Irreducibly Thick Evaluation is not Thinly Evaluative

N. D. Cannon¹

Accepted: 14 July 2020 / Published online: 21 July 2020
© The Author(s) 2020

Abstract

In this paper, I criticize the pairing of irreducible thickness (the view that thick concepts do not reduce to thin evaluations plus some descriptive component) with the traditional view of evaluation which says evaluation is a matter of encoding good or bad in some way. To do this, I first explicate the determination view, which holds that irreducibly thick concepts are to thin concepts as determinates are to determinables. I then show that, even if the determination view did establish irreducible thickness, it would not have proven that the evaluative is well understood as being an instance of the determination relation; in order to do that, the determination view needs to show how the evaluative fit a general analysis of the determination relation. However, when the determination view attempts to fill in the analysis, we get implausible results—so implausible, I claim, that we should see the results as a *reductio* to the view. To generalize the criticism to *any* view like the determination view, I show that the same results ensue when we model the evaluative on mereology. Finally, I diagnose the general failure by claiming that the evaluative domain, as conceived by the defender of irreducible thickness, just does not have the structure to secure the tight connection between thick and thin concepts while also carving up our conceptual economy in a plausible way.

Keywords Metaethics · Thick Concepts · Evaluation · Determination Relation; · Determinate and Determinable · Ethics · Harcourt, Edward · Thomas, Alan · Antireductionism · Irreducibility of the Thick · Irreducible Thickness

✉ N. D. Cannon
nc402@kent.ac.uk

¹ Department of Philosophy, University of Kent, Canterbury, Kent, UK

1 Introduction

Evaluative concepts tend to be divided into two basic categories: Thin and thick. *Thin* concepts are those like GOOD and BAD, which seem to involve little more than evaluation; *thick* concepts, in contrast, are those like COURAGE and COWARDLY, and seem to involve both evaluative *and* nonevaluative content.¹ This division of evaluative concepts raises a question—how do they relate to one another? For my purposes, responses to this question fall in roughly two different categories.² The first view encompasses those interested in reducing thick evaluative concepts to some (possibly limited set) of thin concepts and is what I call *evaluative reductionism*. Evaluative reductionists suppose thick concepts can be broken up into a (thin) evaluative component and a nonevaluative part (which may involve an embedded thin evaluation). For example, evaluative reductionists maintain that COURAGE can be (roughly) analyzed into GOOD and RISKING HARM FOR A GOOD END, where it is in virtue of satisfying the latter part of the analysis that GOOD applies (Cf. Elstein and Hurka (2009): 526). This view is an important option and is notably defended by Simon Blackburn (1992, 2013) and Allan Gibbard (1992), but I propose to set it aside for now to focus on the opposing view, *irreducible thickness*.³

Irreducible thickness is the view that the thick do not break up in the way the evaluative reductionist suggests. Instead, thick concepts are unitary, wholly evaluative concepts—they cannot be analyzed into their component parts because there are no component parts to begin with.⁴ This bare-bones irreducible thickness can get fleshed out in a few different ways. Until recently, the typical way irreducible thickness gets spelled out is by trying to maintain that, despite the *sui generis* nature of thick concepts, they nevertheless necessarily entail thin evaluations (e.g., COURAGE necessarily entails GOOD, but is not reducible to it). This, I think, is an instance of a more general tendency to pair antireductionism about the thick with a traditional view of evaluation. However, a defender of irreducible thickness might instead propose the thick are indeed wholly evaluative, but not in virtue of entailing some thin evaluation; instead, thick concepts are evaluative *in their own right*—to evaluate something as courageous just is to employ the concept, COURAGE, without thereby employing GOOD. In other words, defenders of irreducible thickness can deny the traditional view of evaluation.⁵

In this paper, I will argue we should be pessimistic about views which spell out irreducible thickness in terms that also commit us to the traditional view of evaluation. To do this, I will primarily target the *determination view*, an especially prominent way of attempting to spell out the dual commitments of irreducible thickness and the traditional view of evaluation. Very roughly, the determination view holds the thick and thin relate as determinate to determinable—applying thick concepts are just determinate ways of applying more abstract (determinable) thin concepts. My criticism will center around Edward Harcourt and Alan Thomas' (2013) articulation of the view.⁶ As a start, I show their defense of the determination

¹ Throughout this paper, I will use capitalized words (e.g. COURAGE and cognates) to indicate a concept; for kinds, I will italicize the word (e.g., *red* and cognates).

² Though, for a view that denies thick concepts are evaluative in the way required by both evaluative reductionists and defenders of irreducible thickness, see Väyrynen (2013).

³ I will sometimes refer to defenders of irreducible thickness as antireductionists.

⁴ For a nice overview of the issues involved in irreducible thickness see Kirchin (2013a); For some suggestions in favor of irreducible thickness, see Roberts (2011; 2013b).

⁵ For views like the one I labeled “traditional,” see Harcourt and Thomas (2013). For views like the latter one, see Kirchin (2017) and Roberts (2011, 2013a).

⁶ Oddie (2005) and Tappolet (2004) also defend this view. Though I do not think they are antireductionists.

view leaves something to be desired; for all they have said, it may turn out that irreducible thickness is not well understood in terms of the determination relation at all. In order to show that irreducible thickness should be so-understood, they need to justify their commitment to the determination relation. To demonstrate this cannot be done, I will focus on an analysis of the determination relation proffered by Eric Funkhouser (2006, 2014). After seeing how the analysis is supposed to work, I will argue that the only available way to fill in the determination view according to the analysis fails for a significant reason—there are not enough resources in the evaluative domain to at once secure irreducible thickness and the traditional view of evaluation. In section three, I generalize the argument to another view (what I call the *composition view*) and draw out some general lessons for defenders of irreducible thickness.

