THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Shapelessness and the Thick

Citation for published version:

Roberts, D 2011, 'Shapelessness and the Thick', *Ethics: An International Journal of Social, Political, and Legal Philosophy*, vol. 121, no. 3, pp. 489-520. <https://doi.org/10.1086/658898>

Digital Object Identifier (DOI):

[10.1086/658898](https://doi.org/10.1086/658898)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Ethics: An International Journal of Social, Political, and Legal Philosophy

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Shapelessness and the Thick

Author(s): by Debbie Roberts

Reviewed work(s):

Source: *Ethics*, Vol. 121, No. 3 (April 2011), pp. 489-520

Published by: [The University of Chicago Press](#)

Stable URL: <http://www.jstor.org/stable/10.1086/658898>

Accessed: 25/10/2012 08:40

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at
<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

The University of Chicago Press is collaborating with JSTOR to digitize, preserve and extend access to *Ethics*.

Shapelessness and the Thick*

Debbie Roberts

This article aims to clarify the view that thick concepts are irreducibly thick. I do this by putting the disentangling argument in its place and then setting out what nonreductivists about the thick are committed to. To distinguish the view from possible reductive accounts, defenders of irreducible thickness are, I argue, committed to the claim that evaluative concepts and properties are nonevaluatively shapeless. This in turn requires a commitment to (radical) holism and particularism. Nonreductivists are also committed to the claim that a thick concept is in itself evaluative, and not evaluative because of any link to thin evaluation.

I. INTRODUCTION

There has been a recent surge of interest in thick concepts.¹ My aim in this article is to clarify a prominent position in the debate: the view that thick concepts are irreducibly thick. This is usually further characterized metaphorically and negatively as the view that there can be no disentangling of the deeply entangled evaluative and nonevaluative elements that make up thick conceptual content. My strategy is to

* A previous version of this article was presented at the 2009 BSET conference at the University of Reading, and I am grateful to the participants for helpful discussion. For helpful comments on previous drafts, I am grateful to Max de Gaynesford, Matti Eklund, Daniel Elstein, Brad Hooker, James Lenman, David McNaughton, Derek Parfit, and Pekka Väyrynen, and to two referees and seven editors for this journal. I am also grateful to Jamie Dreier for editorial guidance. I am especially grateful to Jonathan Dancy and Guy Fletcher.

1. See, e.g., Daniel Y. Elstein and Thomas Hurka, "From Thick to Thin: Two Moral Reduction Plans," *Canadian Journal of Philosophy* 39 (2009): 515–35; Pekka Väyrynen, "Objectionable Thick Concepts in Denials," *Philosophical Perspectives* 23 (2009): 439–69; Simon Kirchin, "The Shapelessness Hypothesis," *Philosophers' Imprint* 10 (2010), <http://www.philosophersimprint.org/010004>; Matti Eklund, "What Are Thick Concepts?" *Canadian Journal of Philosophy* (forthcoming).

Ethics 121 (April 2011): 489–520

© 2011 by The University of Chicago. All rights reserved. 0014-1704/2011/12103-0001\$10.00

unpack this metaphor and thereby give a positive statement of the view.

In the literature, the claim that thick concepts are irreducibly thick is held to be important because of what is supposed to follow from it. It is irreducible thickness that, some think, wreaks havoc with the fact-value distinction. It is also supposed to create problems for noncognitivism, problems for those who think that all moral judgments can be expressed using a few thin moral concepts such as GOOD, BAD, RIGHT, and WRONG, and even problems for the idea that the evaluative in general supervenes on the nonevaluative. But exactly what irreducible thickness is remains elusive. Many admit to finding the notion somewhat dark.²

This is an unfortunate state of affairs. Understanding the nature of the thick is important for understanding the nature of the evaluative: for answering such questions as what it is for a concept to be evaluative, what the evaluative should be contrasted with, what it is to be competent with an evaluative concept, and what we should take the evaluative to supervene and depend on. If thick concepts are irreducibly thick, this has significance wherever there is thick evaluation.³ Moreover, if thick concepts are irreducibly thick, we cannot assume that their relation to certain thin concepts is analytic or conceptual. This raises potentially unsettling questions for all those normative theories formulated using only thin moral terms. In addition, the accounts we can give of particular ethical concepts like FRIENDSHIP, JUSTICE, COURAGE, and IMPARTIALITY will be dramatically affected if these concepts are irreducibly thick.⁴ Given these implications, and the prominence of the

2. These include people currently working on thick concepts. See, e.g., Eklund, "What Are Thick Concepts?" Elijah Millgram makes a related point in "Thick Ethical Concepts and the Fact-Value Distinction," published as "Inhaltsreiche ethische Begriffe und die Unterscheidung zwischen Tatsachen und Werten," in *Zum moralischen Denken*, ed. C. Fehige and G. Meggle (Frankfurt: Suhrkamp, 1995), 354–88. Millgram argues that the "Thick Ethical Concept Argument" against the fact-value distinction, while frequently invoked, has never actually been made. He does not go on to give an account of the irreducibly thick, however. Jonathan Dancy gives the most comprehensive account of the irreducibly thick in "In Defense of Thick Concepts," in *Midwest Studies in Philosophy XX: Moral Concepts*, ed. P. A. French et al. (Notre Dame, IN: University of Notre Dame Press, 1995), 263–79, but this article is both extremely dense and insufficiently discussed.

3. Beyond metaethics, thick concepts are discussed in aesthetics. Increasingly, they are also discussed in epistemology; see, e.g., Jeremy Wanderer and Ben Kotzee, eds., "Thick and Thin Concepts in Epistemology," special issue, *Philosophical Papers* 37 (2008).

4. Proper discussion of this point is beyond the scope of this article, but it is worth pointing out that the nature of thick concepts also affects whether normative ethical theories should aim to be universal, or whether they should be rooted in particular forms of life. The issue thus has implications beyond normative ethics, for example, for debates about communitarianism in political philosophy.

view in the debate about the thick, it is worth getting clear on just what irreducible thickness is.

The discussion proceeds as follows. Sections II and III set the scene with some preliminary remarks and an initial outline of the distinction between thick and thin. In Section IV, I discuss the disentangling argument (DA).⁵ The disentangling argument, often attributed to John McDowell, is taken to be the main argument in the literature against reductive views of thick concepts—views that hold that the evaluative and nonevaluative elements that make up thick conceptual content can be disentangled by analysis. Those who think that the disentangling argument succeeds take it to be a reason for holding that thick concepts are irreducibly thick.⁶ If this argument does succeed, we might expect it to provide insight into what irreducible thickness is. I argue, though, that the disentangling argument does not succeed as a general argument against reductive views of the thick and thus that it is not a reason to hold that thick concepts are irreducibly thick. More importantly for my purposes, the disentangling argument does not serve to clarify what irreducible thickness is.

However, help is at hand if we return to the original remarks that gave rise to the disentangling argument. Section V shows, contrary to what appears to have been assumed in the literature on the thick, that McDowell's antidisentangling remarks in his "Non-cognitivism and Rule-Following" do not especially concern the thick; their primary point is not about the analyzability of thick concepts, but it is about the nature of evaluation in general. The fact that these antidisentangling remarks are about the nature of evaluation in general, however, does not mean that they cannot be useful in explicating the notion of the irreducibly thick. Section VI shows how this is so, and develops in detail just what a nonreductivist about the thick is committed to. It turns out, I argue, that commitment to irreducible thickness involves commitment to a view of the nature of evaluation in general. In Section VII, I explore some of the implications of this account, and I conclude by considering the implications for the fact-value distinction.

5. "Disentangling argument" is perhaps an unfortunate label, as it is supposed to be an argument against disentangling, rather than for it. I use this label, however, as it is established in the literature.

6. See, e.g., Bernard Williams, *Ethics and the Limits of Philosophy* (Cambridge, MA: Harvard University Press, 1985); Dancy, "In Defense of Thick Concepts"; Hilary Putnam, "The Entanglement of Fact and Value," in *The Collapse of the Fact/Value Dichotomy and Other Essays* (Cambridge, MA: Harvard University Press, 2002), 28–45; and Charles Taylor, "Ethics and Ontology," *Journal of Philosophy* 100 (2003): 302–20.

II. PRELIMINARY REMARKS

In order to clear the way for my discussion, it will be helpful to indicate how certain complexities will be negotiated.

A. The Relation of the Debate about the Nature of Thick Concepts to the Debate between Cognitivism and Noncognitivism

In the early literature on thick concepts, the disentangling argument is primarily deployed within the context of the broader debate between cognitivism and noncognitivism. There are two points to note here. First, while it is important to keep this context in mind in order to appreciate the primary point of McDowell's antidisentangling remarks, the primary point of this current article is not to offer an argument against noncognitivism. The primary point of this article is to clarify the view that thick concepts are irreducibly thick, and to indicate what defenders of this view are committed to.

Second, it is not obvious that the view that thick concepts are irreducibly thick must be a cognitivist view, as has been assumed. Indeed, we cannot legislate in advance that cognitivists must be non-reductivists and noncognitivists reductivists about the thick. Christine Tappolet, Daniel Elstein, and Tom Hurka have recently pointed out that both cognitivists and noncognitivists can be reductivists.⁷ Furthermore, it may be that noncognitivists can be nonreductivists about the thick. Whether they can or not depends, as I argue in Sections VII and VIII, on whether noncognitivists can be particularists, and on whether a nonreductivist noncognitivist view can account for the supervenience and dependence of the evaluative on the nonevaluative.

