

promoting access to White Rose research papers



Universities of Leeds, Sheffield and York
<http://eprints.whiterose.ac.uk/>

This is an author produced version of a paper published in **Noûs**.

White Rose Research Online URL for this paper:

<http://eprints.whiterose.ac.uk/73423/>

Paper:

Vayrynen, PP (2012) *Shapelessness in Context*. Noûs.

<http://dx.doi.org/10.1111/j.1468-0068.2012.00877.x>

Shapelessness in Context*

Pekka Väyrynen

University of Leeds

(Forthcoming in *Noûs*, penultimate draft)

Abstract

Many philosophers believe that the extensions of evaluative terms and concepts aren't unified under non-evaluative similarity relations and that this "shapelessness thesis" (ST) has significant metaethical implications regarding non-cognitivism, ethical naturalism, moral particularism, thick concepts and more. ST is typically offered as an explanation of why evaluative classifications appear to "outrun" classifications specifiable in independently intelligible non-evaluative terms. This paper argues that both ST and the outrunning point used to motivate it can be explained on the basis of more general factors that have nothing in particular to do with being evaluative. Insofar as ST is plausible, a wide variety of non-evaluative terms will also be such that the extension of a term T isn't unified under similarity relations specifiable in purely T -free terms. If so, there is no reason to expect ST to carry the sorts of metaethical implications that get attributed to it. I also show that my main argument is robust across certain complications that are raised by the context-sensitivity of many evaluative terms but have so far been ignored in discussions of ST and related matters.

*I am grateful to helpful feedback from audiences at the 7th European Congress of Analytic Philosophy in Milan, the 2012 Meeting of the Pacific Division of the American Philosophical Association in Seattle, University of Nebraska, Lincoln, University of Leeds and Kings College London. Special thanks are due to Jamie Dreier, Matti Eklund, Ulrike Heuer, Simon Kirchin, Debbie Roberts, Preston Werner and anonymous referees for useful written comments or correspondence. Work on the paper was supported by the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 231016 and the Arts and Humanities Research Council (UK) under grant agreement no. AH/H038035.

1. Introduction

Many philosophers believe that the evaluative is “shapeless” with respect to the non-evaluative. What they mean by this shapelessness thesis is roughly that the extensions of evaluative predicates or concepts aren’t unified under non-evaluative similarity relations. This idea is recruited to do significant work in a number of metaethical contexts. The shapelessness thesis is supposed to undermine non-cognitivism.¹ It is supposed to undermine at least reductive naturalism in ethics.² It is supposed to illuminate the nature of such “thick” terms and concepts as *selfish* or *courageous*.³ More recently it is also claimed to imply moral particularism and the “holism” of reasons and value.⁴ Yet the shapelessness thesis itself went long without sustained attention and clarification – until recently.⁵ The main argument of this paper is that the notion of shapelessness operative in these discussions neither is characteristic of the evaluative in particular nor (partly for that reason) supports the sorts of significant metaethical implications that tend to get attributed to it. The phenomena that are used to support the non-evaluative shapelessness of the evaluative can be explained on the basis of factors that have nothing in particular to do with being evaluative and have no distinctive metaethical implications. After responding to an objection and explaining why my argument is robust across certain complications, I conclude by explaining how my main argument challenges certain claims about the implications of shapelessness regarding thick terms and concepts and bears on what notions count as evaluative in the first place.⁶

¹McDowell (1981), Williams (1985, 140-1), McNaughton (1988, 60-2), Hurley (1989, 13), Dancy (1993, 74; 1995) and Kirchin (2010).

²Dancy (1993, 76), Wiggins (1993) and Little (2000, 279).

³Williams (1985), Hurley (1989), Dancy (1995), Kirchin (2010) and Roberts (2011).

⁴Dancy (1993) and Roberts (2011).

⁵See especially Miller (2003, 244-54), Kirchin (2010) and Roberts (2011).

⁶I’ll mostly speak of terms and concepts in one breath; but see note 37 below for an important exception. It is common in metaethics to assume that the meanings of predicates are much the same thing as the contents of the concepts for which those predicates stand, where concepts are understood as something like senses or modes of representation as these notions are meant to contrast with reference. (Put more pithily, the assumption is that predicates *have* meanings

2. The Shapelessness Thesis

The idea that the evaluative is shapeless with respect to the non-evaluative is commonly credited to John McDowell and was given the label “shapelessness” by Simon Blackburn.⁷ The idea that the evaluative is non-evaluatively shapeless is commonly attributed to a passage in which McDowell argues that evaluative supervenience (the idea that there can be no evaluative difference without a non-evaluative difference) leaves open the following possibility:

[H]owever 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. Hence there need be no possibility of mastering, in a way that would enable one to go on to new cases, a term that is to function at the level supervened upon, but is to group together exactly the items to which competent users would apply the supervening term. Understanding why just those things belong together may essentially require understanding the supervening term. (McDowell 1981, 202.)⁸

In other words, supervenience doesn’t guarantee that the extensions of evaluative terms or concepts are “shapely” in the sense that they “correspond to kinds into which things can in principle be seen to fall independently of an evaluative outlook” (McDowell 1981, 216). The “shape” of a concept may here be understood as the real similarity between all and only the things that fall under it (Roberts 2011, 505). So what the shapelessness thesis says generically is that the extensions of evaluative terms and concepts have an evaluative shape but whereas concepts *are* such meanings. Such an assumption may be a significant simplification.