2 The Determination View

In this section, I spell out the determination view in some detail. In order to show how we might intuitively understand the determination view, I contrast it with the genus-species model of evaluation, which nicely captures the evaluative reductionist line. Following this, I briefly note a few extra features of the determination view and then conclude with a general outlining of the key commitments incurred by those who defend the view.

The determination relation is meant to be a specification relation distinct from the *genus-species model* for specification. On the genus-species model, species are specifications of the genus along logically independent and prior criteria called *differentia* (Tappolet (2004) and Kirchin (2017): 44–46); accordingly, evaluative reductionists suppose that analyses of thick concepts will break into a genus (thin) evaluative concept and some (nonevaluative) *differentia* which distinguishes one species (thick) evaluative concept from another. For example, the evaluative reductionist would give the following analysis of COURAGE: COURAGE = GOOD + RISKING HARM FOR A GOOD END. As the analysis suggests, the genus-species model of specification invites the picture that thick concepts are “built” from prior thin evaluative concepts and from independent nonevaluative matters which make something good in particular circumstances.

In contrast, the determination relation represents a specification relation between more general and more specific things which does not entail reduction. A canonical instance of the determination relation is *colored-red*. Even though *redness* is a specification of *coloredness* and therefore entails that, anytime something is red, it is also colored, there is no analysis of *redness* in terms of “x is red iff x is colored and F.” In other words, *being such that it is red* necessarily entails *being such that it is colored*, but *being red* does not thereby reduce to *being colored*.⁷ In the same way, claim defenders of the determination view, thick concepts are the determinates of thin determinables. In our example, COURAGE does not specify GOOD according to some reductive analysis like the one given above; rather, COURAGE is *sui generis* and picks out a specific way some person (or whatever) can be good (Harcourt and Thomas (2013): 26–27).

There are important features of the determination relation which secure it as a view friendly to the line taken up by defenders of irreducible thickness. The first is that, in order for

⁷ There is some reason to think, however, that the reduction *does* work in the other direction. See Gillett and Rives (2005) and Schroer (2011) for discussions of how determinates explanatorily (causally, etc.) exclude their determinables. I think Harcourt and Thomas (2013: 26–27, 31) would accept the reduction in the other direction.

something to possess a determinable, it must possess a determinate. This just means that determinables are not logically prior to the determinates. Second, and in virtue of there being no analysis along the lines of “*x* is red iff *x* is colored and *F*,” there are no logically prior *differentia* we can appeal to in order to explain why the specification happens in just the way it does; in fact, appealing to the features which specify a determinate just is to appeal to the determinate. For instance, appealing to some brightness, saturation, and hue values to specify *redness* just is to appeal to *redness*—to be colored just is to be such as to satisfy any of those values and *being red* just is to satisfy a subset of those values.⁸ Evaluative concepts, according to Harcourt and Thomas (2013: 26–27),⁹ work in an importantly similar way:

[W]e agree with Judith Jarvis Thomson when she claims that [“it seems very plausible to think that a thing’s being good must consist in its being good in some way... if that is the case then there is no metaphysically mysterious property goodness.”] Compare: no object just is coloured without being some particular color...

... If, as on our view, to be good is always to be good in some particular way, in coming to know (by whatever epistemic route) that something is that particular way one also comes to know (by that route) that it is good.

Indeed, this commitment is just a combination of irreducible thickness with the traditional view of evaluation.

Defenders of the determination view hold that evaluative concepts are *nonevaluatively shapeless*; that is, the determination view (probably in accord with other views describable as antireductionist)¹⁰ hold that evaluative concepts do not have any unifying nonevaluative specifications on their extension. This commitment is evident in Harcourt and Thomas’ (2013: 33) assertion that, “If the disentangling argument is right, there will not be anything non-evaluative that unifies [a thick concept] (the extension will be, as Elstein and Hurka seem to agree, ‘shapeless’). So, any opinions that users of a thick term may have about what non-evaluatively unifies its extension will be false.” Their contention is that we cannot disentangle the components of a thick concept (if it even makes sense to talk of there being components in the first place)—there is no way in principle to “pull apart” a thick concept into its evaluative and nonevaluative components. Transparently, so the thought goes, if there is no specifiable nonevaluative core to a given thick concept (COURAGE, say), then COURAGE is disentangleable because there will be no specification of the nonevaluative application conditions of COURAGE that can be pulled apart from the evaluative component. If the evaluative is all there is, then there is nothing to analyze in COURAGE (or any other thick concept).

⁸ Of course, brightness, saturation, and hue may not be what color is. The only thing I wish to do here is have a convenient example; in pursuit of that, I am following Funkhouser (2006, 2014) in stipulating that color just consists in brightness, saturation, and hue.

⁹ Harcourt and Thomas assume there is enough symmetry between color properties and evaluative concepts to get the analogy between the color case and the evaluative case going. This can be demonstrated in their (2013: 24–25) argument for seeing the determination view as a serious contender to evaluative reductionism. In that passage, they move from talking about evaluative terms, to color properties, and then to evaluative concepts. As I will show later, the analogy actually needs to be really strong if they are going to push through the idea that the determination view is the correct view to take when it comes to irreducible thickness.

¹⁰ It does not seem as though Roberts (2011, 2013a) and Kirchin (2013b, 2017) deny shapelessness.

From the above discussion, we can extract the following general picture for the determination view. First, the determination view holds that thick concepts are evaluative as a matter of being a determinate of some thinly evaluative determinable; for instance, COURAGE is an evaluative concept in virtue of being a specific way someone can possess the concept of GOOD.¹¹ Second, though thick concepts are evaluative in virtue of specifying GOOD or BAD, they are not thereby reducible to them. Third, the determination view is committed to nonevaluative shapelessness. It is these three commitments, I think, that are central to the determination view.