B. The 'Evaluative' and the 'Descriptive'

The debate about the nature of thick concepts has been cast in terms of a contrast between the 'evaluative' and the 'descriptive'. While we do need to contrast the evaluative with something, it is far from obvious that the descriptive is the right candidate. If we contrast the evaluative with the descriptive, we obscure the possibility that to evaluate something might also be to describe it, when for the purposes of the debate we should use a way of drawing the contrast that is neutral between cognitivism and noncognitivism. Also, I will need to talk of the features of the world on which the evaluative is held to depend. While some will be happy to call these features 'descriptive'

7. Christine Tappolet, "Through Thick and Thin: Good and Its Determinates," *Dialectica* 58 (2004): 214; Elstein and Hurka, "From Thick to Thin," 516–17. However, while these writers acknowledge, contrary to what the literature on the thick seems to assume, that cognitivists can be reductivists about the thick, they also hold that it can be legislated in advance that noncognitivists must be reductivists.

there is something odd about this: it is bits of language that describe, and are descriptive, rather than features of the world. It is similarly odd to call properties descriptive. Similar considerations apply to a factual/evaluative contrast.

It may seem as if the obvious way to proceed here is to contrast the 'evaluative' with the 'natural', but this is also problematic. First, while there is nothing wrong with calling features or properties of the world 'natural', it is odd to talk of the natural, as opposed to the evaluative, content of a concept. Second, as with the contrast between the evaluative and the descriptive, the distinction between the 'evaluative' and the 'natural' obscures, at least, the possibility that the evaluative might be natural. And third, there are difficult issues regarding how the natural should be characterized. While these issues are interesting and important, they are also a complication that can be ignored for present purposes.

A further option is simply to draw the contrast with the nonevaluative, and this is the option that I take. This version of the contrast has a number of advantages in the present context. First, it is exhaustive and exclusive where the others are not. Second, we can talk of both nonevaluative content and nonevaluative features without strain. Third, there are no separate difficulties about how to characterize the nonevaluative in the way that there are about the natural.

C. Dependence and Supervenience

Most metaethicists agree that some dependence relationship holds between the evaluative and the nonevaluative: things get to be the way they are evaluatively because of, or in virtue of, the way they are nonevaluatively.⁸ In metaethics, this dependence relationship is often expressed as the claim that the evaluative supervenes on the nonevaluative. Strictly speaking, however, supervenience claims are covariance rather than dependence claims. In this context, the supervenience claim in slogan form is that there can be no evaluative difference without a nonevaluative difference.⁹

8. This is sometimes expressed by saying that evaluative properties are not free-floating. Different theories give different accounts of the nature of the dependence. One might hold that evaluative properties ontologically depend on nonevaluative properties or that evaluative responses somehow depend on beliefs about the nonevaluative, to mention only two examples (and these are not meant to be mutually exclusive). What I say should be construed as neutral between different accounts of the dependence relationship, except when otherwise indicated.

9. This by itself is not sufficient to establish that the evaluative depends on the nonevaluative. A-properties (A-judgments) can supervene on B-properties (B-judgments) without the A-properties (A-judgments) depending on the B properties (B-judgments), since all that is required is covariance. See Brian P. McLaughlin, "Varieties of Supervenience," in *Supervenience: New Essays*, ed. Elias E. Savellos and U. Yalcin (Cambridge: Cam-

One of the issues that will be raised by the discussion of this article is whether irreducible thickness presents a challenge to the dependence and supervenience theses. Without wishing to spoil the ending, it will turn out that irreducible thickness (and its attendant view of evaluation) is no threat to the dependence thesis if that thesis is construed ontologically. Whether it is a threat to the supervenience thesis depends too on how that thesis is construed, but the issues here are trickier. I discuss this in detail in Section VII.¹⁰

D. Concepts and Meanings

Finally, the issue regarding the reducibility of thick concepts is sometimes characterized as the issue of whether a thick concept can be analyzed into a nonevaluative component and a thin evaluative component, and sometimes as the issue of whether the meaning of a thick term can be analyzed into a nonevaluative meaning component and a thin evaluative meaning component. Herein I assume that concepts are at the level of sense rather than at the level of reference. I also assume that thick terms express thick concepts, and that the thick concept is what we grasp when we understand the thick term. Thus, for the purposes of this article, I take it that the meaning of a thick term, such as 'tactful', is much the same thing as the content of the concept—*tactful*—expressed by that term.

With these preliminaries out of the way, we can now move onto an initial characterization of the distinction between thick and thin evaluative concepts.

III. THICK AND THIN: AN INITIAL SENSE OF THE DISTINCTION

1. The recent film adaptation of Evelyn Waugh's *Brideshead Revisited* is bad.
2. "Sadly, Julian Jarrold's big-screen *Brideshead Revisited* . . . is *dull*, *perfunctory* and *moderately efficient*. . . . At the centre is a *dim* performance from Matthew Goode as Charles Ryder."¹¹

bridge University Press, 1995), 16–59. Compare Nicholas Sturgeon, "Doubts about the Supervenience of the Evaluative," in *Oxford Studies in Metaethics*, vol. 4, ed. Russ Shafer-Landau (Oxford: Oxford University Press, 2009), 60.

10. Up until Sec. VII, what I say about supervenience should be understood to be neutral between ascriptive (evaluative judgments supervene on nonevaluative judgments) and ontological supervenience (evaluative properties supervene on nonevaluative properties), since the discussion is neutral between cognitivism and noncognitivism. See James C. Klagge, "Supervenience: Ontological and Ascriptive," *Australasian Journal of Philosophy* 66 (1988): 461–70. I discuss both ascriptive and ontological supervenience in detail in Sec. VII. My thanks to a referee for pointing out the importance of this distinction for my discussion.

11. Philip French, "Brideshead Revisited," *Observer*, October, 5, 2008, <http://www.guardian.co.uk/film/2008/oct/05/periodandhistorical.drama> (emphasis added).

“It’s a *dutiful*, good-looking slice of heritage cinema so *empty* that you expect Keira Knightley to show up at any minute.”¹²

We can call the evaluative concept in (1) ‘thin’ and those in (2) ‘thick’. It is clear that (1) and (2) are both evaluations, but it is also clear that they are different. Thick evaluations are more specific than thin evaluations. It is perhaps tempting to attempt to explain this difference by saying that (1) is a thin evaluation because it uses a concept that is purely evaluative and that (2) is a thick evaluation because it uses concepts like DULL, PERFUNCTORY, EMPTY, and DIM that both evaluate the relevant object and ascribe certain non-evaluative features to it.¹³ Those who are so tempted are usually further tempted to say that the distinction between thin and thick can be drawn in terms of content: thin concepts have a unitary character in virtue of having purely evaluative content; thick concepts, on the other hand, have a dual character in virtue of having some combination of evaluative and nonevaluative content.

This seemingly dual character has led some philosophers to think that thickness is generated by a combination of elements, elements that can in principle be separated by analysis. According to these philosophers, thickness is reducible in the sense that the elements that generate thickness can be disentangled.¹⁴ The simplest form of such a reduction would be a two-component analysis of a thick concept. One component would be a thin evaluation, the other a specification of the nonevaluative features things must have if they are to merit the thin evaluation. However, two-component analyses do not simply conjunctively combine nonevaluation and evaluation. The view is not that a thick evaluative term like ‘tactful’ can be analyzed into some non-evaluative predicate ‘T’ (ascribing a specific nonevaluative property or set of properties) and some evaluative predicate, for example,

12. Sukhdev Sandhu, “Brideshead Revisited: Brideshead Redecorated,” *Telegraph*, October, 2, 2008, <http://www.telegraph.co.uk/culture/film/filmreviews/3561525/Brideshead-Revisited-Brideshead-redecorated.html> (emphasis added).

13. I will use small caps to indicate concepts and single quotes to indicate terms. Throughout the article, I talk of thick concepts and thick evaluative concepts as if thick evaluative concepts exhaust all the thick concepts there are. There may also be thick deontic concepts, however.

14. This quick sketch of a reductive view is not meant as an accurate representation of all reductive views. There are those, notably Simon Blackburn, whose reductive view of the thick is not accurately represented as a view about the analyzability of thick concepts. Blackburn holds that there are no thick concepts, only loaded terms, that is, that the evaluative aspect of a so-called thick term is no part of the content of the concept expressed by that term. However, it is still accurate to characterize Blackburn as holding that the elements that generate the appearance of thickness can be disentangled. Given the aims of this article, I will largely ignore the complications introduced by Blackburn’s view.

'good'. This is because an analysis that was merely conjunctive would not capture an aspect of what is conveyed by 'tactful': that it is *due to its being T* that the action is good. Thus, a two-component analysis of 'tactful' would have to say something like: "x is tactful if and only if x is T and is good in virtue of being T."

This is the simplest reductive account of the thick, and it is against this that Bernard Williams and Hilary Putnam, among others, have wielded the disentangling argument.