Following along with it will, however, be largely harmless here.

⁷See McDowell (1981, 202, 216) and Blackburn (1981, 167).

⁸I’ll leave to one side various complications about this passage, such as whether the notion of supervenience with which McDowell operates is supposed to be global or local, ontological or ascriptive and so on.

no non-evaluative shape. This is meant to be an in-principle point about evaluative classifications, not merely a point about contingent expressive limitations of our non-evaluative language. I'll assume that the point is consistent with the plausible idea that a "non-evaluative" relation is best characterized as one that can be ascribed in purely non-evaluative terms. I'll also assume that the relevant notion of extension is the set of all actual and possible things to which a predicate correctly applies. Otherwise shapelessness will be cut off from such items as properties.

The idea that the evaluative is non-evaluatively shapeless in this sort of sense garners wide sympathy.⁹ Its appearance in a number of metaethical contexts is unsurprising. For although McDowell asks us to consider as an example "a specific conception of some moral virtue" (1981, 201), it is clear from context that his discussion is meant to apply not only to thick terms and concepts (such as *cruel*, *generous* or *tactful*) but to the evaluative in general, including thin terms and concepts (such as *intrinsically good* or *morally wrong*).¹⁰ For McDowell is concerned with "whether, corresponding to any value concept, one can always isolate . . . a feature that is there . . . independently of anyone's value experience being as it is" and whether ascriptions of value can accordingly be "disentangled" into two distinct components of sensitivity to a non-evaluative aspect of the world plus an attitude or sentiment that constitutes the experience of value (McDowell 1981, 200-1; cf. 1981, 216).¹¹

⁹See McDowell (1981), McNaughton (1988, 60-2), Hurley (1989, 13), Dancy (1993, 84-6; 1995), Wiggins (1993), Little (2000), Kirchin (2010), and Roberts (2011).

¹⁰I'll use italics to denote both terms and concepts (recall note 6 above) and capital letters to denote properties. I'll follow my opponents in writing as if thick terms and concepts are evaluative in meaning or content until §7 where I suggest that nothing in considerations of shapelessness supports this assumption.

¹¹McDowell is here talking about the evaluative in general despite his example of virtue terms; see Lang (2001, 197) and Roberts (2011, 503). As Roberts notes, this notion of disentangling must be distinguished from the notion that comes into play when we ask whether thick terms and concepts can be understood in terms of distinct non-evaluative and thin evaluative content; on this latter issue, see e.g. Williams (1985, 140-1), Blackburn (1992), Dancy (1995) and Elstein and Hurka (2009).

The shapelessness relation may be understood in many different ways when it comes to the semantic, metaphysical, and epistemological features of the relation and its (fundamental) relata. Is it evaluative terms, concepts or properties that are shapeless (or shapely), and are they supposed to be shapeless relative to non-evaluative terms, concepts or properties? And, whatever the relata may be, is the relationship itself a semantic or conceptual matter, or a matter of metaphysics which might have no reflection at the level of language or concepts? Similarly, is the relationship supposed to be knowable through reflection on concepts or some other general *a priori* method, or only through substantive evaluative theorizing (be this *a priori* or not)? It is often unclear just how the notion of shapelessness is supposed to be understood in these respects.¹²

A crucial question about the interpretation of the shapelessness thesis concerns the relationship between evaluative terms and concepts and the non-evaluative classifications to which their extensions would correspond if the shapelessness thesis were false. Does the shapelessness thesis say merely that no such non-evaluative classifications are built into the meanings of evaluative terms or concepts? This is the reading we get from McDowell's claim that understanding why certain non-evaluatively described items belong together may require understanding an evaluative term if the relevant notion of understanding an evaluative term is a matter of grasping its sense. Or does the shapelessness thesis say more strongly that evaluative terms and concepts have no non-evaluative shapes at all? This is the reading we get if the relevant notion of understanding an evaluative term requires not merely grasping its sense but also some further not merely conceptual competence.¹³ While the former option is perhaps a more natural way to interpret "understanding a term" and the reasoning that McDowell him-

¹²The semantic and epistemological features of the relation are unclear in McDowell (1981, 200-2) and Dancy (1993, 76, 85). Its fundamental relata are unclear in Little (2000, 279). Some of these choice points, and some of the available options, are registered and discussed by Miller (2003, 250-4), Kirchin (2010) and Roberts (2011).

¹³This distinction is closely related to the distinction between "intensional" and "extensional" readings of shapelessness and its negation in Miller (2003, 252-3). Both Miller and Lang (2001, 203-4) argue that McDowell conflates this distinction.

self deploys, the latter option fits better with the possibility that two thinkers who grasp the same sense can assign it different extensions and thereby disagree about its correct extension. These readings are thus distinct: while grasping the sense of an evaluative term is plausibly necessary for grasping its extension, the latter cannot in general be part of what explains the former.¹⁴ Conceptually competent speakers can in general be wrong about the correct extensions; evaluative terms and concepts are no exception.