3 Analysis and Determination

In the last section, we saw how the determination view proposes to give flesh to the idea that thick concepts are irreducibly thick. Relying on an intuitive appeal to the relation between *being red* and *being colored*, the determination view proposed that we see the relation between the thick and the thin as one of determinate to determinable. If they are correct, then we have a plausible way of spelling out how thick concepts are inherently evaluative while also being irreducible—they are evaluative as a matter of being good (or bad) in a particular way, where the specification is not spelt out in terms of logically prior nonevaluative criteria in virtue of which something is made good (or bad) in that particular way. What I want to do in this section is show that the determination view has not done enough to prove that evaluative concepts are well understood in terms of the determination relation; further, I will argue that the only available way to answer my challenge is unworkable for them in principle.

The contention I am making starts with the thought that there being some intuitive sense to the idea that evaluative concepts are an instance of the determination relation does not just entail that evaluative concepts in fact are such an instance. First, consider the three truisms Harcourt and Thomas (2013: 25–26) pull from Funkhouser’s (2006: 548–549):

- (1) The determination relation holds between pairs of concepts. So, BRUTAL is a determinable relative to THUGGISH but a determinate relative to BAD.¹²
- (2) Any determinate concept is a specification of a determinable concept, where that specification is distinct from the genus-species model.
- (3) Any determinate concept entails all the determinable concepts it falls under. So, tokening THUGGISH entails BRUTAL and BAD.

At the stage in which Harcourt and Thomas introduce these truisms, there is no argument for thinking that the determination relation is *in fact* the correct view of evaluative concept relations. The clencher here is (2)—what reason do we have for supposing the specification relation between thin and thick is distinct from the specification adverted to by the species-genus model? The species-genus model, at least if we are thinking about (1) and (3) is at least as initially intuitive as the determination view—it is not, so I say, a completely crazy idea. But if that is right, then we need some reason, in addition to a few initially plausible intuitions

¹¹ This characterization is equivalent to the one I give below—viz., that to be a thick evaluative concept just is to satisfy some values of the determination dimensions of a thin evaluative concept. This turns out to be a crucial (and crucially damning) commitment for the determination view, as will be made clear below.

¹² In the original passage, Harcourt and Thomas talk about properties; but, since their usage of ‘property’ and ‘concept’ is fairly loose throughout, I have taken the liberty of changing their talk of properties to talk of concepts.

about how the evaluative case can preserve some (cherry-picked) truisms,¹³ to think that evaluative concepts are an instance of the determination relation—what we need is some way to prove that the evaluative are in fact an instance of the determination relation.

Of course, Harcourt and Thomas (2013: 29–36) *do* try to prove such a thing by targeting Elstein and Hurka's (2009) view, which is essentially a species-genus model of evaluative specification. Recall Harcourt and Thomas' commitment to nonevaluative shapelessness; in their reply to Elstein and Hurka, Harcourt and Thomas attempt to leverage shapelessness against them. The thought seems to be that, if shapelessness is true, then it seems like some view along the lines of the determination view is correct since there is just no nonevaluative specifications we could make on GOOD in order to individuate thick concepts. In other words, if shapelessness is right, according to the determination view, then the (evaluative reductionist) genus-species model must be false. What more could you want?

There are two general points to make to this objection. The first is that shapelessness, all on its own, does not guarantee the falsity of the genus-species model. In principle, the evaluative reductionist can offer the following analysis of COURAGE which respects shapelessness:

x is the concept, COURAGE, [iff] x is of nonevaluative sort T (specified) and x has some components (unspecified) in virtue of being of nonevaluative sort T and is good in virtue of having those components where there is no nonevaluative shape to those components across a range of instances of the concept's correct application.¹⁴

This analysis would respect the genus-species model of reduction as well as respect shapelessness by maintaining that every occurrence of COURAGE is of a general nonevaluative sort—say, the nonevaluative sort of requiring action in light of danger and (or?) fear—while also making it a matter of (perhaps extreme) context dependence whether the particular components in question belong to the specified nonevaluative sort. If the evaluative reductionist can make adequate sense of shapelessness, then the mere fact that evaluative concepts seem intuitively to fit (1) and (3) would do nothing for the determination view because they fit the genus-species model just as well (depending, anyway, on your prior metaethical commitments). If that is right, then the additional argument that the evaluative are best made sense of in terms of the determination relation is needed, and defenders of the determination view have not done this.

The second point to make here is that, even if shapelessness did ultimately decide matters in favor of the antireductionist, it would still not be obvious that the determination relation is the right answer to give. This is because making sense of the shapelessness hypothesis does not rely on the evaluative being an instance of the determination relation. To see this, consider Harcourt and Thomas' (2013: 31) defense of the disentanglability argument, the crux of which can be captured by this quote: “[W]hat we have in judgments of courage here is not evidence for the priority of the thin, but one point of entry into an inescapable holism of the thick.” Harcourt and Thomas intend for this line about the holism of the thin to support the determination relation (and hence (2)); indeed,

¹³ Funkhouser lists eight truisms in his (2006) and nine in his (2014: 33–34). Examination of those truisms shows that Funkhouser thinks the truisms set criteria for an adequate analysis of the determination relation rather than conditions on instances of the determination relation, as Harcourt and Thomas seem to think. The conditions might well provide some guidance for particular putative instance of the determination relation, but it would only do so indirectly by first confirming the analysis that admits the putative instance of the determination relation. See especially Funkhouser (2006: 562ff.) for an example of how this might go.

¹⁴ This analysis is borrowed from Roberts (2011: 510). I have change some aspects of the example, however; for instance, whereas Roberts' original analysis was for predicates, I have adapted it for concepts and, further, I have changed the relevant example to fit my use of COURAGE throughout.

making adequate sense of the holism of the thick *can* be done by positing the determination view since, plausibly, the determination relation will make the thick prior to the thin. But, one *could also* make sense of the holism of the thick by positing a no-priority view according to which neither the thick nor the thin are conceptually prior; such a view could make sense of holism by positing that, in any given application of an evaluative concept, there will be other evaluative concepts in play grounding the application of the concept to an individual (say, by predicating *courageousness* to a person).¹⁵ But, if that is right, then nonevaluative shapelessness does not clench the determination view as the only view able to fill in the outlines of irreducible thickness. Even granting the shapelessness line clenches the case for irreducible thickness, then, we need more argument for thinking irreducible thickness should be spelt out as the determination view.