IV. THE DISENTANGLING ARGUMENT

The disentangling argument as it is presented in the literature on the thick is supposed to scupper the chances of such two-component analyses.¹⁵ Some have even suggested that it shows that evaluation and nonevaluation are so deeply entangled in the content of thick concepts that the separation a reductive analysis requires is impossible.¹⁶ However, the argument achieves neither of these things. It is only if this two-component analysis is combined with a certain view of the nature of evaluation that it is vulnerable to the disentangling argument.¹⁷ Moreover, there are other forms that reductive analyses of the thick can take that are not vulnerable to the disentangling argument.¹⁸

In this section, I set out the argument and then go on to discuss three different versions of reductive analyses for thick concepts. I argue that the disentangling argument succeeds against two-component reductive analyses that hold the nonevaluative component of the analysis to be extensionally equivalent to the thick concept (RA1). These analyses accept universalizability and codifiability. However, the argument fails against two-component reductive analyses that deny universalizability and codifiability (RA2). And it also fails against reductive analyses that endorse universalizability and codifiability but allow that

15. This is the disentangling argument as it is presented in the literature on thick concepts by, for example, Williams, *Ethics and the Limits of Philosophy*, 140–42; Putnam, "Entanglement of Fact and Value," 38–39; Taylor, "Ethics and Ontology," 360; Millgram, "Thick Ethical Concepts," sec. 2; Dancy, "In Defense of Thick Concepts," 263; and Elstein and Hurka, "From Thick to Thin," 518. McDowell is cited as the source of the argument, except in Dancy's case. I argue below that it is misleading to suggest that McDowell himself makes this argument.

16. Putnam, "Entanglement of Fact and Value," 34. Compare Williams, *Ethics and the Limits of Philosophy*, 129.

17. I disagree here with Elstein and Hurka, "From Thick to Thin," who suggest that all versions of this form of two-component analysis will be vulnerable to the disentangling argument.

18. It is worth pointing out that I am not claiming that these reductive strategies are ultimately successful, merely that they do not run aground on the disentangling argument.

the nonevaluative meaning component of an analysis of a thick term does not fully specify the set of nonevaluative properties that make the application of the relevant thick predicate appropriate (RA3). To see all this, we must first consider the argument.

A. *The Disentangling Argument*

If the two-component reductive view of the thick were correct, it is said, an outsider who did not share the relevant evaluative point of view would be able to grasp the extension of the thick concept (and so come to mimic the practice of insiders regarding the application of this concept) merely by observing the nonevaluative features of different cases of the concept's correct application. But, it is contended, it is not plausible to suppose that an outsider could achieve this. The support for the claim that it is not plausible to suppose that an outsider could achieve this comes from the plausible thought that the extension of a thick concept cannot be determined without making some evaluations, that is, without sharing the insiders' evaluative point of view. If this is the case, the argument concludes, it is not plausible to suppose that the two-component reductive view of the thick is correct.¹⁹

If we grant that it is indeed the case that we cannot determine the extension of a thick concept without making some evaluations,²⁰ this argument is successful against accounts of the thick like those of R. M. Hare and C. L. Stevenson, which hold that nonevaluative extensionally equivalent terms are available for thick terms, and that it is the nonevaluative meaning of thick terms that drives their application from case to case. To use Williams's terminology, on these accounts it is the nonevaluative meaning component *alone* that guides the application of these terms to the world.²¹

Consider the two-component analysis of 'tactful' again, the first and simplest form of a reductive analysis for thick concepts:

RA1: x is tactful if and only if x is T (where 'T' is a nonevaluative

19. The argument is deliberately cautiously phrased. This is faithful to its presentation in the literature and is what allows the different interpretations of it that I go on to discuss. It strikes some as obviously a bad argument. It is certainly not a good argument if it is meant to be an argument against reductivism about the thick in general, or an argument for irreducible thickness.

20. For the purposes of this article I take it that this is indeed the case. All of the positions discussed here except for those of R. M. Hare and C. L. Stevenson accept this claim, though they do not each interpret it in the same way.

21. C. L. Stevenson, *Ethics and Language* (New Haven, CT: Yale University Press, 1944), chap. 3; R. M. Hare, *The Language of Morals* (Oxford: Clarendon Press, 1952), 121; Williams, *Ethics and the Limits of Philosophy*, 141. Given universalizability, this is true once the evaluation has been made for the first time. See n. 22.

predicate referring to a specific set of nonevaluative properties) and x is good in virtue of being T.

On Hare's view, given his commitment to universalizability, once I have judged that something is good in virtue of being T, I am logically bound to judge good anything else that I judge to be T.²² 'T', on Hare's view, can be employed to pick out all those things to which I would apply 'tactful'. It follows that, for Hare, a nonevaluative extensionally equivalent term is available for any thick term. This version of the reductive view thus holds exactly what the disentangling argument says is impossible: an outsider could come to mimic the insiders' practice with a thick term simply by attending to the nonevaluative features of different cases of the thick term's correct application. Since these features would neatly repeat from case to case, judgments regarding when a particular thick term applies would be easily codifiable.

B. A Version of the Two-Component Analysis That Escapes the DA

However, a defender of the two-component analysis is not necessarily committed to holding that the nonevaluative component of a thick concept can be expressed by a nonevaluative term that is extensionally equivalent to the thick term. According to the general version of the analysis, a person or an action must not only be T but must also be good in virtue of being T, in order to count as tactful. The two-component analysis is therefore compatible with the plausible thought that we cannot determine the extension of a thick term without making some evaluations. It is possible, on this analysis, that an action can be T (and even good) and yet not be tactful because it is not also good in virtue of being T.²³ A defender of the two-component analysis could drop Hare's universalizability requirement and claim that, if something is T, whether that thing is also good in virtue of being T is dependent on context. Thus, if I judge that an action is T, I am not logically bound to judge that it is good. However, the nature of the contextual dependence of the goodness of T would have to be uncodifiable. If it were codifiable into a manageable and useful set of finite principles, there would be no reason to hold that the nonevaluative component of the concept was 'T' rather than 'T+', where

22. For R. M. Hare, a judgment is universalizable when "it logically commits the speaker to making a similar judgment about anything which is either exactly like the subject of the original judgment, or like it in the relevant respects" (*Freedom and Reason* [Oxford: Oxford University Press, 1963], 139). I take it that "like it in the relevant respects" means in the case of evaluative judgments that the thing shares the nonevaluative features that were the speaker's reasons for making the original judgment.

23. Compare Tappolet, "Through Thick and Thin," 214–15.

'T⁺' would be 'T' plus the principles detailing the nature of the contextual dependence. 'T⁺' would once again be a nonevaluative component extensionally equivalent to 'tactful', and thus the analysis would again fall foul of the disentangling argument. Dropping both universalizability and codifiability would thus result in a two-component analysis that was compatible with the plausible thought that we cannot determine the extension of a thick term without making some evaluations.

I will label this form of reductive analysis, a two-component analysis that denies universalizability and codifiability, RA2.

RA2: x is tactful if and only if x is T (where 'T' is a nonevaluative predicate referring to a specific set of nonevaluative properties) and x is good in virtue of being T, and whether x is good in virtue of being T is uncodifiably dependent upon context such that universalizability with regard to T is not possible.

This is roughly the form that Blackburn's response to the disentangling argument takes (he combines it with a claim about evaluation being no part of the meaning of thick terms). Fattists, he says, may choose to overlook the fact Pavarotti was fat to an extent that would normally repel them because he was "so transcendently uncontentable in other ways." The evaluation, in other words, can play a role in determining the extension of the term 'fat↓' as said with a sneer by fattists. A person must not only be fat; it must also be the case that a '↓' is called for: "It is quite compatible with attitude being carried by the tone, that it then play a role in determining the extension and ruling out of the extension things which, for quite different reasons, escape the attitude."²⁴ For it to truly capture that it is *evaluation* doing the work here, the analysis must have it that the things that can count as reasons for ruling some fat people out of the extension of 'fat↓' are uncodifiable. If the analysis held that these reasons were codifiable, the nonevaluative component would be 'fat⁺' rather than 'fat', and evaluative judgments with respect to 'fat⁺' would be universalizable. We would again have an analysis that produced a nonevaluative term extensionally equivalent to the thick term, and the disentangling argument would again have purchase.

C. Other Reductive Analyses That Escape the DA

A defender of a reductive account of the thick need not deny universalizability and codifiability. She need not deny that there is some

24. Simon Blackburn, "Through Thick and Thin," *Proceedings of the Aristotelian Society, Supplementary Volume* 66 (1992): 290. Of course, we need not give the expressivist spin to this point.

finite set of nonevaluative properties such that whenever they are instantiated, the thing they are instantiated by is good (or bad) in virtue of those properties, and therefore that the thick predicate in question applies to that thing.²⁵ She need only deny that that set of nonevaluative properties is specified in the meaning of a thick term. She thus denies that a reductive analysis is possible that would yield a nonevaluative component with the same extension as the thick term.²⁶

If this strategy is to be compatible with, and so evade, the disentangling argument, then that argument needs to be interpreted in a particular way. According to that argument, it is implausible that an outsider (someone who did not make the relevant evaluations) could come to mimic the practice of insiders with regard to a particular thick concept. The way this point is most commonly made is to question the possibility of there being a nonevaluative term that is extensionally equivalent to the thick term.²⁷ On this second strategy for avoiding the disentangling argument, it would not be the case that there could never be a nonevaluative term that is extensionally equivalent to the thick term, nor could it be that it is impossible for an outsider to come to mimic the practice of insiders. (For, according to this strategy, there is a set of nonevaluative properties such that each time it is instantiated it makes the thick concept applicable.) Instead, the idea must be that it is very unlikely that an outsider could come to mimic the practice of insiders while remaining an outsider.