We can use the broader metaethical context to guide our interpretation of the shapelessness thesis. If the thesis were taken in the former way, as denying only that grasping the sense of an evaluative predicate provides access to a kind recognizable as such at the non-evaluative level, then it would fail to threaten non-cognitivism about evaluative judgment in the way it is sometimes claimed to do.¹⁵ Non-cognitivists needn't claim that the connection between non-evaluative and evaluative qualities is conceptual or fixed by linguistic convention; as Simon Blackburn puts it, "to tell which moral quality results from a given natural state means using standards whose correctness cannot be shown by conceptual means alone. It means moralizing, and bad people moralize badly, but need not be confused" (Blackburn 1984, 184). Interpreting shapelessness merely as a thesis about the meanings of evaluative predicates would similarly have no bite against reductive evaluative naturalism according to which evaluative terms (such as *intrinsically good*) ascribe properties that can also be ascribed in independently intelligible non-evaluative terms (such as *pleasant*, in the crude case of metaethical hedonism) but this property identity is a synthetic truth knowable only *a posteriori*. Yet the shapelessness thesis is typically presented as incompatible

¹⁴If grasping a concept or the sense of an expression cannot plausibly be a matter of grasping its extension, then by Kirchin's own lights his whole line of argument in favor of shapelessness is wrong (Kirchin 2010, 25). The alternative is implausible especially in the case of context-sensitive terms, since competence with them cannot plausibly require the ability to grasp their extensions relative to an arbitrary context.

¹⁵McDowell himself appeals to shapelessness to support skepticism about a commitment he attributes to non-cognitivism, namely that ascriptions of value can be "disentangled" in the way described in the text above.

with such reductionism. Nor would this interpretation support reasons holism, the thesis that any non-evaluative consideration that is a reason to ϕ in one context may be no reason at all, or even a reason not to ϕ , in another context. Whether reasons for action exhibit this kind of sensitivity to non-evaluative context concerns the nature of reasons themselves.

What these examples illustrate is that the sorts of purposes for which the shapelessness thesis gets invoked require interpreting it as a claim about the semantic values, not just the senses, of evaluative predicates. Not even substantive evaluative theorizing about the correct extensions of evaluative predicates will uncover independently intelligible non-evaluative relations of real similarity under which those extensions are unified.¹⁶ This is the shapelessness thesis on which I'll focus. It may be summarized as follows:

(ST) The extensions of evaluative terms and concepts aren't unified under independently intelligible non-evaluative relations of real similarity, irrespective of whether such similarity relations were supposed to be built into the meanings of evaluative terms or concepts.

In what follows I'll assume that the notion of shapelessness in ST – a relation that relates terms or concepts to something like properties and may hold as a matter of synthetic truth – is what is required by the uses to which shapelessness tends to get put in metaethics.¹⁷

¹⁶Compare Miller (2003, 250, 253) and Roberts (2011, 505-6).

¹⁷ST should be distinguished from a yet different thesis, one that drops any reference to items like predicates and concepts and instead concerns properties on *both* sides of the shapelessness relation. This property-property notion of shapelessness generates the claim that evaluative properties (not predicates or concepts) are such that all the actual and possible things that bear a given evaluative property aren't unified under any non-evaluative relation of real similarity. The difference between this purely metaphysical shapelessness thesis and the shapelessness thesis I'll be focusing on is that the set of things which the purely metaphysical thesis is about needn't be the extension of any evaluative predicate or concept, so long as it corresponds to a property. Whether this distinction makes for a robust difference regarding the putative metaethical implications of shapelessness depends on further issues. For instance, is the connec-

What a positive argument for ST would look like is unclear. This may be partly because shapelessness is characterized as a feature of extension and reference-fixing for evaluative terms is generally a controversial and murky issue in metaethics. But several potential avenues of argument are quickly ruled out. ST is supposed to be an in-principle claim that won't follow simply if evaluative classifications correspond only to non-evaluative patterns that are too complex to be captured in language or thought. Nor does ST follow from the view that the evaluative is autonomous with respect to the non-evaluative in the sense that there is no reasonable inference from purely non-evaluative premises to any evaluative conclusion. Nor does it follow from non-naturalism in metaethics. Non-naturalists can hold, for instance, that all and only intrinsically good things increase utility (where utility is defined in non-evaluative terms), so long as they don't identify one with the other; this would be to say that intrinsic goodness has a non-evaluative shape.¹⁸ And although ST would seem to follow from any response-dependence view of evaluative concepts or properties on which the responses and conditions which determine extension cannot be characterized in purely non-evaluative terms, it is ST rather than the response-dependence account that is the less controversial of the two.¹⁹ Nor do we find any direct argument between properties and extensions close enough for the distinction to make little difference? Is non-evaluative property or relation best understood as a property or relation that can be ascribed in non-evaluative terms, so that a reference to something conceptual or linguistic is required after all? However these chips may fall, I'll set aside the purely metaphysical thesis in this paper because the discussions of the shapelessness thesis that can be found in the literature tend to characterize it explicitly as a claim about the extensions or correct applications of evaluative predicates or concepts, just as ST has it. See McDowell (1981, 200-2), Dancy (1993, 76, 85), Kirchin (2010, 7, 12) and Roberts (2011, 505-6).