Thus, the determination view, just in virtue of showing the evaluative *can* be made sense of in terms of determination, does not thereby justify the claim that it *is* an instance of the determination relation; therefore, the determination view has not justified adherence to the traditional view of evaluation and has therefore not justified its preferred way of thinking about evaluative concept relations (as specification via determination). What the determination view must do is show how the evaluative case fits a general analysis of the determination relation. If it can do this, then it will have done enough to show that the thick, if irreducible, are an instance of the determination relation.¹⁶

3.1 The Determination Relation and Color

The objection to the determination view so far can be summed up as follows: Even supposing Harcourt and Thomas have done enough to show that we should take the determination view seriously, they have not yet shown that the determination view is correct. The next step is to try and provide the missing justification for the determination relation. One obvious way they can prove their view is the correct one is to show how the evaluative case fits with a general analysis of the determination relation (this is even suggested by (2)). The aim of the next two subsections is to show that the determination view cannot be correct—there is no plausible story to tell for determination relation. I will show this by explicating Funkhouser's (2006, 2014) treatment of the determination relation, and then offer a counterexample to the determination view in light of this analysis.

Funkhouser's analysis of the determination relation can be given in brief form:

(D) Some B determines some A iff:

- (1) A and B have the same determination dimensions,
- (2) B has the absolute non-determinable necessities of A, and
- (3) the range of determination dimension values for B is a proper subset of the range of determination dimension values for A.¹⁷

Where 'determination dimensions' are the features which specify *redness* (etc.), and 'absolute non-determinable necessities' are those features *redness* and *coloredness* must have in order to be an instance of the kind in question, but nevertheless do not play a part in determination. Crucially, the *absolute* non-determinable necessities *do not vary* along with the determination dimensions.

It seems clear to me this analysis makes good sense of how *being red* (or whatever) relates to *being colored*. Suppose that, for something to be red, that thing must have a certain range of values

¹⁵ See Kirchin (2017: 72–79, 139–142) for relevant discussions here.

¹⁶ I am grateful to an anonymous reviewer for pressing me in ways that sharpened this part of my argument.

¹⁷ Cf. Funkhouser's (2014: 39) original statement of the analysis.

with respect to brightness, saturation, and hue which jointly constitute the determination dimensions of *redness* (not to mention any number of the other various color kinds) and *coloredness*; thus, condition (1) is satisfied. Condition (2) is also satisfied since, it seems, neither *redness* nor *coloredness* have any non-determinable necessities. Finally, condition (3) is met in virtue of the determination dimension values for *redness* being a proper subset of the determination dimension values for *being colored* since it is the case that the determination dimension values for *colored* also contain the determination dimension values for *green*, *yellow*, *orange*, and so on. Importantly, since the *red-color* case is a canonical example of the determination relation, I think we have good reason to suppose Funkhouser's analysis is correct.¹⁸

3.2 Determination and Evaluative Concepts

The discussion above shows us (D) is a good analysis of the determination relation and therefore provides hope for the determination view to secure (2). It should therefore be useful for seeing whether the determination view can provide the missing justification I claim they need. My contention is that the only likely way for the determination view to answer my challenge is via appeal to purely evaluative determination dimensions since to appeal to nonevaluative determination dimensions would be to make the extension of thick evaluative concepts a matter of nonevaluative specificity—a clear violation of shapelessness and the thesis of irreducible thickness.

So, can the determination view meet my challenge with appeal to purely evaluative determination dimensions? To get a start, consider Pekka Väyrynen's (Unpublished) suggestion that we look to an evaluative concept's flavor, valence, and strength. *Flavor* refers the sort of evaluative concept it is; i.e., when we consider what flavor a concept has, we are asking whether it is a moral evaluative concept, a prudential evaluative concept, or something along these lines. *Valence* refers, unsurprisingly, to whether a given evaluative concept is positive or negative. Finally, *strength* refers to the kind of reasons entailed by an evaluative concept. We might think of these reasons being divided into contributory reasons and overall reasons—only the latter being decisive reasons for action. In addition to these three dimensions, I propose to add two more. The first is *focus*; when we inquire about the focus of a given evaluative concept, we are asking whether the concept in question is relational-focused or individual-focused. An example of the former kind would be GENEROUS while an example of the latter kind would be TEMPERATE. The second addition is *motivation*. The motivation dimension will determine whether the concept in question picks out an internal or consequential motivation; for example, one might think TEMPERATE picks out a trait or property motivated by the consequences it brings about while HONEST picks out a trait or property motivated internally.

Graham Oddie (2005: 162) can be construed as suggesting that we include the criterion of *intensity*, which would contain values ranging from 'mild' to 'extreme', and which would refer to the amount of (dis)value ascribed by an appropriate application of a thick concept.¹⁹ However, the amount of (dis)value denoted by the application of a thick concept is likely to be highly context responsive. Since we are interested in individuating concept types (rather than tokens), I take it the

¹⁸ See Funkhouser (2014) for a fuller defense of this contention.

¹⁹ From other work (e.g., Oddie (2018)), I suspect that Oddie is an evaluative reductionist (though, I wonder if this is a stable position), and so I do not want to hold him to actually advocating this as a criterion for individuation. In fact, I think the context of Oddie's discussion (and other parts of his (2005): e.g., 207) of thick concepts lends itself to reading him as suggesting (probably) purely descriptive individuating criteria (but, possibly, he could be suggesting mixed descriptive/evaluative criteria). Since the purely descriptive option is incompatible with irreducible thickness, I will not consider it in the main text.

context responsiveness of intensity would make it basically useless for individuating concepts and we should therefore not include it in the determination dimensions for evaluative concepts.