On this strategy, the claim must be that it is very likely that an outsider will have to engage with the evaluations of the insiders in order to be able to apply the concept correctly, because it is too difficult to detect the set of nonevaluative properties present in each

25. For simplicity's sake, I have put this in terms of properties, but the same point could just as easily be made in terms of judgments.

26. Elstein and Hurka, "From Thick to Thin," give two strategies for reduction of this sort, one two-part and one three-part with an embedded thin evaluation. They hold the nonevaluative component "specifies good- or right-making properties to some degree but not completely, saying only that they must be of some specified general type but not selecting specific properties within that type—that is left to evaluation" (*ibid.*, 521). What makes it plausible that their strategies are reductionist yet nonetheless accept universalizability and codifiability is that *none* of the nonevaluative meaning component is held to be determinate or specific. Once *some* is, as in Blackburn's 'fat!' example, if codifiability is accepted, it is not clear what rationale there is for restricting the nonevaluative meaning component to 'fat' rather than allowing it to be 'fat+'. Thus, it is not clear what rationale there is for rejecting universalizability with regard to 'fat+', and for rejecting the claim that there is a nonevaluative extensionally equivalent term for 'fat!'.

27. David Wiggins makes it briefly and with a pleasing rhetorical flourish: "In any case, who can seriously believe that there is any chance of giving the 'ethically uncommitted' or 'pure' description that 'scoundrel' and 'knave' piggy-back upon?" (*Ethics: Twelve Lectures on the Philosophy of Morality* [Cambridge, MA: Harvard University Press, 2006], 379).

case of the concept's correct application directly (that is, without going through the evaluations the concept is used to make). This is compatible with being able to reductively analyze the content of a thick concept into a thin evaluative component (or components) and a nonevaluative component. The nonevaluative component, however, will not fully specify the set of nonevaluative properties things must have in order to fall under the concept. It will only characterize what nonevaluative sort they must be.

Sticking with 'tactful' as an example, 'tactful' on this view might be analyzed as something like:

RA3: *x* is tactful if and only if *x* is good and there are properties *P*, *Q*, and *R* (unspecified) of nonevaluative sort *T* (specified) such that *x* has properties *P*, *Q*, and *R*, and *P*, *Q*, and *R* make anything that has them good.

'*T*' might be something like "having something to do with showing concern for the feelings of others."²⁸

To sum up the discussion so far, the disentangling argument is supposed to undermine the plausibility of reducing thick concepts to thin evaluative content plus nonevaluative content. I have argued that the disentangling argument does not achieve this much. It only succeeds against two-component reductive analyses that hold the nonevaluative component of the analysis to be extensionally equivalent to the thick term (RA1). However, it fails against two-component reductive analyses that deny the nonevaluative component is extensionally equivalent to the thick term by denying universalizability and codifiability (RA2). It also fails against reductive analyses that endorse universalizability and codifiability but deny that the nonevaluative component is extensionally equivalent to the thick term, by allowing that the nonevaluative meaning component of an analysis of a thick term does not fully specify the set of nonevaluative properties that make the application of the relevant thick predicate appropriate (RA3).

Contrary to what some have assumed, then, the disentangling argument does nothing to show that evaluation and nonevaluation are so deeply entangled in the content of thick concepts that the separation a reductive analysis requires is impossible. It thus does nothing to clarify what this 'deep entanglement' is supposed to be. It is my proposal that we can find such clarification if we turn to what is cited as the source of the disentangling argument.

28. Compare Elstein and Hurka, "From Thick to Thin," 521, 526. I am grateful to Daniel Elstein for helpful discussions about this reductive strategy.

V. DISENTANGLING AND SHAPELESSNESS (1)

I will continue to use 'disentangling argument' to refer to the argument discussed in the previous section, that is, the argument given by Williams and Putnam and those others who have written on the thick. A passage in McDowell's "Non-cognitivism and Rule-Following" is commonly cited as the source of this argument.²⁹ In fact, McDowell himself does not make the disentangling argument as it is set out above. What he does say, however, will be useful for making clear in the next section just what it is that nonreductivists about the thick have to hold.

It is helpful to begin by asking two questions:

1. What is 'disentangling'?
2. What is the reason for thinking that disentangling is not, or not likely to be, possible?

Answers to these questions diverge depending on whether we are considering the (a) disentangling argument or (b) McDowell's antidisentangling remarks. The answers found in b pave the way for an account of the irreducibly thick.

McDowell uses 'disentangling' to refer to separating an evaluative judgment into two components. The first is sensitivity to a non-evaluative feature of the world, and the second is a 'propensity to a certain attitude'. What thwarts disentangling, McDowell claims, is that evaluative concepts are nonevaluatively shapeless. By contrast, according to the disentangling argument, 'disentangling' refers to the reductive analysis of thick concepts into thin evaluation and nonevaluation, and the reason that disentangling is thwarted is that if it were possible, an outsider would be able to grasp the extension of a thick term without making any evaluations. But it is not plausible to think that an outsider would be able to grasp the extension in this way. Although McDowell does make this point about the graspability of the extensions of evaluative terms, this is not his primary point, as will become clear below.

29. John McDowell, "Non-cognitivism and Rule-Following," in *Mind, Value, and Reality* (Cambridge, MA: Harvard University Press, 1998), 198–220. Unlike others (see n. 15), Williams does not cite this article but cites McDowell's "Are Moral Requirements Hypothetical Imperatives?" (*ibid.*, 77–96) and "Virtue and Reason" (*ibid.*, 50–76) as the sources of the disentangling argument. Williams says that the idea (which he finds in these articles) that it might be impossible to pick up an evaluative concept unless one shared its evaluative interest is basically a Wittgensteinian idea that he first heard expressed by Iris Murdoch and Philippa Foot in a seminar in the 1950s. In my view, the disentangling argument found in "Hypothetical Imperatives?" and "Virtue and Reason" does not concern the acquisition of evaluative concepts, and is distinct from the one in "Non-cognitivism and Rule-Following."

McDowell's targets in "Non-cognitivism and Rule-Following" are Hare and J. L. Mackie. He takes them both to hold that values are not genuine features of the world but are projected by us onto a value-neutral screen of how things really are. Hare and Mackie, says McDowell, are committed to explaining value experience using a disentangling maneuver: they must hold that when we ascribe value to something, what is actually happening can be disentangled into two components. Since they think there are no evaluative features in the world, competence with an evaluative concept must involve sensitivity to a genuine feature of the world, on the one hand, and a propensity to a certain attitude that gives that feature of the world the appearance of having the value in question.³⁰

McDowell's aim is to undermine the Hare/Mackie account of value experience in general, and he nowhere makes the distinction between thick and thin evaluative concepts. In the passages that McDowell references specifically, Mackie is writing about thin normative concepts, and Hare about both thick and thin.

Crucially, McDowell's objective here is not to undermine the possibility of the reductive analysis of thick concepts, but to undermine the view that for any evaluative term there must be a corresponding nonevaluative classification. Hare's version of this view was that for any evaluative term there must be a corresponding value-neutral classification that is a part of the meaning of the evaluative term (the 'descriptive' meaning). But Mackie's view was that, while there must be such a nonevaluative classification, this is no part of the meaning of an evaluative term.³¹ That McDowell takes himself to be undermining both versions of the view, and that he is concerned with evaluative concepts in general, thick and thin, should be sufficient to show that his notion of disentangling is not the notion of disentangling elements of conceptual content. Thin evaluative concepts, after all, are not supposed to have any nonevaluative elements to their content.³²

The main point of "Non-cognitivism and Rule-Following" is thus to undermine the plausibility of holding that there are nonevaluative classifications corresponding to evaluative classifications—whether or

30. McDowell, "Non-cognitivism and Rule-Following," 200–201. McDowell recognizes that this projectivist formulation does not strictly fit Hare's account. However, he says that the differences between Hare and Mackie are not relevant for the purposes of his article, and that the disentangling maneuver can easily be redescribed to reflect Hare's position. *Ibid.*, 201 n. 7.

31. Hare, *Freedom and Reason*, chap. 2, esp. secs. 2.4–2.6; J. L. Mackie, *Ethics: Inventing Right and Wrong* (Harmondsworth: Pelican, 1977), 86.

32. McDowell, "Non-cognitivism and Rule-Following," 200–203, esp. notes 4, 5, 7, 8, and 10. Of course, none of this should be taken to imply that McDowell's account, if successful, does not also undermine the possibility of the reductive analysis of thick concepts.

not we consider the nonevaluative classifications to be a part of the meaning of evaluative terms. The antidisentangling remarks are supposed to achieve this.³³ It is true that McDowell says that if (his sort of) disentangling were possible, then an outsider would be able to grasp the extension of an evaluative term without making any evaluations. But, unlike the disentangling argument, he does not stop at the graspability point. He goes on to offer a reason why the outsider will not be able to grasp the extension of an evaluative term without making some evaluations. We get to what this reason is by thinking about supervenience.