¹⁸Among non-naturalists, G. E. Moore allows this as a possibility (Moore 1903, §6). Also note that (non-reductive) evaluative naturalism is compatible with ST. Properties ascribed by using evaluative terms can be natural even if they cannot be ascribed also by using non-evaluative terms (Sturgeon 2003).

¹⁹For the relevant kind of response-dependence account, often called "sensitivity theory", see McDowell (1985) and Wiggins (1993).

ment for ST in the literature.²⁰ McDowell claims only to be removing a “prejudice” against the non-evaluative shapelessness of the evaluative, not giving a positive argument for it, when he argues that it isn’t ruled out by supervenience.

Whatever the argument for ST might be, the strongest conclusion that I have seen explicitly advocated (rather than merely suggested) is the epistemic possibility claim that the evaluative is non-evaluatively shapeless for all we know (Kirchin 2010, 19). I suppose Shapelessness may well be true for all I know. But by the same token it is false for all I know. What I am interested in here, however, is less the truth of ST itself than whether the notion of shapelessness found in ST is characteristic of the evaluative in particular or carries the sorts of broader metaethical implications that tend to get attributed to ST.

3. Implications of ST: A Dilemma

Why does it matter whether shapelessness is characteristic of the evaluative in particular? We have seen that ST is often taken to carry significant metaethical implications. Now ST is restricted by definition to the evaluative. But suppose that the notion of shapelessness deployed in ST is applicable to terms and concepts in many non-evaluative domains as well. And suppose that its application in these other domains can be explained by noting that the extensions of terms in these domains are also not unified under similarity relations of other sorts. That would seem to provide a general explanation of shapelessness phenomena across a variety of domains. The availability of such an explanation might well make us wonder why the shapelessness of the evaluative should have significant metaethical implications. We should wonder about this especially if shapelessness phenomena in other domains didn’t have comparably significant

²⁰McDowell appeals to Wittgensteinian rule-following considerations (McDowell 1981). Those considerations may support the idea that evaluative terms and concepts cannot be attributed a non-evaluative shape simply on the basis of conceptual or other general *a priori* reflection. As noted earlier, however, this falls short of ST (Miller 2003, 253). In a discussion of the role of rule-following considerations in metaethics, Lang (2001) argues that such considerations don’t support the view that the extensions of evaluative terms are essentially uncodifiable.

implications in these other domains.

To illustrate this prospect, consider that one might hold that the classifications drawn by mental or psychological terms aren't unified under non-psychological similarity relations.²¹ I suppose it may be true, for all I know, that the extensions of such psychological terms as *pain* or *belief* aren't unified under non-psychological similarity relations. Or false for all I know. But either way the shapelessness of the mental would seem to lack significant consequences in philosophy of mind which would be distinctive to shapelessness. The shapelessness of the mental doesn't rule out non-reductive physicalism about the mental because it is compatible with functionalism and anomalous monism, which don't entail physicalism but are compatible with it.²² Shapelessness of the mental is compatible with functionalism because it is compatible with the multiple realizability of the mental by a variety of physical media; multiple realizability allows that these physical media have in common nothing other than a capacity to support a distinctive functional profile (Fodor 1974; Block 1980). And it is compatible with anomalous monism because anomalous monism claims that there are no strict psychophysical laws but every token mental event is token-identical to some physical event; this is compatible with a lack of real non-psychological similarity among the physical token events that are identical to different tokens of the given type of mental event (Davidson 1970). Nor would the shapelessness of the mental open up any distinctive "explanatory gap" between the mental and the physical, be required to explain why mental events cannot be explained solely by physical events without auxiliary assumptions concerning the relation between the mental and the physical or rule out non-strict psychophysical laws.

If the shapelessness of the mental lacks distinctive significant implications in philosophy of mind, one begins to wonder why it should nonetheless have such implications in metaethics. The sentiment grows stronger once we notice

²¹For a sympathetic discussion of this sort of idea, see Child (1993).

²²Although the falsity of reductive physicalism would be a significant consequence, I doubt that shapelessness provides the best, let alone the only, route to it. But I acknowledge that my move in this paragraph may not move those physicalists whom the shapelessness of the mental would make really nervous.

that the above points about the shapelessness of the mental parallel the earlier observations that the shapelessness of the evaluative has no distinctive consequences for the autonomy of ethics or the debate between ethical naturalism and non-naturalism.²³ Avoiding this parallel reasoning would seem to require that the relations between the mental and non-mental and the evaluative and non-evaluative be disanalogous not just in any old respect, but specifically regarding the implications of shapelessness. It is at least not clear whether that is so.

