka’s paper appeals to a further putative advantage that their view possesses over other views, namely, that it allows one to diagnose a kind protracted disagreement that they take ethical disagreements to exemplify. Our response is that W. B. Gallie has identified certain kinds of concept as essentially contest in the sense that [disagreements over evaluative concepts can be peculiarly](#)

intractable because they are evaluative “all the way down”. (Gallie, 1956) Once again, our view that the thin stands in no privileged relationship to the verdictive plays a role here. The grounds of concept application can, on our view, be characterized using either thick or thin ethical concepts. There is nothing in our use of ethical concepts to sustain the rationalist hope that greater abstraction correlates with greater suasive power, and with a special role for thin vocabulary. (Yet it is those assumptions that shape Elstein and Hurka’s discussion.) It is not always possible to identify a locus for a disagreement in non-evaluative terms; we do not believe that this is always the case. But the main point is that you can believe that thick concepts are unanalysable but still explain reasonable disagreement between competent users of the concept by appealing to the concept versus conception distinction.

Our disagreement with Elstein and Hurka is a radical one: we claim that a set of concepts is in principle unanalysable while they claim actually to have provided the correct analysis. Are we embarrassed by the fact that an analysis of thick concepts seems available to such competent users of ethical concepts as Elstein and Hurka?

Does any plausibility accrue to their view simply because their own linguistic intuitions suggest to them the correctness of their analysis? No, because the situation is directly analogous to other areas of philosophy, notably epistemology. Those who deny that knowledge can be analysed are very well aware of the large number of putative analyses of knowledge that epistemologists have been putting forward over a considerable period of time.

Nor do we believe that we have to enter into the game of putative-analysis-meets-counter-example. We have explained why the data cited by Elstein and Hurka is fully explained by our own view. Once again, we would draw a direct analogy with the putative definition of knowledge as justified true belief plus contested further

factor x (where each reductively minded epistemologist solves this equation for x in his or her own proprietary way). Proponents of unanalysability are under no obligation to enter the over-familiar dialectic where putative analysis meets putative counter-example, resulting in an *ad hoc* clause added to the analysis, which meets with a further counter-example – and so on. Our aim throughout this paper has been well-captured by P F Strawson when he contrasted conceptual analysis with conceptual elucidation. The latter seeks the identification of "general, pervasive, and ultimately irreducible" components in "a structure which constitutes the framework of our ordinary thought and talk and which is presupposed by (our theories)" (Strawson, 1998, p. 24). Ethical judgements that use thick ethical concepts seem to us to be an irreducible component of ordinary thought and talk of precisely this kind.

5 Thick Concepts and Presupposition

One further objection to our non-reductionist cognitivism about thick concepts, due to Pekka Väyrynen (Väyrynen, 2009), deserves extended consideration. Suppose two speakers take different evaluative attitudes to items to which they agree in applying a given thick term. If the difference in attitude induces a difference in what the speakers express, such a theory implies a mere illusion of communication. Väyrynen argues that anyone who agrees, as Blackburn has persuasively argued, that in reality there is communication and indeed disagreement here (Blackburn, 1992) must also concede that "thick terms (do not) have evaluative content as part of their conventional meaning" (Väyrynen, op. cit, p. 442), that is, that thick terms do not express thick concepts in our sense.^v

Let us call the agreed fact that 'the evaluative contents associated with (thick expressions) enjoy a certain autonomy with respect to the attitudes and intentions of

particular thinkers' (Väyrynen, op. cit, p. 444) the *slack* between attitudes and evaluative contents. The most oft-cited example of slack in this sense is Oscar Wilde: if the evaluative content of "blasphemous" had depended solely on Wilde's own attitudes, he need not have worried that by using that word he would have expressed condemnation of the thing, so he would not have needed to say that 'blasphemous' was 'not a word' of his. There's also another kind of case, which is very common, where – unlike Wilde - we use an expression in line with the usage of someone fully signed up to the evaluations associated with it, but while our own attitudes are still work in progress.