'Disentangling' on McDowell's view, as I have said, requires that there are nonevaluative classifications corresponding to evaluative classifications. And it makes no difference whether the evaluations we are concerned with are thick or thin. If we allow that evaluative classifications supervene on nonevaluative classifications, some might think that this points in favor of disentangling: that supervenience gives us reason to think that there must be nonevaluative classifications corresponding to evaluative classifications. According to McDowell this would be a mistake:

Supervenience requires only that one be able to find *differences* expressible in terms of the level supervened upon whenever one wants to make different judgements in terms of the supervening level. *It does not follow from the satisfaction of this requirement that the set of items to which a supervening term is correctly applied need constitute a kind recognisable as such at the level supervened upon.* In fact supervenience leaves open this possibility *which is just the possibility my scepticism* [that the disentangling maneuver can always be brought off] *envisages*: however long a list we give of items to which a supervening term applies, described in terms of the level supervened upon, there may be no way, expressible at the level supervened upon, of grouping just such items together.³⁴

The idea that there is no matching nonevaluative classification for an evaluative classification (whether or not it is a part of the meaning of the evaluative term) is the idea that the evaluative classification is nonevaluatively shapeless. If evaluative concepts are nonevaluatively shapeless, then disentangling will not be possible, for there will be no 'genuine feature of the world' for the outsider to cotton onto. Thus, the outsider will not be able to succeed in coming to mimic the insiders' practice with regard to a particular evaluative concept. If eval-

33. In the rest of the article, McDowell takes himself to be removing a prejudice, that is, the prejudice that such nonevaluative classifications are required in order for our evaluative practice to count as rational.

34. *Ibid.*, 202 (emphasis added).

uative concepts are shapeless in this way, this explains why it is that outsiders cannot grasp the extension of a thick term without engaging with the evaluations of the insiders.

An evaluative concept is nonevaluatively shapeless if there is no nonevaluative classification corresponding to the evaluative classification. It is tempting to think that a classification is something done by us, something effected by, for example, the predicates we apply to the world. But this is not all that McDowell has in mind. His claim is not merely that evaluative classifications may be nonevaluatively shapeless in the sense that we lack a term for the matching nonevaluative property (or ‘genuine feature’) that is nonetheless there in the world: “The point is not merely that the language may lack such a term, a gap that might perhaps be filled by coining one.” Even if we were to coin such a term, he says, “such a coinage might not be learnable except parasitically on a mastery of the full-blown evaluative expression.”³⁵

What, then, is shapelessness supposed to be? The ‘shape’ of a concept, we might say, is what all the things falling under the concept have in common: the unifying feature or the real resemblance.³⁶ If evaluative concepts are nonevaluatively shapeless, then that unifying feature or real resemblance is not nonevaluative: it is not simply that we lack a term for the nonevaluative feature of the world that we are nonetheless sensitive to in applying our evaluative concept—it is not there. Evaluative concepts are nonevaluatively shapeless, if McDowell is correct, because the commonality that unifies instances across a range of cases of the concept’s correct application is evaluative.³⁷ This

35. Ibid., 202 n. 8. One might think that a crucial challenge to McDowell is why sharing the relevant evaluative perspective should accommodate the shapelessness problem. McDowell’s view is that without that perspective the relevant unifying (‘shapely’) feature is not discernible. The nonreductivist need not adopt McDowell’s particular formulation, however. We can recast this more simply as the point that if one lacks the evaluative concept, one will be unable to identify instances of the evaluative property. See Wiggins, *Ethics*, 334. I thank an editor for pressing me to make this point clear.

36. It may be the case that many nonevaluative predicates that apply in virtue of the application of other nonevaluative predicates (e.g., ‘animal’) are shapeless at the base level. However, that is not sufficient to make these predicates shapeless in the sense that is being discussed here, for the concepts expressed by such predicates do have a nonevaluative shape: the real resemblance or unifying feature between different instances of the concepts application is nonevaluative. In the animal case, the shape is the nonevaluative property of animality.

37. Another point in favor of my claim that McDowell is not concerned specifically with the thick is that this shapelessness point has been used in the debate in metaethics about the so-called pattern problem, where the focus has been on thin normative concepts. Compare Frank Jackson, Philip Pettit, and Michael Smith, “Ethical Particularism and Patterns,” in *Moral Particularism*, ed. Brad Hooker and Margaret Little (Oxford: Oxford University Press, 2000), 79–99; Jonathan Dancy, “Can the Particularist Learn the Difference

is nonetheless consistent with supervenience because supervenience merely requires that there be (that we judge there to be) non-evaluative differences where there are (where we judge there to be) evaluative differences between cases. By itself, supervenience cannot generate the requirement that there be matching nonevaluative classifications for evaluative classifications.

McDowell's antidisentangling remarks can thus be recast as follows: if you hold that values are not genuine features of the world, you are committed to disentangling. To be committed to disentangling is to hold that correctly making an evaluative judgment has two components. The first involves detecting a nonevaluative property in the world; the second involves taking that property to ground the value in question. Commitment to disentangling thus commits you to the claim that for every evaluative classification there is a corresponding nonevaluative classification (whether or not you take that nonevaluative classification to be a part of the meaning of the evaluative term). If there were such matching nonevaluative classifications, it would be the case that an outsider could come to mimic the practice of insiders with regard to a particular evaluative concept, coming to learn the extension of the concept without engaging with the evaluations the insiders use it to make. But it is not plausible to suppose that this is possible. Why not? One ready explanation is that the nonevaluative classification, the nonevaluative property, is not there for the outsider to detect: evaluative concepts are nonevaluatively shapeless.³⁸

Clearly nothing has been said that establishes that evaluative concepts are nonevaluatively shapeless.³⁹ At this point, I am merely mak-

between Right and Wrong?" in *The Proceedings of the Twentieth World Congress of Philosophy*, vol. 1, *Ethics*, ed. K. Brinkmann (Bowling Green, OH: Philosophy Documentation Center, Bowling Green State University, 1999), 59–72; David McNaughton and Piers Rawling, "Naturalism and Normativity: Descriptivism, Normativity and the Metaphysics of Reasons," *Proceedings of the Aristotelian Society: Supplementary Volume 77* (2003): 23–44; Sabina Lovibond, "Naturalism and Normativity: Reply to McNaughton and Rawling," *Proceedings of the Aristotelian Society* 105 (2004): 187–203; and Brad Majors, "Moral Discourse and Descriptive Properties," *Philosophical Quarterly* 55 (2005): 475–94.

38. Another way to put this point is to say that nonevaluative shapelessness has as a consequence that evaluative judgments cannot be codified and are not universalizable. Codification will not be possible, since shapelessness means there can be no (useful) rule telling us what something has to be like, nonevaluatively speaking, to be courageous. Universalizability is not possible, for shapelessness means that the features that here are your reasons for correctly judging the action to be courageous may recur without you being logically required to make the same judgment. See n. 22.

39. McDowell does not even purport to provide such an argument. But, although there is no knockdown argument for shapelessness, there is equally no knockdown argument for nonevaluative shapefulness. As I say below, the rest of "Non-cognitivism and Rule-Following" is taken up with the attempt to remove a prejudice McDowell sees as operating against shapelessness.

ing clear what the notion of shapelessness is. Armed with this notion of shapelessness, we can now return to making sense of the notion of the irreducibly thick.

VI. DISENTANGLING AND SHAPELESSNESS (2)

The view that thick concepts are irreducibly thick is the view that evaluation and nonevaluation are deeply entangled in the content of thick concepts, and that the content of thick concepts cannot be disentangled into a thin evaluative component (or components) and a nonevaluative component. As we saw in Section IV, the disentangling argument did nothing to further the understanding of this view. All it told us, in effect, was that to grasp the extension of a thick concept, an outsider needs to engage with the evaluations of the insiders. But this graspability point works against only one version of the reductionist view. Two other versions of the reductive analysis of thick concepts are compatible with the graspability point. Using the shapelessness claim, however, we are now in a position to see what the nonreductionist about the thick needs to hold, in addition to the graspability point, in order to distinguish her view from a reductionist account of the thick. We can start by comparing her view with the type of account I labeled RA3.

A. Shapelessness, Irreducible Thickness, and RA3

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

RA3 holds that while there is a nonevaluative classification corresponding to a thick evaluative term, that nonevaluative classification is not part of the meaning of the thick term. It is compatible with the graspability point. According to this account, it would not be impossible for an outsider to grasp the extension of a thick concept without making any evaluations. It would, however, be very unlikely.

The view that thick concepts are irreducibly thick is distinguished from RA3 by the shapelessness claim. The defender of the irreducibly thick is committed to the claim that if a thick concept is irreducibly thick, then it is nonevaluatively shapeless.⁴⁰ If a thick concept is irre-

40. I have previously implied that the nonreductionist view of the thick is that all thick concepts are irreducibly thick. Strictly speaking, however, the defender of irreducible thickness need not commit herself to the claim that all the concepts that we have termed thick are irreducibly thick, just that at least some of them are, and those that are are nonevaluatively shapeless.

ducibly thick, it will be impossible, and not merely very difficult, for the outsider to detect a nonevaluative classification corresponding to the thick evaluative classification. The claim is not merely that such a classification is not a part of the meaning of the thick term, but that it is not there in the world to be detected. On this view, what something must be like in order to be, for example, tactful, cannot be specified nonevaluatively. It cannot be so specified because what it is to be tactful is not a matter of having a nonevaluative property. This means that the nonreductivist about the thick is committed to a certain view of the nature of properties. I return to this point in Section VII below.