With the exception of intensity, then, I think the above dimensions are initially appealing and therefore give the determination view some initial hope of meeting my challenge. To see how, consider the structure of determinates and determinables, starting at what I will assume is the highest level of abstraction. So, PRO, we might think, is the highest-level determinable we have as an evaluative concept.²⁰ A determinate of PRO would be GOOD (in addition to PRUDENT and PLEASING). The determination dimensions of GOOD would be ‘moral’, ‘positive’, and be neutral with respect to strength, focus, and motivation. A determinate of GOOD would be GENEROUS and would have the following dimensions: ‘moral’, ‘positive’, ‘contributory’, ‘relational’, and ‘internal’. Further, we can suppose the non-determinable necessity each evaluative concept has is something like APPLIES TO CERTAIN NONEVALUTIVE CIRCUMSTANCES.²¹ So, intuitively, there seems to be a way to fill out the analysis in purely evaluative terms.²²

Despite its initial appeal, however, I think this option fails to distinguish thick concepts we should maintain are in fact distinct. Looking back at the example, it appears that many determinates of GOOD would fall along the same subset of determination dimension values. To see the problem with this, consider the determination dimensions for COURAGE and HONEST. The determination dimensions for COURAGE would be ‘moral’, ‘positive’, ‘contributory’, neutral with respect to ‘focus’, and ‘internal’. HONEST has exactly the same

²⁰ An anonymous reviewer pointed out that some might think of ‘pro’ as an attitude. While I take the point, I think the objection does no real damage to the example here. Suppose Sarah thinks, ‘I am pro women’s rights.’ Assuming a cognitivist understanding of evaluative discourse, we can construe Sarah as saying that she ascribes some positive evaluative status on women’s rights, where the ascription of such a status is construed as attributing (via predication) some property to the state of affairs in which women are treated equally. But, plausibly, predicates (again, given a cognitivist understanding of such things) are filled in by our concepts. So, thinking of pro stances as (noncognitive?) attitudes might just mean that someone does not go in for cognitivism. However, and as the present context assumes, if we go in for cognitivism, then PRO is a very general concept about positive evaluative status. Perhaps we never do use PRO without using some more determinate concept (e.g., GOOD), but that is all to the good on the view we are currently considering.

There could be a complication here about how we would then make sense of noncognitive states like desires and intentions as being an instance of *being pro something*; for example, intending to travel to Scotland might be construed (inadequately) as ‘being pro travelling to Scotland’. My general feeling of this, however, is that we could have the concept, PRO, which included both potential (cognitivist) predications and (noncognitivist) expressions of desire-like states. A rendering of ‘pro’ which makes it a concept does not preclude this possibility.

²¹ There may be a question here of whether this could work since there is supposed to be only an evaluative component. But I think that would be a mistake. That a concept is essentially evaluative does not entail that it has no nonevaluative application conditions. As an example, even on an evaluative reductionist understanding of evaluation, no one ever thought thin (purely evaluative) concepts like GOOD utterly lacked nonevaluative application conditions.

²² Oddie (2005: ch. 6, 203–206) conceives of thick concepts as determinates along convex regions in a quality space (what others call a conceptual space), where a set of concepts are convex (and so mark out a genuine conceptual domain) just in case they satisfy *betweenness*—any concept falling between two concepts in the collection is itself part of the collection. It seems to me that the idea of a quality space which is marked out by convexity is no different from the sum of the determination dimensions for a determinable (what I would dub a determination space). This can be seen by noticing that quality spaces are, like determination spaces, plottable on multidimensional planes (or single-dimensional planes, if it turned out thick concepts only had one determination dimension). At any rate, I think what I say here will apply equally to how (a suitably bastardized—cf. n. 19) Oddie spells out his own position. This is because, even if betweenness does tell us something about the presence of a conceptual domain, it presupposes that we can meaningfully demarcate the various concepts within the domain (this is confirmed by his comments on p. 204). Thus, my challenges can be seen as demanding (in addition to Harcourt and Thomas) that (some version of) Oddie tell us more about how the concepts relate to one another than merely that they satisfy betweenness. I should say: My criticisms will not apply to his use of the determination relation as it is utilized against naturalistic reductions of value.

determination dimensions as COURAGE. The difficulty here is that our concepts, HONEST and COURAGE, when applied to various situations, do not always overlap. It is easy enough to identify situations in which COURAGE is the appropriate concept to apply, but HONEST is not—say, when a soldier runs into battle despite fear for life and limb. Conversely, we can describe situations in which HONEST is the appropriate concept to apply but COURAGE is not; for example, we might think one should be honest to tourists when they ask for directions, but we should not think such an action takes courage. Here, it seems, we have a clear counterexample to the determination view spelled out in this way.

The issue here is not so much that defenders of the determination view cannot fill in their view such that it coheres with (D); this much, it seems clear, they certainly can do. The issue, rather, is that the determination view can only get this result at the expense of collapsing obviously disparate evaluative concepts into the same concept. Since, on the determination view, being an evaluative concept just is to possess certain “values” for flavor, valence, strength, focus, and motivation, concepts that have the same values for these determination dimensions *in fact are the very same concept*; thus, as the counterexample shows, HONEST and COURAGE turn out to be the same concept. Since that is obviously not the case, I conclude that, even if the determination view does fill in the general analysis in some *prima facie* plausible way, it generates what is to my mind a *reductio* for the determination view. There is, it seems, not enough structure. Since this is the only option available, I conclude that the determination view is false.