Väyrynen seems to be interested in cases of a third kind, where speakers have an expertise in the use of an expression comparable to that of someone fully signed up to the associated evaluations, but reject those valuations wholesale – as Väyrynen puts it, they find the concept "objectionable". We say that Väyrynen's 'objectors' must have an expertise comparable to that of someone fully signed up to the associated evaluations because he stresses that "non-evaluative constraints" on the application of a thick expression are accepted both by the objectors and by those who are fully signed up to the evaluations associated with the expression (the 'non-objectors'). (Väyrynen, 2009, p. 442) In Väyrynen's example the concept under dispute is CHASTE. Now in some sense or other, Väyrynen has to be right about non-evaluative constraints: no one could count as understanding the term "chaste" unless they understood that chastity has got something to do with sex, for example. But in Väyrynen's view, the constraints are much more specific than this. Although he does not claim that the relevant non-evaluative predicate "is anything like extensionally equivalent to 'chaste' or sufficient for it to apply, or even that it captures the full non-evaluative meaning of 'chaste'" (p.443), he does claim that all will agree that "if ...

conduct is to count as satisfying ‘chaste’, it must have some properties which signal some kind of dedication to not being sexually provocative”. (Väyrynen , 2009, p.442).

So it seems as if these properties will be possessed – and agreed to be so by objectors and non-objectors alike - by everything in the extension of “chaste”. To see how this is relevant to the level of linguistic expertise Väyrynen expects of his objectors, consider the fact that devotees of the concept CHASTE usually – as we understand it - acknowledge the possibility of chaste behaviour in marriage (and we do not mean *mariages blancs*). Now we are not quite sure what “dedication to not being sexually provocative” actually means (nor whether “provocative” is really non-evaluative). But let us say it means making an effort not to arouse another sexually. Assuming married persons may attempt to arouse one another sexually without offending against chastity, the case of chastity in marriage seems like a counterexample to Väyrynen’s non-evaluative constraint because it is a necessary (though not a sufficient) condition. But to figure out whether it is a counterexample or not, one has to be rather expert – more expert than we are – at applying the term “chaste”. And if the non-evaluative condition must be agreed to apply by non-objectors and objectors to “chaste” alike, the objectors must have the same considerable level of expertise with the term as the non-objectors.

This combination – wholesale rejection of the concept plus unimpeachable expertise in how to apply it - must surely be something of a rarity. Wilde, after all, rejected any evaluations the word “blasphemous” might be used to express, but refused even to use the word. If one really thinks there is nothing worthwhile to be said by means of a given thick expression, why retain the same interest in the niceties of its usage as would be shown by someone fully signed up to the associated evaluations? Indeed many speakers who are far from any such wholesale rejection –

for instance, because they think the concept is an attractive one but they want to find out more – cannot match the expertise in use of those speakers who are fully signed up to it. This is not to say that we think the combination is impossible: it might be found in someone, for example, who was fully identified with the evaluations associated with a concept and had for some reason become suddenly and totally disillusioned. And even if it were impossible, that fact wouldn't magic away all the difficulties of occupying our position, since some of those stem from the phenomenon of slack generally, not from Väyrynen's very specific variant of it. However, the difficulties Väyrynen himself raises do arise from the specific variant, and not from the general phenomenon. We confine ourselves to noting this point for now; it will be important later on.

We turn first, however, to Väyrynen's main argument for the claim that the evaluative contents associated with a given thick expression are not part of the concept it expresses. His argument depends on the behaviour of "objectionable thick concepts" - i.e. thick expressions where the speaker retains full competence with the expression but objects wholesale to the attitude it expresses - in denials. The focus of Väyrynen's discussion is the sentence

(5) Abstinence from extramarital sex is chaste.

He claims that a wholesale objector to the concept "chaste" will disagree with (5), but will not 'typically ... be willing to express their disagreement with (5)' by uttering:

(6) Abstinence from extramarital sex is not chaste.

Formatted: Bullets and

Formatted: Bullets and

On the contrary, someone who asserts (6) will typically be a non-objector, who in addition thinks that abstinence from extramarital sex is not (for example) enough for chastity. (Väyrynen, 2009, p. 448) Again, argues Väyrynen, “if we regard ‘chaste’ as an objectionable concept, our objection to (5) isn’t (merely) that it’s false”, for “if (5) were false, then (6) should be non-problematically true” (though perhaps pragmatically odd without further explanation). But if you believe (6) is non-problematically true, then why be unwilling to say it?