If thick concepts are irreducibly thick, however, this is nonetheless consistent with there being some rough but not extensionally equivalent purely nonevaluative characterization of the sort of things the concept applies to. To that extent, the original outline of the distinction between thick and thin evaluation that I gave in Section III above is accurate: thick concepts are more specific than thin concepts. They are so because they narrow down the sorts of things that the concept can apply to. The defender of irreducible thickness could thus say “*x* is tactful if *x* is of nonevaluative sort *T*,” where “*T*” means “having something to do with having or showing concern for others feelings.” But, even if we can give such a rough nonevaluative characterization, the nonreductivist is committed to the view that the exact properties of sort *T* that make things tactful are nonevaluatively shapeless, and thus that judgments regarding these properties cannot be codified, and thus that there is no set of properties of type *T* with respect to which judgments that an act is tactful can be universalized.

However, strictly speaking, the nonreductivist about the thick is not committed to there necessarily being even a rough but not extensionally equivalent purely nonevaluative characterization of the sort of thing a thick concept applies to. For one thing, it may be that we lack the appropriate (nonevaluative) terminology to give this characterization.⁴¹ For another, it may be that shapelessness means that we will be unable to come up with such a nonevaluative classification that is to function even as a necessary condition for the appropriate application of the concept. It might be possible to come up with an example of a tactful action, say, which does not show concern for

41. Compare James Griffin, “Values: Reduction, Supervenience, and Explanation by Ascent,” in *Reduction, Explanations, and Realism*, ed. David Charles and Kathleen Lennon (Oxford: Clarendon, 1992), 297–322; and Joseph Raz, “The Truth in Particularism,” in *Engaging Reason* (Oxford: Oxford University Press, 1999), 218–46. Both argue, though in different ways, that this point undermines supervenience. I argue below that it possibly undermines ascriptive supervenience, but not ontological supervenience.

others' feelings. The nonreductivist view, however, is compatible with the thought that we can narrow down the sorts of things that thick concepts apply to, since we can do this without committing to a nonevaluative shape for the concept.⁴²

B. Shapelessness, Irreducible Thickness, and RA2

RA2: *x* is tactful if and only if *x* is *T* (where '*T*' is a nonevaluative predicate referring to a specific set of nonevaluative properties) and *x* is good in virtue of being *T*, and whether *x* is good in virtue of being *T* is uncodifiably dependent upon context such that universalizability with regard to *T* is not possible.

RA2 is compatible with the claim that thick concepts are nonevaluatively shapeless. To return to Blackburn's example, 'fat' in 'fat↓' does not amount to a nonevaluative shape for the concept, for it does not amount to a matching extensionally equivalent nonevaluative classification. An outsider who latched onto the fact that every individual that 'fat↓' applied to was fat would not thereby be able to master the use of 'fat↓'.

To distinguish her view from RA2, the nonreductivist about the thick cannot merely claim that thick concepts are nonevaluatively shapeless. That is, her view cannot be merely that there are no extensionally equivalent matching nonevaluative classifications for thick concepts since this is compatible with RA2. The nonreductivist's claim is that across a range of cases of a thick concept's application there need not be a particular nonevaluative property or set of nonevaluative properties instantiated in each case, even when that does not amount to a nonevaluative shape for the concept. This is not an implausible claim; think of all the different actions that can be courageous, or all the different actions that can be tactful: leaving a conversation without drawing attention to oneself can count as tactful, but so can staying in a conversation and drawing attention to oneself. Being tactful can involve being polite, but it need not. It may even involve being impolite: interrupting someone while they are speaking may be impolite but tactful if it saves them from making the blunder you realize they are just about to make. (And, of course, politeness may be an evaluative property.)

This is compatible with saying that, for any particular thick concept, there may be such a nonevaluative property or set of properties instantiated in each case, as there is in the case of 'fat↓', where that

42. Compare Alan Gibbard, "Thick Concepts and Warrant for Feelings," *Proceedings of the Aristotelian Society: Supplementary Volume* 66 (1992): 275–77; Dancy, "In Defense of Thick Concepts," 277–78.

doesn't amount to a nonevaluative shape for the concept. The non-reductivist distinguishes her view from RA2 by denying that there must be such a nonevaluative property or set of properties instantiated in each case.

There is, moreover, an additional feature that distinguishes the nonreductive view from RA2 (and RA3). To see this, we need to consider a version of the reductive view not yet discussed.

C. Shapelessness, Irreducible Thickness, and RA4

There is still more to do to mark out irreducible thickness, for there is space for a further reductive analysis.⁴³ Again, I use 'tactful' as an example.

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

The defender of irreducible thickness, then, cannot simply rest her view on the claim that thick concepts are nonevaluatively shapeless, for RA4 accommodates shapelessness. Nor can she merely add to this the claim that there are not necessarily any particular nonevaluative properties instantiated in each instance of the correct application of the concept, for RA4 accommodates this point too.

To distinguish her view from RA4, the nonreductivist must make the further claim that the evaluative content of thick evaluative concepts is not thin evaluation. The nonreductivist about the thick is committed, in other words, to the claim that there can be evaluation without thin evaluation. To call an action tactful is to evaluate it. But, on this view, it is not to evaluate it by calling it tactful as a way of calling it good, or as a way of indicating that it is good to some extent in some particular way. The nonreductivist about the thick is committed to the claim that to call it tactful is in itself an evaluation. I return to this point in Section VII.

This is not the same as Susan Hurley's view that thin evaluative concepts are neither conceptually prior to nor independent of thick evaluative concepts, but it chimes nicely with that view.⁴⁴ To be a reductivist about the thick is to hold that thick concepts are reductively analyzable into a thin evaluative component (or components) and a

43. This analysis is reductive in the sense that it reduces thick concepts to a combination of thin evaluation and nonevaluation.

44. S. L. Hurley, *Natural Reasons: Personality and Polity* (New York: Oxford University Press, 1989), chap. 2.

nonevaluative component. This is presumably to hold that the thin is conceptually more basic and thus prior to the thick, and thus to be a centralist in Hurley's terms. To be a nonreductivist about the thick is to deny the conceptual priority of the thin, and to be a noncentralist in Hurley's terms.⁴⁵ However, this is not necessarily to make the further claim that the thick is prior to the thin. A nonreductivist about the thick can hold a no-priority view: thick and thin ethical concepts might be interdependent.⁴⁶

At this point we can take stock. If thick concepts are irreducibly thick, then no reductive analysis of any thick concept will be possible. But this does not mean that there is nothing to be said in the way of elucidation of irreducible thickness in general (or of particular thick concepts—but that is not the primary task of this article). From the above discussion, we can characterize irreducible thickness in the following way:⁴⁷

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

Of this account we can say the following: the first condition is required to distinguish the nonreductivist account from both RA2 and RA3, as we saw above. The second is required to distinguish nonreductivism from RA4: the object must have features of the relevant sort that ground the evaluative property; features that make the object, for example, tactful (thick evaluation), rather than features that make

45. Williams endorses this notion of noncentralism in the context of rejecting reductivism about the thick. Williams, *Ethics and the Limits of Philosophy*, 184 n. 7.

46. Hurley, *Natural Reasons*, chap. 2. On Hurley's view, the various ethical concepts get their sense through their positions in a network.

47. I call this a characterization and not an analysis because the view is that, strictly speaking (i.e., if we take as criteria of analyses that they be noncircular and that they involve the breaking down of a concept into simpler elements that together compose the whole), an analysis of these concepts is not possible.

the object good (thin evaluation).⁴⁸ Condition 3 is a further spelling out of both 1 and 2. “Determined . . . by evaluation on a case-by-case basis” must be understood both ontologically and epistemically. That is, whether or not a token action is, for example, tactful will depend ontologically on the particular nonevaluative features of the case in context, and there will be no nonevaluative shape for tactful actions across different token instances. Moreover, the user of the concept will have to make an evaluative judgment as to whether, for example, ‘tactful’ applies here, which will involve attending to the particular features of the case, and not applying any general principles, for there will be no such principles to be applied.

A nonreductivist about the thick is thus both a holist and a particularist. However, this is not a sufficient condition for nonreductivism—the positions I labeled RA2 and RA4 would also count as both holist and particularist, since they deny universalizability and codifiability. In this context, holism is the view that the features that in this case make an action tactful may in other cases play no such role, or they may even make an action tactless. The role that these features play will be dependent on context. Particularism here is the view that there will be no true and useful general principles telling us what actions have to be like to be tactful.

VII. IMPLICATIONS

The nonreductivist view as it has been developed here has a number of implications. I explore some central ones in this section.