There is a potential hiccup at this point, however. A defender of the determination view could try and salvage it *qua* a theory amenable to irreducible thickness by proposing there are in fact determination dimension values for a given thick concept, and these determination dimensions are not at odds with shapelessness. This is so because the determination dimension values are themselves thick evaluations.²³ The suggestion, as best I can tell, is meant to salvage the determination relation between GOOD and the thick concepts we intuitively appeal to (e.g., COURAGE, JUSTICE, GENEROSITY) by adding a layer of complexity I claim the determination view does not have.

My worries with this suggestion are connected. The first worry is that I am not totally sure how this gets spelt out in a way that seems plausible and friendly the determination view’s main tenets. This uncertainty derives from a deeper worry: This fix leaves the determination relation with respect to evaluative concepts unmotivated. To address the deeper worry first, it seems to me that, whatever the thick determination dimensions are, they would themselves need to be either determinates of GOOD or not. If they are, then they cannot be determination dimensions, since this would violate the structure of the determination relation. To see the problem, recall that what it is to be an evaluative concept, on the determination view, just is to satisfy certain values for a range of determination dimensions. But, then, the determination dimensions are not themselves evaluative concepts since satisfying any one dimension (which, perhaps, determination dimensions trivially do) does not confer the status of being an evaluative concept. If that is right, however, we would have to give a different analysis of the essences of those thick evaluations serving as determination dimensions—what it is to be those thick evaluative concepts would be something different from what it is to be a thick evaluative concept that has its evaluative nature in virtue of being a determinate.

²³ Something like this objection was presented to me by D. Gene Witmer in conversation. It is possible Oddie (2005: 162) has suggested something like this in the past. Witness his use of ‘fun’ and ‘frustrating’. These are not obviously purely descriptive—though, see my n. 19.

So, let us suppose that thick evaluative determination dimensions are not themselves determinates of GOOD and can therefore receive a separate analysis of what it is to be an evaluative concept; once a defender of irreducible thickness admits this is possible, however, then she has picked out another way in which evaluative concepts necessarily encode evaluation. This is so because, according to the determination view, being a determinate of GOOD is constitutive of what it is for a given concept to be evaluative.²⁴ If the thick determination dimensions are outside of the determination relation, then there must be some other way in which evaluative concepts can be evaluative. But once we have picked out a different way the thick are essentially evaluative—say, by being evaluative in their own right—then positing the determination relation would be unmotivated. I think this is the root of the problem for why there seems to be no plausible and friendly way to spell out the suggestion of a couple paragraphs back. So, it seems, this will not do as a fix for the determination view. My conclusion about the falsity of the determination view stands.²⁵

4 Generalizing the Problem

The last section showed us the determination view lacked the resources to fill in the determination relation whilst holding on to irreducible thickness. The trouble for this view, it seems, comes from depending on a traditional view of evaluation paired with the idea that evaluative concepts—whether thick or thin—are purely evaluative. Given these two commitments, there just is no credible way to fill in the general analysis of the determination relation which does not get us some deeply implausible results; this, in my mind, amounts to a *reductio* for the view. In this section, I will consider a test case for generalizing my criticism: the composition view.²⁶ This view maintains that the thick are to the thin as parts are to wholes. My criticism here will be the same as it was for the determination view: there just is no plausible way to fill in the analysis which remains faithful to irreducible thickness and the traditional view of evaluation. What will emerge, if I am right, is that any view purporting to maintain a traditional view of evaluation while also being antireductionist *will have to* answer the challenges made in this paper (perhaps suitably adjusted to fit the particular case).

4.1 The Test View

The composition view holds that GOOD is composed of thick concepts. To get started, it will be helpful to have a few (evaluative) compositional principles in place. For our purposes, and allowing the variables to stand in for evaluative concepts,²⁷ the following will do:²⁸

²⁴ See Harcourt and Thomas (2013: 27): “If, as on our view, to be good is always to be good in some particular way, in coming to know... that something is that particular way one also comes to know... that it is good.”

²⁵ Thanks to an anonymous reviewer for pressing me to develop my response to this objection in more detail.

²⁶ No one I know of has defended this view of thick concepts. However, it is, I think, a natural view to take with respect to irreducible thickness given a compositional mereology, so I plan to give it a good read. Much of what I say in developing the composition view is adapted from Varzi (2019).

²⁷ I take it parthood satisfies a transitivity constraint so that it is open for thick concepts to have parts (and proper parts) as well.

²⁸ I include (EQ) for the purposes of assessing whether any thick concepts which should not satisfy (EQ) in fact do. In other words, (EQ) will help us in discussing a possible counterexample to the view.

$$(EQ) \quad EQ_{xy} =_{df} (P_{xy} \& P_{yx}).$$

$$(PP) \quad PP_{xy} =_{df} (P_{xy} \& \neg(x = y)).$$

$$(O) \quad O_{xy} =_{df} \exists z(P_{zx} \& P_{zy}).$$

On the composition view, then, to be a thick concept is just to satisfy (PP) with respect to some (or other, or multiple) thin concept(s). Overlap (O) is necessary as a way of making sense of evaluations which seem to include similar considerations. It also could allow for defenders of the traditional view of evaluation to allow for thick evaluative concepts to be evaluatively flexible—GOOD and BAD overlap at certain points in our conceptual economy (say, at certain points in our concept of TRUSTING or HARSH).

The above thoughts on (O) seem at least coherent, but what we need are plausible individuating criterion which can help us make good on the idea that thick concepts relate to thin concepts in just the way parts relate to wholes.²⁹ Seeing if we can individuate concepts according to the purely evaluative individuating criterion we considered for the determination view immediately suggests itself. On that option, evaluative concepts are individuated by satisfying certain values for valence, strength, flavour, motivation, and focus. What would need to be the case in order for (O) to deliver on the possibility of evaluative flexibility would be for some thick concepts to be neutral on valence—COURAGE, say, would need to be neutral with respect to positive and negative valences. If it makes sense to say of COURAGE (or some other concept) that it is neither definitely positive nor definitely negative, then there could be hope for overcoming some suitably adjusted counter example to the determination view based on these individuating criteria.