Now there’s a danger of misunderstanding Väyrynen at this point. If objectors to the concept CHASTE really are unwilling to register their objection by uttering (6), then there should be no account to give of what they mean when they utter it. But that is exactly the thing Väyrynen seeks to give an account of. In fact, however, it is not that objectors do not utter (6): it’s rather that they don’t utter (6) on its own, but accompanied by some follow-up. The difference between objectors and non-objectors to chaste is that the former must “hesitate to express disagreement with (5) by uttering an expansion of (6) that uses negation truth-conditionally”. (Väyrynen, 2010, p.448) Why? Because truth-conditional negation “takes scope over the assertion that sentence S would have made if uttered, but not over other types of information its utterance would have conveyed (sc. non-semantically)” (ibid.) So “assertions of ‘A is chaste’ and of ‘A is unchaste’ both convey the evaluative content which is associated with ‘chaste’” and which is exactly what objectors want not to convey. Hence, on the lips of an objector, the negation in (6) and expansions of it need to be interpreted not as “ordinary truth-conditional” negation but metalinguistically. As an analogy, compare negation directed towards someone’s pronunciation: “He didn’t call the POlice, he called the poLICE”, (Väyrynen, 2009, p.449)

In reply, let us first of all explore the idea that objectors to CHASTE should after all be willing to express their objection by uttering (6), with negation understood truth-conditionally – that is, the idea that objectors regard (5) as straightforwardly false. Many virtue and vice terms have naturally associated contrasting terms that are not their contradictories, like “just” and “unjust” or “brave” and “cowardly”; the same goes for “chaste” and “unchaste”. As a result, (5) has two ways of being false: either because abstinence from extramarital sex is unchaste, or because – presumably like having books overdue at the library – it is neither chaste nor unchaste. Now if the non-objector – at least, the non-objector with the sort of fanatical first-order views Väyrynen has in mind – utters (6), it seems plausible that he would also be prepared to utter:

(5*) Abstinence from extramarital sex is unchaste.

Certainly the objector would have no truck with that. But mightn't the objector nonetheless be willing to assert (6), because he thinks that sexual abstinence is neither chaste nor unchaste? Suppose he thinks that sexual conduct is a free-for-all, or at least, a free-for-all as long as other requirements are met (e.g. absence of harm, or consent – in any case, no requirement it takes a distinctive concept of chastity to articulate). With an attitude like that, the objector would certainly satisfy the requirement, which Väyrynen rightly insists on, that there be genuine disagreement between the objector and the non-objector. What is more, the attitude would explain the objector's holding that no sexual conduct falls into the class either of the chaste or the unchaste: as far as chastity goes, all sexual conduct is like having books overdue at the library. This possibility opens the question whether objectionable thick concepts are to be thought of as empty, which we will consider in a moment.

But barring worries from that quarter, if the objector does assert (6) with that sort of attitude in mind, then there seems to be no difference in meaning between the kind of negation involved in his utterance of it, and the kind of negation involved in the non-objector's utterance of the same sentence. If the test for the detachability of attitude from the concept expressed by thick expressions is that, in denials involving objectionable thick concepts, negation has to be interpreted as metalinguistic – because denial “is directed specifically at the associated evaluative content” (Väyrynen , 2010, p. 446) - on the present proposal denials involving objectionable thick concepts fail the test.

We now need to confront the question that we have deferred up to now, namely whether objectionable thick concepts are empty, and this brings us also to the question whether the simple solution we have proposed so far to Väyrynen's challenge – namely, that the objector (of the niche variety Väyrynen has in mind) is disposed to utter the truth-conditional negation of (5) because he thinks precisely that nothing is either chaste or unchaste – is really a solution we can endorse ourselves, or just a dialectical move against Väyrynen.