A. *The Distinction between Thick and Thin*

This account of irreducible thickness is likely to be tied to a nonreductive account of the nature of evaluation in general. The nonreductivist is likely to hold that thick and thin evaluative concepts are both nonevaluatively shapeless. She is also likely to hold that thick

48. I am taking it for granted that evaluative properties are indeed dependent properties, for example, that something is made tactful *in virtue of* other, nonevaluative properties that it has. Since on this view there will be no such thing as *the* features that make objects tactful (no single set of features amounting to a nonevaluative shape), this gives rise to the expression that I employ: that features must be had that ground or make the object tactful or whatever. Someone might object that this is not sufficiently informative, since all it amounts to saying is that objects must be tactful for the concept to apply. But the characterization is not meant as an informative analysis of the concept tactful, or even as a strategy for the analysis of thick concepts in general, but as an elucidation of the nature of irreducible thickness.

concepts are more specific than thin concepts and that the distinction between thick and thin is a matter of degree.⁴⁹

But how is this to be captured if she also holds that a thick concept does not necessarily have any nonevaluative content? I suggest that what the nonreductivist ought to say here is that the distinction between thick and thin can be drawn in terms of the extent to which the sorts of things that the concept applies to can be restricted. But the restriction need not be carried out in nonevaluative terms (it could involve more specific evaluative terms). The narrower the range of things that the concept applies to, the thicker the concept.

It is plausible to think that there are almost no restrictions on what sorts of things can count as GOOD or BAD. Perhaps RIGHT and WRONG are slightly thicker if they only properly apply to actions. OUGHT might be slightly thicker too if it is only properly applied to actions that agents can perform.⁵⁰ CRUEL, and TACTFUL on the other hand, are much thicker, for they are only properly applied to the infliction of suffering and to being sensitive to others' concerns respectively.⁵¹

To illustrate, the particular example I used in the characterization of irreducible thickness above was TACTFUL, but the non-reductivist is likely to say almost the same of a thin evaluative concept, adjusting conditions 1 and 2 to take account of the wider range of the concept:

1. The concept (e.g., GOOD) applies to an object in virtue of it having other features, but where there is no nonevaluative shape for the concept.
2. The object has features that ground the relevant property (e.g., makes it good) and thus merits the application of the thin concept.
3. Precisely which features ground the property (e.g., make the object good), and thus make the application of the concept appropriate, are determined, in a way that cannot be specified in advance, by evaluation on a case-by-case basis.

49. Compare Samuel Sheffler, "Morality through Thick and Thin: A Critical Notice of *Ethics and the Limits of Philosophy*," *Philosophical Review* 96 (1987): 411–34. I say that it is likely that the nonreductivist will hold these things, because it may be possible to combine a nonreductive view of the thick with some other account of the thin.

50. It may seem as though I am now saying that the nonreductivist holds that these notions have necessary nonevaluative content, and thus that I am denying something I said above. My claim above though was that the nonreductivist holds that a thick concept does not necessarily *qua thick concept* have any nonevaluative content, not that no thick concept has any necessary nonevaluative content.

51. I take it that for the nonreductivist both 'suffering' and 'sensitivity' are evaluative notions.

B. Supervenience

Does the nonreductivist about the thick deny the supervenience of the evaluative on the nonevaluative? The answer to this question depends on how supervenience is construed. Two distinctions are relevant here: the distinction between local and global supervenience and the distinction between ontological and ascriptive supervenience. I discuss each of these in turn.

The claim that the evaluative locally supervenes on the nonevaluative is the claim that no two individuals (persons, actions, states of affairs) can differ evaluatively but not nonevaluatively: if two individuals are identical in all of their nonevaluative features, they must be identical in evaluative respects as well.⁵² Global supervenience obtains if no two possible worlds can differ evaluatively but not nonevaluatively: two worlds that are nonevaluatively identical must be identical in evaluative respects as well.

Supervenience, whether local or global, is most commonly held to be an ontological relation holding between classes of properties. However, since for the most part the discussion above relied on claims about supervenience that were supposed to hold for both cognitivists and noncognitivists, I did not assume that the supervenience in question was ontological supervenience.⁵³ Instead, the claims made about supervenience were neutral between ontological supervenience and ascriptive supervenience.⁵⁴ Ascriptive supervenience is the term given to a relation holding between classes of judgments. In this context, ascriptive supervenience is the view that, logically speaking, a person's evaluative judgments about things (or worlds) cannot differ unless her nonevaluative judgments about those things (or worlds) differ.⁵⁵

Given these distinctions—between local and global supervenience and ontological and ascriptive supervenience—it turns out that

52. I will assume that the supervenience in question is strong supervenience, that is, holding across possible worlds rather than only weakly, that is, within each possible world. Nothing in the discussion hangs on this assumption, however. (It is possible to extend the distinction between weak and strong supervenience to ascriptive supervenience: see Sturgeon, "Doubts about Supervenience," 59 n. 14.)

53. Noncognitivists are antirealists about the evaluative and so strictly speaking cannot accommodate an ontological construal of the supervenience claim. See Alan Gibbard, *Thinking How to Live* (Cambridge, MA: Harvard University Press, 2008), 89.

54. For this terminology and the distinction between ontological and ascriptive supervenience, see Klagge, "Supervenience: Ontological and Ascriptive."

55. *Ibid.*, 462. I interpret ascriptive supervenience to be the claim that we cannot make different evaluative judgments about two things/worlds without actually making different nonevaluative judgments about these two things/worlds. But this is controversial. Some think that ascriptive supervenience is merely the weaker claim that we at least judge that there are different nonevaluative judgments to be made in these cases, even if we are not actually equipped to make them.

irreducible thickness undermines the plausibility of local supervenience, both ontological and ascriptive. Moreover, if thickness is irreducible, it may also give us reason to doubt that global ascriptive supervenience obtains. The only kind of supervenience not threatened by irreducible thickness is global ontological supervenience.

Shapelessness, I said above, is compatible with supervenience. More specifically, it is compatible with global supervenience, for shapelessness does not require the possibility of two worlds differing evaluatively without differing nonevaluatively. Shapelessness is not compatible with local supervenience, however. At least, it is not compatible with it unless we understand the local supervenience base to be the entire world. Typically, however, local supervenience doctrines work with a conception of a restricted supervenience base. Recall that local supervenience is the view that no two individuals (persons, actions, states of affairs) can differ evaluatively but not nonevaluatively: if two individuals are identical in all nonevaluative respects, they must be identical in evaluative respects as well. To avoid triviality, this doctrine must give some principled way of restricting the supervenience base of an individual such that it does not include the nonevaluative features of the entire world in which that individual is situated.

However, if shapelessness is true of thick concepts, then they have no nonevaluative shape. Shapelessness thus leaves open the possibility that two individual actions can share all of their nonevaluative properties (however this is determined) and yet not both be tactful. In the one case, for example, some feature of the context might make it the case that the action is no longer tactful.⁵⁶ We can avoid this result by widening the supervenience base. However, if shapelessness is the case, we would have to widen the supervenience base to include the nonevaluative features of the entire world in which the action is situated to guarantee that an action identical to this one in all nonevaluative respects would be identical to it in evaluative respects. But this would render local supervenience a trivial doctrine. This point applies to both ontological and ascriptive forms of local supervenience.

Irreducible thickness is entirely compatible with global ontological supervenience, but what of global ascriptive supervenience? Global ascriptive supervenience holds that we cannot make different evaluative judgments about two worlds without making different nonevaluative judgments about these worlds, and that two worlds that we judge to be nonevaluatively identical we must judge to be evaluatively identical. This requires not merely that there is a nonevaluative sub-

56. Compare Brad Majors, "The Natural and the Normative," in *Oxford Studies in Metaethics*, vol. 4, ed. Russ Shafer-Landau (Oxford: Oxford University Press, 2009), 36–40.

vening base but that we are able to give a complete characterization of it in nonevaluative terms, that is, a characterization that is sufficient to "fix" the evaluative.

It seems to me that a nonreductivist about the thick, who holds that there is no necessary nonevaluative content to evaluative concepts, and who holds that evaluative concepts are nonevaluatively shapeless, should leave open the possibility that we are not equipped to give such a complete nonevaluative characterization of the subvening base. I do not have the space here to pursue this point further, however. For the moment, I merely want to note the interesting result that, according to the nonreductivist view, there is no guarantee that global ascriptive supervenience will hold. There is no guarantee, in other words, that we will be able to make different nonevaluative judgments about two worlds that we make different evaluative judgments about.

What of the dependence thesis? Nonreductivism about the thick is compatible with this thesis ontologically construed: that the evaluative depends on the nonevaluative, for evaluative properties must be grounded in nonevaluative properties. It is consistent with this thesis because it is consistent with the nonreductivist view that the features that make actions tactful, say, on a case-by-case basis, are nonevaluative features. The nonreductivist simply denies that in any particular case the nonevaluative features making an action tactful are the property of tactfulness.

Defenders of nonreductivism are not committed to anything that would require them to deny that in each token case the features that make an action tactful are nonevaluative features. Things are less clear if we consider dependence construed ascriptively. There is no guarantee on the nonreductive view that we will be able to give a purely nonevaluative characterization of the things the thick concept applies to, at the level of type, even when that nonevaluative characterization does not amount to a nonevaluative shape for the concept. Perhaps the nonreductivist can say that it is likely that we will be able to, for each token case, but there is nothing in the view that guarantees that we will.

C. Properties

The nonreductivist view is consistent with the dependence thesis because it is consistent with holding that the features that make actions tactful, say, on a case-by-case basis, are nonevaluative features. Because of shapelessness, however, the nonreductivist denies that in any particular case the nonevaluative features making an action tactful are the property of tactfulness. The nonreductivist about the thick is thus committed to denying that the long disjunction of all the non-

evaluative features of different instances of tactfulness constitutes the property of tactfulness.