But does it make much sense to say COURAGE is neither definitely positive nor definitely negative on a traditional view of evaluation? Recall the traditional view of evaluation says that to be evaluative just is to be good or bad in some way; if evaluation just consists in being good or bad in some way, then it seems natural for us to pick a particular valence and say that a concept is evaluative in virtue of encoding that valence in a particular way. The normal way of doing this for COURAGE is to say that it is good-in-a-way, but, it seems to me, this way of explicating the evaluative nature of COURAGE is consistent with saying that courage is also bad-in-a-way. But even if that sounds odd, we might just abstract how we explain evaluative concepts by saying that COURAGE is evaluative in virtue of being a part of good or bad, where the 'or' is read inclusively.³⁰ It is not obvious to me why this would not work and I propose we grant the idea in order to assess the composition view of evaluation in what might be thought to be its strongest form.

There is some enhanced initial plausibility for the composition view brought about by the considerations on (O). If it is possible to assign a *neutral* value to valence, then we might be able to make good on an irreducible thickness position that holds on to the traditional view of evaluation. To test this, we can think again about the problem concepts for the determination

²⁹ Identical considerations that drove us to considering a general analysis of the determination relation would lead us to exploring individuating criterion for thick parts to thin wholes.

³⁰ It seems to me that this should be perfectly acceptable. Given thick concepts do in fact sometimes get used to evaluate actions (say) sometimes as good and sometimes as bad, then an overlap principle that requires us to see evaluation as a matter of disjunctively encoding thin evaluations would be more-or-less required.

view: COURAGE and HONEST. Recall that the values for COURAGE and HONEST were ‘moral’ (flavour), ‘positive’ (valence), ‘contributory’ (strength), neutral (focus), and ‘internal’ (motivation), and so what needs to be the case with either COURAGE or HONEST (*but not both*) is that the value with respect to valence is ‘neutral’ rather than ‘positive’.

The trouble for the suggestion currently on offer is that both COURAGE and HONEST are plausibly neutral with respect to valence. To see this, consider two sets of utterances a competent speaker of English might utter. The first witnesses the concepts’ application in a clearly positive evaluative manner: “Your courage in that circumstance was heroic,” and, “That you were honest about your mistake is admirable.” Other positive evaluations—‘heroic’ and ‘admirable’—seem to show that appropriate use of these concepts include positively evaluating something. This is perhaps no surprise since, after all, we tend to think of things that are courageous or honest as being good in some way.

Perhaps more surprising (and controversial), is that the following two utterances convey a negative overall evaluation, but do not for that reason seem to involve us in a conceptual mistake with respect to COURAGE and HONEST: “That was courageous, but stupid,” and, “It would be bad if you were honest with the Nazis about where those Jews are hiding.” The first locution, in virtue of including the evaluation that the action (or whatever) was stupid, seems to indicate that ‘courage’ is being used to convey an action that would not be recommended or deemed good; similarly, the second utterance explicitly states that a policy of honesty would be a positively bad thing if that policy was utilized in telling the Nazis the truth about where some Jews were hiding out. I think we should recognize these evaluations as genuine evaluations featuring the use of the concepts COURAGE and HONEST rather than thinking of them as involving some different concept (say, COURAGE-negative and HONEST-negative). But if that is right, then it seems to me that both of the concept types would be neutral with respect to evaluation.³¹

The crucial move here is to notice that, if an evaluative concept can in principle be used to convey both an evaluative and a negative overall evaluation, then that concept would be neutral with respect to valence. If that is right, however, then it seems to me that HONEST and COURAGE would satisfy (EQ); they are equal because, it seems, they both would possess one another as a part in virtue of having the exact same values with respect to the individuation criteria. But if they are equal, then they are identical to one another.³² Thus, the composition view seems to predict exact overlap between concepts we should think are distinct just like the determination view did.

4.2 A General Failure

This concludes my assessment of the composition view. Are there other criteria by which we could individuate thick concepts as proper parts of some (or multiple) thin evaluation(s)? I am not totally sure there is, but there are a few necessary conditions on an adequate attempt. The first thing that will be necessary in choosing the individuating criteria is that they not be purely descriptive since that would plausibly offend the heart of irreducible thickness. A second possibly necessary condition on selecting criteria for individuation would be that they allow for evaluative flexibility. This is suggested by the discussion of COURAGE and HONEST being neutral with respect to the valence they convey. In this respect, the determination view is

³¹ The brief discussion of the last two paragraphs is inspired by Kirchin (2017: 49–53).

³² They would be identical if equal if the parthood relation is antisymmetrical, which I think is uncontroversial.

doubly doomed (if I am right) since it cannot allow for the sort of evaluative flexibility I say is possible.³³ A third condition on the individuating criteria would be that the criteria need to keep obviously disparate concepts (like COURAGE and HONEST) distinct.³⁴

A final necessary condition on adequate individuating criteria would be avoiding triviality. Though obvious, this is crucially important and explains why just saying courage is a way of being good and honesty is a different way of being good is not sufficient for describing the relationship between thin and thick concepts. To see this, notice that evaluative reductionists also say courage is a particular way of being good, yet they conceive of the relation between goodness and courage as being wildly different from the way defenders of irreducible thickness (of any variety) say it is. Notice also that defenders of irreducible thickness who think thick concepts are evaluative in their own right (where their evaluative nature is not a matter of encoding a thin evaluation) can happily say that all good is good in some particular way. So, if the criteria for individuation entail that just about anyone can accept the criteria, then those criteria will not be adequate for proving the hypothesis that irreducible thickness is compatible with adopting the traditional view of evaluation.