In short, those who attribute significant metaethical implications to ST face a dilemma. If the relevant notion of shapelessness isn't characteristic of the evaluative in particular, then it isn't clear why ST should carry the sorts of distinctive metaethical implications that get attributed to it. But if the relevant notion of shapelessness is supposed to be characteristic of the evaluative in particular, then it will be a problem if ST can be explained on the basis of more general factors that have nothing in particular to do with being evaluative.²⁴ In what follows I'll focus on this second horn. We can try to determine whether there is a kind of shapelessness that is characteristic of the evaluative in particular by considering what is supposed to motivate ST. I'll argue that the typical motivation for ST reveals nothing that is characteristic of the evaluative in particular or supports the sorts of metaethical implications that get attributed to ST.

4. Shapelessness and Outrunning

Why is the shapelessness of the evaluative supposed to be attractive in the first place? Consider the following motivations for it, in addition to the passage

²³Nor would shapelessness rule out the existence of certain kinds of non-strict or hedged moral principles. One can deny moral particularism and yet hold that evaluative properties tend to figure in multiple different non-strict principles, such as 'All else equal, N₁s are Es', 'All else equal, N₂s are Es' etc. (where N_{*i*} is a non-evaluative property and E the target evaluative property) such that N₁, N₂, . . . , N_{*n*} aren't unified under real non-evaluative similarity. Therefore I doubt that Roberts (2011) is right that the shapelessness of the evaluative supports moral particularism.

²⁴Simon Kirchin claims that shapelessness is characteristic of the evaluative in particular (Kirchin 2010, 16). Other writers supply little textual evidence one way or the other.

quoted at the beginning of §2 from McDowell (1981, 202):

[Of] the infinitely many ways of being cruel – kicking a dog, teasing a sensitive person, and forgetting to invite someone to a party might each qualify – there is no way of saying what they have in common (and why, say, the pain inflicted during a spinal tap is different) except by helping oneself to the moral concept of ‘cruelty’. (Little 2000, 279.)

[I]t might be cruel to refrain from sharing chocolate with a young child who desperately wants it, but it can be kind if, in addition, we are acting because there is some risk of her teeth rotting in the future. In more complicated situations it might be kinder to share, despite the risk of tooth rot, because, say, someone has hurt her feelings and she needs comforting. . . . [Such cases] motivate us to see that the variation of features relevant to the ethical value of the situations they constitute can continue indefinitely. The key thought is that our concept of kindness might *outrun* any non-ethical characterization one could give of the actions deemed kind. (Kirchin 2010, 6.)

If there were such matching nonevaluative classifications . . . 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 . . . is not there for the outsider to detect: evaluative concepts are nonevaluatively shapeless. (Roberts 2011, 506.)

These statements seem to be aimed at motivating ST as a good explanation of a certain phenomenon. The phenomenon is that the extensions of evaluative terms and concepts seem to “outrun” (as Kirchin puts it) non-evaluative classifications.

Whether a particular item falls under a given evaluative term or concept E cannot be reliably projected just from the non-evaluative features of the other instances of E , without an understanding of E . ST is supposed to provide a good explanation of this “outrunning point”.²⁵

A significant complication arises immediately. Earlier we distinguished two versions of the shapelessness thesis. The weaker, less interesting version says that the meanings of evaluative terms and concepts don’t have built into them non-evaluative similarity relations that would unify their extensions. The stronger, more interesting version laid out in ST says that the extensions of evaluative terms and concepts aren’t unified under non-evaluative similarity relations, irrespective of whether these were supposed to be built into the meanings of evaluative terms or concepts. We also saw that these two options correlate with two different notions of understanding a term - grasping its sense, and some notion stronger than such conceptual competence. To which option does the outrunning point correspond?

If understanding a term is a matter of grasping its sense, then the outrunning point says that whether a particular item falls under a given evaluative term or concept E cannot be reliably projected just from the non-evaluative features of the other instances of E , without having a grasp of the sense of E . This sort of claim might be fairly natural to read into the quotes from McDowell and Little. But this version of the outrunning point can be used to motivate at most the less interesting version of the shapelessness thesis. The stronger, more interesting version of the shapelessness thesis laid out in ST is simply not needed to explain this version of the outrunning point.²⁶ By the same token, if the outrunning point is read in this weak way, then the conclusion that the outrunning point provides no support for the sorts of metaethical implications that get attributed to ST will be relatively uninteresting.

²⁵It is unclear whether the outrunning point is to be understood as (merely) an epistemic claim or more strongly, such as some kind of constitutive claim. Locutions such as “there is no way of saying...” and “cannot be specified...” can be read either way. I hope my discussion to get by without speculation on this score.

²⁶Thanks to the anonymous referee for this journal for pressing me on this point.

The outrunning point can, however, be given a stronger, more interesting interpretation. Kirchin's central claim is that the extension of *kind* is sensitive to potentially indefinite variation in the non-evaluative features relevant to kindness. Roberts appeals to certain conditions for learning the extension of a concept beyond the condition that grasping its sense is necessary for grasping its extension. Neither seems to make any fundamental appeal to a specifically conceptual competence in articulating the outrunning point. A stronger version of the outrunning point would therefore say that whether a particular item falls under a given evaluative term or concept *E* cannot be reliably projected just from the non-evaluative features of the other instances of *E*, irrespective of whether these are built into the meaning of *E* or uncovered through some not merely conceptual inquiry.