First, we need to clarify some issues of terminology. Väyrynen says “what is wrong with wielding such concepts as CHASTE or LUSTFUL isn't merely the sort of fault that is involved in wielding such empty concepts as *phlogiston* or *Bigfoot*”. (Väyrynen , 2009, p.452). He does not say what that sort of fault is, but he is right that these are not helpful analogies. The latter two expressions are putative singular terms, whereas “chaste” is ostensibly a first-level predicate, and emptiness for predicates cannot mean the non-existence of the object to which they purportedly refer, because they do not purport to refer to an object. So what does it mean? “Empty” can – with some artificiality – be used in two ways. An empty term can be one that has nothing

in its extension and, thereby, everything in its anti-extension, for example “unicorn”.
Terms that are empty in this sense nonetheless succeed in dividing (as it were) the
universal domain in two, just as “tiger” does. But a term can be empty in a second
sense, that of failing to determine an extension (or therefore an anti-extension). Race
terms may be of this kind – “Caucasian”, “black”, or whatever the list is supposed to
include: despite seeming to speakers for a good long time as if each such term
determines a principle for dividing the universal domain, it may turn out that none
does (because the terms’ use is in some way incoherent).

Now in what sense, if any, should non-reductionist cognitivists say
objectionable thick concepts are empty? The view we have been canvassing so far is
that “chaste” is empty in the first sense: since nothing is either chaste or unchaste, *a*
fortiori nothing is chaste. So the concept CHASTE has nothing in its extension and
everything in its anti-extension (since the not chaste includes the unchaste). This line
certainly offers some straightforward solutions – (5) is supplied with a truth-value, for
example (the value false). But there is something inappropriate about a wholesale
objector to the concept CHASTE saying *in propria voce* that nothing is either chaste or
unchaste. For we take it that what concerns wholesale objectors to the concept is the
very idea that these terms serve to classify, not that they classify all right but place
things on the wrong side of a boundary. People who deny that there are unicorns do
not reject the concept UNICORN; on the contrary, they need it to articulate the thing
they hold true (that there are not any things of that kind). Similarly, it is compatible
with the free-for-all first-order sexual attitude we envisaged to think that CHASTE is
a perfectly good concept – we can see exactly where it divides the chaste, the
unchaste and the neither-chaste-nor-unchaste. But whereas most people who deploy it
are (according to the free-for-all view) prudes or “repressed” or the like, the free-for-

all user_of it makes all sexual behaviour come out in the category of the neither-chaste-nor-unchaste. Our original “objector” seems after all not to be an objector to the concept “chaste”, but a non-objector with (compared to other non-objectors) non-standard sexual attitudes.

If that really is the kind of “objector” Väyrynen had in mind, of course, then there is nothing wrong with the foregoing as a reply to him. However, it must be at least possible for there to be an objector of a more radical kind, who needs to be characterized as someone who holds that CHASTE is empty in the second sense. Väyrynen sees problems here for our view. If the concept is empty, he says, we must say (5) is neither true nor false. But “lack of truth-value is most typically attributed to expressions with false presuppositions” (452), and that diagnosis sits ill with our contention that the evaluative content associated with a thick expression is a *semantic* aspect of what is said by utterances of it.

However, it is not clear that the link between lack of truth-value and presupposition-failure is as strong as Väyrynen thinks. Certainly in the cases of the singular terms he mentions (“Bigfoot” etc.), it is arguable that utterances involving the expression, far from being apt to express a thought but failing to do so because of a presupposition failure (as on non-Russellian treatments of ‘the King of France is bald’), aren’t even apt to express a thought – which one would it be? When speakers attempt to formulate thoughts involving these terms, there is at best an illusion of a thought entertained. The race example suggests a parallel account might be made out for predicate expressions – and also reasons why we might be chronically unable to ascertain whether or not we are succeeding in expressing senses by means of these terms at all. That gives us an account of lack of truth-value without appealing to presupposition.

Moreover this explanation seems accurately to locate not only what the wholesale objector thinks – that ‘chaste’ does not express a concept – but also what less than wholesale objectors think, namely: here is a term, I can use it as well as anyone else, but am I really classifying anything? It also commits us to an account of what objectors – if their objection is indeed well-founded – are doing with the objectionable terms. It is not that they are classifying in a way that is at variance with mainstream evaluations. All they are doing is applying the term as it is conventionally applied, but not in so doing grasping a sense and withholding an evaluation, because there isn’t a sense to grasp. (That, by the way, is a very good reason why we might expect most wholesale objectors to be like Wilde rather than like Väyrynen’s character: they refuse to use the word, because they believe there are no thoughts it is apt to express.)