Is there a unified explanation of the general failure of views which seek to pair irreducible thickness with a traditional view of evaluation? My suggestion is this: The reason that none of these views work is because the evaluative domain does not have enough structure to (a) secure the required tight connection between thick concepts and thin concepts and (b) ensure that we carve up our conceptual economy in a way that at least preserves the distinctness of obviously discrete concepts. I do not think it is possible for the antireductionists to do both of these things. This, I claim, will be a quite general failure, and so any view which attempts to meet my challenge will have to do it by capturing (a) and (b).

This suggests a general formula for criticizing antireductionist traditional views of evaluation. Start with the purported relation between the thick and the thin; then, given the arguments provided in support of the *irreducible thickness* portion of the view do not fully determine the conclusion that we should adopt a traditional view of evaluation in addition, challenge defenders of the view to provide an individuating criteria for thick concepts that at once meet the challenges represented in (a) and (b). The conjecture of this paper, supported by the evidence of the last two sections, is that this cannot be done.

5 Conclusion

There are, then, two things a defender of irreducible thickness who also wants to hold on to the traditional view of evaluation must do. The first thing is to try and better justify their adherence to the traditional view of evaluation. One way of doing this is to provide a non-trivial connection between thin and thick concepts that does not make obviously disparate concepts identical. If I have been right, this is not likely to be doable, though I admit there could be relations or individuating criteria I have not thought of. Another way of doing this is to come up with a different reason for thinking the traditional view of evaluation is correct, and then conjoining this with the shapelessness hypothesis to fully support their view. This just means there is, ultimately, space to prove me wrong—and I wholeheartedly invite challengers to my conclusion that the pairing of views considered in this paper is not possible.

³³ See Väyrynen (2019) for discussion on this point.

³⁴ Some revision will be expected given an illuminating analysis. But there is a limit to acceptable revision.

There is, of course, another response we can make to the above discussion—drop the traditional view of evaluation. Those of us who are friendly to the idea of irreducible thickness, but find the arguments like the ones I have levelled here compelling, should look elsewhere for an adequate account of evaluation—one that abandons the need to tightly link thin and thick concepts. This route has been taken by Debbie Roberts (2011, 2013a, 2017, 2018) and Simon Kirchin (2013b, 2017). I think their suggestions are interesting and roughly on the right track, though I admit that there is a lot to assess and further work will need to be done before we have a satisfactory account on our hands; for instance, Roberts' notion of *essence* and its relation to necessity needs to be evaluated, as does Kirchin's commitment to thick concepts having unified extensions along the lines of Wittgensteinian family resemblance. Finally, I think we will need to have a hard look at how thick concepts incorporate descriptive content without having it as an additional component. To some degree, this last is dependent on the answer we give to the questions about Kirchin's and Roberts' ways of unifying the extension of a thick concept. But all this is too much for now, and therefore must be the topic of further work.

Acknowledgments Thank you to Simon Kirchin and D. Gene Witmer for their many helpful comments on different drafts of this paper. Thanks also to those at the Florida Philosophical Association and the Southeast Philosophy Congress for their helpful comments and questions. Finally, thank you to my anonymous reviewers for their feedback.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Blackburn S (1992) Morality and Thick Concepts: Through Thick and Thin. In *Proceedings of the Aristotelian Society*, suppl. vol. 66. 285–299
- Blackburn S (2013) Disentangling Disentangling. In *Thick Concepts*. Edited by Simon Kirchin. 121–135
- Elstein D, Hurka T (2009) From thick to thin: two moral reduction plans. *Canadian Journal of Philosophy*, 39 no 4(Dec.):515–536
- Funkhouser E (2006) The Determinable-Determinate Relation. *Nous* 40(3):548–569
- Funkhouser E (2014) The logical structure of kinds. Oxford University Press, Oxford
- Gibbard A (1992) Morality and Thick Concepts. In *Proceedings of the Aristotelian Society*, suppl. vol. 66. 267–283
- Gillett C, Rives B (2005) The Non-Existence of Determinables: Or, a World of Absolute Determinates as Default Hypothesis. *Nous* 39(3):483–504
- Harcourt E, Thomas A (2013) Thick Concepts, Analysis, and Reductionism. In *Thick Concepts*. Edited by Simon Kirchin. 20–43
- Kirchin S (ed) (2013a) Thick concepts. Oxford University Press, Oxford
- Kirchin S (2013b) Thick Concepts and Thick Descriptions. In *Thick Concepts*. Edited by Simon Kirchin. 60–77
- Kirchin S (2017) Thick evaluation. Oxford University Press, Oxford
- Oddie G (2005) Value, reality, & desire. Oxford University Press, Oxford
- Oddie G (2018) Non-naturalist moral realism, autonomy and entanglement. *Topoi* 37:607–620
- Roberts D (2011) Shapelessness and the thick. *Ethics* 121(April):489–520

- Roberts D (2013a) It's Evaluation, Only Thicker. In *Thick Concepts*. Edited by Simon Kirchin. 79–80
- Roberts D (2013b) Thick Concepts. *Philosophy Compass* 8(8):677–688
- Roberts D (2017) Depending on the Thick. In *Aristotelian Society Supplementary Volume XCI*. Oxford: Oxford University Press. 197–220
- Roberts D (2018) Why believe in normative Supervenience? in *Oxford Studies in Metaethics*, vol. 13. Edited by Russ Shafer-Landau. Oxford: Oxford University Press. 1–24
- Schroer R (2011) Can determinable properties earn their keep? *Synthese* 183:229–247
- Tappolet C (2004) Through Thick and Thin: good and its determinates. *Dialectica*, 58 no. 2: 207–221
- Varzi A (2019) Mereology. In *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. (Spring): URL= <https://plato.stanford.edu/entries/mereology/#CorePri>
- Väyrynen P (2013) *The lewd, the rude and the nasty: a study of thick concepts in ethics*. Oxford University Press, Oxford
- Väyrynen P (2019) Thick ethical concepts. In *The Stanford Encyclopedia of Philosophy* (Sum.): URL= <https://plato.stanford.edu/entries/thick-ethical-concepts/#InseView>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.