Another complication is that the notion of extension relevant to ST is correct extension but the outrunning point is often characterized in terms of the set of things to which some group of competent users believe the relevant term to apply. Both McDowell and Roberts claim that "outsiders" to some evaluative practice who don't possess a given evaluative concept or try to engage with the evaluative concerns reflected in the concept won't be able reliably to track how competent "insiders" would apply the term. But insiders can be mistaken in their beliefs about the term's extension. I'll assume, for reasons of charity, that the cases under discussion are those special cases where insiders have got the extension right.²⁷

To grant the outrunning point, on this interpretation, is to grant that one cannot reliably track whether some prospective cases fall under a predicate like *morally good* or *cruel* just on the basis of its other instances, as described without information provided by a grasp of its sense or some more substantive inquiry in which the predicate is deployed. Even if a person without such information

²⁷This is to assume that the stronger outrunning point articulated above can be got off the ground even in cases where someone grasps the sense of a predicate but has an incorrect view of its extension. It seems to me that if the outrunning point is plausible in the case of "outsiders", it should be plausible in this case, too.

recognizes what the actions picked out by descriptions (a)-(c) have in common, and what (d)-(f) don't, and even if she recognizes that competent users of the predicate apply E to (a)-(c) but not to (d)-(f), she may not be able to go on correctly identifying (g)-(i) also as things to which E applies without such information.²⁸

The writers who raise the outrunning point seem to differ as to whether it is characteristic of the evaluative in particular. Some of them think it is. Kirchin, for instance, writes: "Indeed, we are assuming that ethical concepts are special concepts where outrunning occurs, as opposed to other concepts – such as 'is a monarch' – where one might assume this does not happen" (Kirchin 2010, 16). Others suggest that evaluative terms and concepts aren't unique in this respect, but are nonetheless paradigmatic examples of a special class of "interest-relative" concepts "whose point is given by some human interest or concern" and whose extensions cannot be grasped without engaging with those concerns (Dancy 1993, 76; cf. McDowell 1981 and Wiggins 1993). Dancy doesn't say more about what makes a concept interest-relative, but I presume that paradigmatic examples are meant to include concepts like *funny* (reflecting our interest in humor), *tasty* (reflecting our interest in gustatory pleasure) and perhaps functional terms like *car* or *knife* (reflecting our interests in artifacts that perform certain functions). My suggestion is that the outrunning phenomenon seems more general still.

The outrunning phenomenon can be got going with respect to a term or concept T already if the extension of T isn't unified under any similarity relation expressible without T or, as I'll say, expressible in independently intelligible T -free terms. This might happen even if T isn't an evaluative or interest-relative term or concept. One very general view that would deliver this result is that natural language expressions cannot generally be analyzed or understood in terms of independently intelligible necessary and sufficient conditions. Just consider the notorious difficulties of analyzing even *horse* or *red* in such terms. The present point might be simply a local instance of this very general idea.

²⁸My formulation here borrows freely from Sreenivasan (2001, 19).

A less overarching and more concrete way to make the same point is to consider other examples. Thus, for instance, specifying the extension of a psychological term like *painful* seems to require an understanding of *painful* and thus seems like a good candidate for outrunning independently intelligible *painful*-free descriptions. Yet *painful* isn't an evaluative term. Nor does it seem to reflect any specifically human interest or concern. But if there is doubt as to whether psychological terms are interest-relative in this sense, note that some biological terms or concepts can also exemplify outrunning. For instance, the structures to which concepts in morphology refer may have a developmental and functional role that is robust even under changes or breaks in some underlying developmental process or molecular configuration: "Knocking out a gene need not have any impact on development due to genetic redundancy, while eliminating a particular cell or group of cells may dramatically interfere with normal development" (Brigandt and Love 2008, §4.2). The developmental effects of adding a gene or a group of cells may similarly differ across situations. If so, morphological concepts in biology seem to be able to do their job even if their extensions aren't expressible in "morphology-free" terms or unified under non-morphological relations of real similarity. These concepts aren't interest-relative in the above sense.

The foregoing suggests that the outrunning point is a fairly broad phenomenon that is explicable on the basis of fairly general factors that have nothing in particular to do with being interest-relative or evaluative. Thus even if there is no way of saying what the various ways of being cruel have in common without helping ourselves to information described using *cruel*, there is no reason why this should have anything in particular to do with whether *cruel* is an evaluative or interest-relative term. If *cruel* is such a term, then it specifies an evaluative (or interest-relative) shape which isn't unified under non-evaluative similarity relations according to ST. All that is required to explain outrunning, however, is that the extension of *cruel* cannot be expressed in independently intelligible *cruel*-free terms. That is independent of whether *cruel* is an evaluative (or interest-relative) term or concept. If *cruel* isn't an evaluative (or interest-relative) term, then it unifies things under the non-evaluative similar-

ity relation of counting as cruel. But there seems to be no compelling reason why this non-evaluative shape should have, even in principle, a *cruel*-free equivalent.²⁹ Whether there is or must be any independently intelligible *cruel*-free way to specify that non-evaluative shape would be simply a further issue that arises in a wide range of cases. Whether such claims as that “all instances of cruelty will be cruel” specify an *evaluative* shape is simply a further issue about what kind of a term or concept *cruel* is.³⁰