This, ironically enough, chimes with an important aspect of Väyrynen’s account. For if no sense is expressed by “chaste” (or “Caucasian”, etc.), and we believe this, then the correct account of what we say in order to register our objection to the concept, i.e. (6), cannot be that we are straightforwardly asserting the negation of the thought that a non-objector took himself to be asserting by uttering (5). That is because if there is no such thought, there is no such thing as its negation. It looks as if the correct account of (6) must after all be metalinguistic (compare “that thing in the sky is not Vulcan”). Our difference from Väyrynen is just that we see no tension between the claim that negation is metalinguistic and the claim that evaluative content is conveyed semantically, not pragmatically. Whether the objector is to be thought of as holding that everything is neither chaste nor unchaste (so that (5) is straightforwardly false and (6) true), or as holding that sentences featuring the term “chaste” fail to express a sense (so (5) lacks a truth-value and (6) – if, unWildeanly,

they utter it at all - is to be construed as metalinguistic negation), we do not see that Väyrynen has presented compelling reasons against the claim that thick terms express thick concepts as we understand them.

Conclusion

In this paper we have argued that two- (or multi-) component reductive analyses of thick concepts are not required by the surface data, because there is another explanation of the data that suffices to explain it, namely, our own explanation. Reductionists need some independent motivation for their position and all the candidate reductionist proposals currently available are internally incoherent in one way or another. We have tried to go a little deeper into the motivations for reductionisms in this area. They are motivated, at least in part, by the thought that some such view has to be correct simply because of the difficulties facing the non-reductionist alternative. That is why we have related our own non-reductionist account to some of the standard objections to show that it can meet them.

List of Works Cited

Blackburn, Simon (1992) *Spreading the Word: Groundings in the Philosophy of Language*, Oxford University Press.

Bridgman, Percy W. (1927) *The Logic of Modern Physics*, Macmillan.

Chihara, Charles and Fodor, Jerry (1965) 'Operationalism and Ordinary Language', *American Philosophical Quarterly*, vol. 2, no. 4, pp. 281-295.

Dancy, Jonathan (1996) 'In Defence of Thick Concepts', *Midwest Studies in Philosophy*.

Formatted: Font: Italic

Formatted: Font: Italic

Dancy, Jonathan (2004) *Ethics Without Principles*, Oxford University Press.

[Elstein, Daniel and Hurka, Thomas \(2009\) 'From Thick to Thin: Two Moral Reduction Plans'](#), *Canadian Journal of Philosophy*, vol. 39, no. 4.

Ezcurdia, Maite (1998) 'The Concept-Conception Distinction', *Philosophical Issues* vol. 9, pp. 187-192.

Foot, Philippa (2003) *Virtues and Vices: and Other Essays in Moral Philosophy*, Oxford University Press.

[Funkhouser, Eric \(2006\) 'The Determinable-Determinate Relation'](#), *Nous*, vol. 40 no. 3, pp. 548–569.

Formatted: Font: Italic

Gallie, W. B. (1956) 'Essentially Contested Concepts', *Proceedings of the Aristotelian Society*, vol. 56, pp. 167-198.

[Harcourt, Edward \(2005\) 'Quasi-realism and Ethical Appearances'](#), *Mind*, vol. 114, no. 454, pp 249-275.

Formatted: Font: Italic

[Hare, R. M. \(1952\) *The Language of Morals*, Clarendon Press.](#)

Formatted: Font: Italic

Hart, H. L. A. (1961/1994) *The Concept of Law*, Clarendon Press.

[Johnson, W. E. \(1921\) *Logic*, vol I, Cambridge University Press.](#)

Formatted: Font: Italic

[Kirchin, Simon \(2004\) 'Reasons and Reduction', presented to the conference 'The Space of Reasons', Cape Town, July.](#)

[Kirchin, Simon \(2010\) 'The Shapelessness Hypothesis'](#), *Philosophers' Imprint*, vol. 10, no. 4, February.

Formatted: Font: Italic

Lovibond, Sabina (2000) 'Ethical Upbringing: From Connivance to Cognition' in *Identity, Truth and Value: Essays for David Wiggins*, eds Lovibond S. and Williams, S. G., Blackwell Publishers, pp. 76-94.