McDowell does not, strictly speaking, provide an argument for the claim that evaluative concepts are nonevaluatively shapeless. What he does do is point out that supervenience leaves open that possibility. There is no knockdown argument for shapelessness, but apart from appeal to examples, the defender of the irreducibly thick can supplement her case for shapelessness in a number of ways. One is to defend an account of properties that rule out disjunctions such as the one described above from constituting a property.⁵⁷ Another is to give an account that makes clear that nonevaluative shapelessness is no bar to acquiring an evaluative concept, and nor is it a bar to our evaluative practice counting as rational.⁵⁸

D. Evaluation

I have left possibly the most interesting set of implications until last. One part of the nonreductivist account that clearly needs to be fleshed out in more detail is the claim that thick concepts are themselves evaluative, and not evaluative because of any link they have to thin evaluation. While I do not have the space to explore all of the implications of this here, it is worth noting what some of them are.

It is commonly assumed that evaluation must be either positive (good to some extent) or negative (bad to some extent), that is, that evaluation must have a pro or con flavor. Use of a thick concept that conveys a positive evaluation is thought to entail 'good to some extent or in some respect', and use of a thick concept to convey a negative evaluation is thought to entail 'bad to some extent or in some respect'. Furthermore, it is assumed, this entailment licenses the view that thick concepts have thin evaluative content. But if thick concepts are themselves evaluative, and not evaluative because of any link they have to thin evaluation, then this assumption is not warranted. Deciding that an action is tactful will be a matter of evaluating it, but deciding that it is tactful need not involve deciding that it is good or bad to some extent. Thick concepts are evaluative in their own right, on this view, just like thin concepts are.

Of course, this flies in the face of, to use Hume's words, thick terms "force an avowal" of their merit or demerit, as the case may be.

57. Much work has already been done on this point. See, e.g., Majors, "Moral Discourse," and Jonathan Dancy, "Nonnaturalism," in *The Oxford Handbook of Ethical Theory*, ed. David Copp (New York: Oxford University Press, 2006), 122–43. For a relatively new account, see Graham Oddie, *Value, Reality, and Desire* (Oxford: Oxford University Press, 2009).

58. Compare Jackson, Pettit, and Smith, "Particularism and Patterns," and Dancy, "Can the Particularist Learn the Difference?"

Courage and tact are thought to be necessarily good, cruelty and brutality necessarily bad. But the nonreductivist is on the side of Blackburn and others who deny that thin evaluation is a part of the semantic content of thick terms.⁵⁹ Nonreductivism is thus compatible with the thought that the thin evaluative valence of a thick concept is variable. This is the claim whether a particular instance of, for example, tact is good to some extent, bad to some extent, or neither will depend on the particular features of the case and will be a matter for substantive judgment.⁶⁰

Those who argue that thin evaluative content is not part of the semantic content of thick terms also argue that thick concepts are not inherently evaluative. The nonreductivist thus owes us an account of what it is for a concept or property to be evaluative in the first place.

One potentially fruitful avenue of pursuit would be to consider essential contestability as a mark of the evaluative, and to pursue the thought that evaluative concepts, as opposed to nonevaluative concepts, are concepts that characteristically admit of substantive disagreement that cannot be settled by fiat.⁶¹ We might find such essential contestability in nonevaluative domains, but in such cases it is likely that disputes could be settled by fiat, if settling them is required for some practical purpose.⁶² In the case of evaluative concepts, this is not possible.⁶³ Thick concepts pick out features that matter, normatively speaking, for practical, aesthetic, and perhaps epistemic evaluation.⁶⁴ These are features that have practical, aesthetic, or epistemic relevance. Precisely what relevance they have will depend, no doubt, on other features of the context. In advance of a particular case, however, it is true to say that where thick properties arise, they will be relevant.

59. Blackburn, "Through Thick and Thin"; Simon Blackburn, *Ruling Passions* (Oxford: Clarendon, 1998); Väyrynen, "Objectionable Thick Concepts."

60. See Blackburn, "Through Thick and Thin," and *Ruling Passions*; Dancy, "In Defense of Thick Concepts"; and Jonathan Dancy, *Ethics without Principles* (Oxford: Clarendon, 2004).

61. Compare Williams, *Ethics and the Limits of Philosophy*, chap. 8; Hurley, *Natural Reasons*, chap. 3, esp. 45–50; and David Wiggins, *Needs, Values, Truth*, 3rd ed. (Oxford: Clarendon, 1998), 314, 318. I don't mean to suggest that this would be the only mark of the evaluative, merely that this would be a good place to start.

62. Compare Hare, *Freedom and Reason*, 28.

63. It may be that an evaluative matter may have to be settled by fiat, for legal purposes, for example. But this is not reason to think that the evaluative dispute will have been settled.

64. Compare Williams, *Ethics and the Limits of Philosophy*, 129–30, 140. I do not mean to suggest here that the practical and the aesthetic necessarily exhaust the realm of the evaluative. There may be epistemic thick concepts that characteristically matter for epistemic evaluation.

A further potentially fruitful avenue of pursuit would be to spell out precisely what is involved in what everyone party to the thick concepts debate seems to accept, namely, 'evaluation drives extension'. This would involve a careful account of what it is to share an evaluative outlook. Finally, the nonreductivist could also pursue the possibility that there is something distinctive about the "in virtue of" relation that holds between evaluative properties and their grounds.

VIII. CONCLUSION

There is much that still needs to be done to flesh out and defend the nonreductive view, as the previous section makes clear. However, there has been a significant advance on the negative and metaphorical formulation of irreducible thickness that we met at the beginning of this article.

We can now characterize irreducible thickness positively and less metaphorically as follows:

1. The concept (e.g., TACTFUL) applies to an object in virtue of its having features of a certain sort. We may be able to give a rough purely nonevaluative characterization of the sort these features belong to in advance, but there is no nonevaluative shape for the concept. The concept will, however, have an evaluative shape: all instances of tactfulness will be tactful, for example.
2. The object has features of that certain sort that ground the relevant property (e.g., make it tactful) and thus merits the application of the thick concept.
3. Precisely which features ground the property (e.g., make the object tactful), and thus make the application of the concept appropriate, are determined, in a way that cannot be specified in advance, by evaluation on a case-by-case basis. 'Determined by evaluation' is to be understood both epistemically and ontologically.

It is now possible to see both why it is that nonreductivists about the thick deny that all moral judgments can be expressed using a few thin concepts such as GOOD, BAD, RIGHT, and WRONG (for they hold 2). Sections VI and VII flesh out to some extent what it is they need to defend in order to uphold this denial. Moreover, we now see the extent to which irreducible thickness, if it obtains, creates problems for the idea that the evaluative supervenes on the nonevaluative. Irreducible thickness confounds both local and ascriptive supervenience and may also create obstacles for the thesis that the evaluative depends on the nonevaluative, where that is construed ascriptively.

Whether a successful defense of irreducible thickness serves to undermine the fact-value distinction is more complicated. The fact-value distinction involves the view that values are not facts in some full-blooded, realist sense. A successful defense of irreducible thickness would thus undermine the fact-value distinction if it also showed that antirealists about the evaluative could not be nonreductivists.

The rationality and consistency of evaluative practice is at the core of the issue of whether a nonreductive account of the thick undermines the fact-value distinction. It is true that irreducible thickness is most comfortably at home in a realist non-naturalist setting. Non-evaluative shapelessness is, in that setting, no bar to the existence of evaluative properties. Rationality and consistency in evaluative practice are a matter of using evaluative concepts to correctly pick out evaluative properties. But that this is where irreducible thickness is most at home should not lead us to assume that no other metaethical view can accommodate it. While the commitment to shapelessness and the attendant view of properties rule out certain kinds of naturalist realists from being nonreductivists, it is not clear that it would necessarily rule out naturalism tout court.⁶⁵

More importantly, as far as the implications of irreducible thickness for the fact-value distinction go, it has not yet been established that antirealists cannot be nonreductivists. There seems to me to be no bar to being both a nonreductivist about the thick and an error theorist. If there is, however, it will be the same one that is a bar to noncognitivism: that the commitment to nonevaluative shapelessness (the denial of universalizability and codifiability) leaves the antirealist with no way to account for evaluative practice as a process of genuine, rational, concept application.⁶⁶ But this issue is far from decided. The question that needs investigating is whether antirealists can be both holists and particularists about the thick.

Of course, if nonreductivism about the thick makes trouble for ascriptive dependence and supervenience, as I have suggested it does, then it may prove difficult for the noncognitivist to account for dependence and supervenience. But even if a successful defense of irreducible thickness establishes this much, this would not by itself undermine the fact-value distinction. What it would do, as Williams suspected, is show that any fact-value distinction that there is does not lie at or near the surface of language.⁶⁷

65. Compare Dancy, "Can the Particularist Learn the Difference?" and Jackson, Pettit, and Smith, "Particularism and Patterns."

66. McDowell, "Non-cognitivism and Rule-Following," 217. Compare Dancy, "In Defense of Thick Concepts," 272–75; Blackburn, *Ruling Passions*, 308, 315; and Gibbard, *Thinking How to Live*, 106.

67. Williams, *Ethics and the Limits of Philosophy*, 130.