Although I have illustrated my argument with the thick term *cruel*, the basic idea works the same with respect to thin evaluative terms such as *wrong* or *morally good*. Philosophers hold different views concerning whether some thin concepts are to be analyzed in terms of others, and if so, which are to be analyzed in terms of which. All the same it is widely thought that if some evaluative terms are more basic than others, then the basic ones outrun the descriptions or classifications that can be drawn without deploying those terms. The intuition here is that there is no entry to evaluative distinctions without some grasp of notions that are basic to those distinctions. This implies some claims of the form taken by the claim that *morally good* outruns descriptions or classifications drawn in *morally good*-free terms. But again the outrunning point as such can be explained without any reference to whether *morally good* is an evaluative term. Although most of us think, of course that *wrong* and *morally good* are evaluative terms, the outrunning point arises irrespective of this fact.

The proponents of ST seem thus to misjudge the implications of the outrunning point they present as motivating ST. What the outrunning point regarding some term or concept *T* shows is at most that the extension of *T* or the conditions for falling under it cannot be captured in *T*-free terms. That shows at most that T-hood (the property to which *T* refers) isn't reducible to something non-

²⁹Skepticism about such descriptive equivalence in the case of thick terms and concepts is well formulated by Sreenivasan (2001, 19). See also Sturgeon (2003) and Elstein and Hurka (2009).

³⁰This is intended to contrast with the remarks on *tactful* in Roberts (2011, 519).

T-involving.³¹ Since nothing about evaluativeness (or interest-relativity) need come into this, nothing here is characteristic of the evaluative (or the interest-relative) in particular.

If the most that can be got out of the outrunning point is some kind of non-reductionism, then the outrunning point would seem to lack the sorts of metaethical implications that tend to get attributed to ST. Ethical naturalism has prominent non-reductive variants that are fully consistent with the autonomy of ethics. One can be a non-reductionist about value but reject moral particularism and go either way on the holism of reasons. Non-reductionists about value can even accept any form of “quasi-realist” non-cognitivism on which evaluative terms and concepts may designate natural properties in addition to their primary non-cognitive function.³² Yet non-reductionism can accommodate the sorts of concerns about “disentangling” that animate writers like McDowell in allowing that there is no entry to evaluative classifications from non-evaluative starting points. Nor does the outrunning point provide the most effective way of defending non-reductionism in ethics. It is open to a reductionist to reply that we should focus on getting the evaluative facts right and remain open to the possibility that this will require revising our views about the application conditions or even the extensions of our evaluative terms or concepts.³³

This completes my main argument for the claim that the outrunning point that is commonly used to motivate ST can be explained by more general factors that have nothing in particular to do with being evaluative. The outrunning point thus fails to support taking shapelessness to be characteristic of the evaluative in particular or attributing to ST the sorts of metaethical implications that get attributed to it. We have yet to see a notion of shapelessness that is

³¹I say ‘at most’ because this upshot seems to require taking the outrunning point as an in-principle point that doesn’t merely reflect the expressive limitations of language. My own preference would be to treat reduction as a relation between items like properties which may hold as a matter of a necessary synthetic *a posteriori* truth.

³²See e.g. Gibbard (2003, 88-102) for this kind of view.

³³See e.g. Railton (1986) for this kind of reductive naturalism in ethics.

plausibly characteristic of the evaluative in particular.³⁴

5. A Quietist Shuffle?

One might object that I have misstated what the foregoing line of argument shows. Suppose I am right that other areas beyond the evaluative which exhibit shapelessness or at least outrunning will be quite widespread. If so, then what my argument really shows is that one cannot take ST to have significant metaethical implications for topics like non-cognitivism, ethical naturalism, reasons holism or the like unless one is prepared to draw parallel conclusions in all of the other areas. The problem is that at least some proponents of ST might be quite happy to do this. McDowell, for instance, argues that the domains in which understanding a term requires immersion in a practice are widespread indeed. It isn't merely that grasping the extension of an evaluative term requires immersion in an evaluative practice. Even in mathematical cases like expanding an arithmetical series the rule governing the expansion cannot be viewed as something that can be grasped independently of the inclinations, acquired when being taught the rule, to go one way rather than another when we reach a certain point in the expansion of the series; human mathematical practice cannot be codified in independently accessible terms (McDowell 1981, 203-11).³⁵

Call "quietism" about some domain the view that the practice of engaging in thought and discourse in that domain is answerable only to standards internal to it which need (and perhaps can have) no independently intelligible external vindication. The objection may now be put by saying that if one is prepared to accept some very extensive form of quietism, then the argument that I am advancing doesn't stop one from accepting the putative metaethical implications

³⁴The potential objection that the shapelessness of the evaluative might be overdetermined by some more general factors and some special feature of the evaluative strikes me as unmotivated by considerations of theoretical economy. In any case, no such special feature seems to have come into view yet.