[Lovibond, Sabina \(2002\) *Ethical Formation*, Harvard University Press.](#)

[McDowell, John \(1982\) 'Criteria, Defeasibility and Knowledge', *Proceedings of the British Academy*, London, Volume LXVIII.](#)

Formatted: Font: Italic

[Moore, A. W. \(2005a\) 'Maxims and Thick Ethical Concepts', paper presented to the Eastern Division of the APA.](#)

[Moore, A. W. \(2005b\) *Noble in Reason, Infinite in Faculty: Themes and Variations in Kant's Moral and Religious Philosophy*, Routledge.](#)

Formatted: Font: Italic

Payne, Andrew (2006) 'A New Account of Thick Concepts', *The Journal of Value Inquiry*, vol. 39. no. 1, pp. 89-103.

[Prior, Arthur \(1949\) 'I. Determinables, Determinates and Determinants', *Mind*, vol. LVII, pp. 1-20.](#)

Formatted: Font: Italic

Rawls, John (1971/1999) *A Theory of Justice*, Belknap Press.

[Roberts, Debbie \(2009\) 'Shapelessness and the Thick', paper presented to B.S.E.T., 13-15 July.](#)

Scheffler, Samuel (1987) 'Morality Through Thick and Thin: A Critical Notice of Ethics and the Limits of Philosophy', *The Philosophical Review*, vol. 96, no. 3, pp. 411-434

Shope, Robert K. (1983) *The Analysis of Knowing: a Decade of Research*, Princeton University Press.

Stratton-Lake, Philip (2000) *Kant, Duty and Moral Worth*, Routledge.

Strawson P. F. (1998) *Analysis and Metaphysics*, Oxford University Press.

[Tappolet, Christine \(2004\) 'Through Thick and Thin: Good and its Determinables', *Dialectica*, vol. 58, pp. 207-220.](#)

Formatted: Font: Italic

[Thomas, Alan \(2006\) *Value and Context: the Nature of Moral and Political Knowledge*, Clarendon Press.](#)

Formatted: Font: Italic

Thomson, J. J. (1996) 'Moral Objectivity' in G. H. Harman and J. J. Thomson, *Moral Objectivity*, Blackwells.

Formatted: Font: Italic

Väyrynen, Pekka (2009) 'Objectionable Thick Concepts in Denials', *Philosophical Perspectives*, 23, pp. 439-469.

Wiggins, David (2006) *Ethics: Twelve Lectures on the Philosophy of Morality*, Penguin Books.

Formatted: Font: Italic

Williams, Bernard (1983) *Ethics and the Limits of Philosophy*, Fontana.

Williamson, Timothy (2000) *Knowledge and Its Limits*, Oxford University Press.

Wittgenstein, Ludwig (1977) *Remarks on Colour*, Blackwell.

ⁱ Cf. David Wiggins (2006), p. 378: "the commendatory or approbatory force of (ethical) terms (where that force is present and in operation) (is) not an extra input into their full meaning but ... results(s) from what they already mean".

ⁱⁱ The point is formulated with admirable clarity by Blackburn in his (1992). For further recognition of the difference between these two strategies, one eliminative, and one reductive, see Väyrynen (2009).

Formatted: Font: Times New Roman, Italic

ⁱⁱⁱ We follow the convention of referring to concepts in small capitals.

^{iv} In order fairly to represent the Elstein and Hurka view we will engage with it in this form. But we do not think that it is, in fact, true that anyone in the existing debate thought that the two components of a putative analysis were merely conjoined. Blackburn and Gibbard, for example, take the attitudinal component to be attitudes directed to the extension falling under the associated description. That is not a mere conjunction. It is true that Blackburn and Elstein and Hurka disagree about whether this extension can be captured in non-evaluative terms. But the difference between them is one of degree, not of kind.

Formatted: Font: Times New Roman, Italic

^v There is of course a larger issue not far off here: if we are to argue that difference in attitude does not induce difference in concept expressed, we also need to show – on pain of conceding externalism in the theory of motivation – how a deployment of the concept in question can on its own bear the right sort of connection with action, given that – by hypothesis – sometimes it does not. This is the large issue that we said in the introduction that we would put off till another occasion. Väyrynen's attack on non-reductionist cognitivism about the thick is independent, however, of any issues in the theory of motivation: independent, that is, of any further debts we stand to incur in defending that view.

Formatted: Font: Times New Roman, Italic