³⁵Here I grant for argument's sake that rule-following considerations can bear on extensional matters in some way that connects with ST. But compare note 20 above.

of ST.³⁶ I have no space adequately to assess the plausibility of such an extensive form of quietism. Two more limited replies to the quietist objection – one of them critical, the other more conciliatory – will have to suffice for the present purposes.

The critical reply is that the objection doesn't fully connect with my argument. I'll give two examples. First, quietism is usually raised in the context of debates about realism in some domain to make room for a moderate form of realism. But shapelessness appears to have no essential connection to realism. Shapelessness is characterized as a feature of extension. Anti-realists about a domain can surely allow that terms in that domain have extensions, and as far as I can see, they needn't require that those extensions are unified under independently intelligible similarity relations. Issues about quietism and shapelessness thus strike me as distinct in this respect. Second, my claim isn't that ST in fact has certain metaethical implications which we should reject because we should reject ST. Rather, my claim is that ST doesn't have those metaethical implications in the first place, perhaps apart from some very general kind of non-reductionism. My argument for doubting that ST has the sorts of metaethical implications regarding, for instance, non-cognitivism, ethical naturalism and reasons holism which get attributed to it is that shapelessness and the outrunning point lack parallel implications in many other domains where they have some plausibility. I noted that the shapelessness of the mental has no distinctive implications regarding (non-reductive) physicalism, the explanatory gap and the like. Similarly, debates about reductionism in biology have been running along independently of our asking such questions as whether morphological classifications correspond to kinds into which things can in principle be seen to fall independently of a biological practice. Whether this means that the integrity of the practice needs (or can have) no external justification strikes me as a further question. Of course an extensive form of quietism would *itself* have significant implications. But that is neither here nor there for the present purposes, since

³⁶Thanks to the anonymous referee for this journal raising this objection. The term 'quietism' has many senses in philosophy, of which my use is but one.

shapelessness doesn't support quietism in particular.

The conciliatory reply is that even if my critical reply falls short, my argument would still carry considerable interest. Extensive quietism would be a highly controversial position and (McDowell apart) most proponents of ST don't express sympathy for it. Those who attribute significant metaethical implications to ST thus face a choice if the quietist objection is correct: either give up on the sorts of metaethical implications that get attributed to ST or accept extensive quietism. This would be a significant choice point. So if my critical reply fell short, I would still be happy enough with a conditional version of my argument: one cannot accept the sorts of metaethical implications that get attributed to ST unless one is prepared to accept some very extensive form of quietism.

6. Complications of Context-Sensitivity

I'll now complement my main argument by showing how it is robust across certain complications that reveal weaknesses in existing discussions of shapelessness. The complications I have in mind arise from the fact that a wide range of evaluative terms are semantically context-sensitive. Terms from *morally good* to *courageous*, *generous* and *tactful* are all gradable because they measure qualities of which things can have more or less. But gradable terms – paradigmatic examples of which are terms like *tall* and *old* – are widely regarded as context-sensitive. One evaluative example is that *good* makes a different contribution to what utterances of sentences of which it is a constituent express in different contexts. For instance, what counts as a good guitar player is different from both what counts as a good philosopher, what counts as good with children or what counts as morally good or aesthetically good. Similarly, what counts as courageous is different when standing up to primary school bullies and saving the civilization are in question and what counts as generous is different when millionaires and paupers are in question; similarly, behaviors that count as selfish,

cruel or okay to do relative to one context may not do so relative to another.³⁷ If there is to be a shapelessness thesis that is characteristic of the evaluative in particular, it must accommodate context-sensitivity.

Context-sensitivity is relevant to shapelessness and outrunning because both are typically characterized as features of extension and the extension of a term varies with context if the term is context-sensitive. Because *old* is a context-sensitive term, there is no such thing as *the* set of old things, period. There is only the contextually shifting set of the things of which *old* is true relative to a (non-defective) context. Nor may the things correctly called *old* in one context (e.g. when redwood trees are under discussion) share age with the things correctly called *old* in another (e.g. when bumblebees are under discussion).³⁸ Thus the things correctly called *old* in one context may lack any real similarity with the things correctly called *old* in another. No wonder, then, that one may not be able to track what is correctly called *old* in a given context on the basis of a grasp of various other instances!

The same will go for any context-sensitive evaluative term. Thus the very question whether one can predict whether some prospective case falls under a term on the basis of what things fall under it in other contexts seems ill-defined when extension varies with context. But recall that existing discussions of outrunning or shapelessness lay considerable stock on whether one can track what prospective cases fall under an evaluative term on the basis of non-evaluative descriptions of its other instances. Insofar as these discussions presuppose that evaluative terms have context-invariant extensions, they will fail to apply to

³⁷I won't speak of concepts as context-sensitive because the label is more appropriate for linguistic expressions. Insofar as concepts are closely related to the senses of linguistic expressions, one might understand them as functions from contexts to contents, or what David Kaplan calls the "character" of an expression (Kaplan 1989).

³⁸Let me forestall some potential misunderstandings. The semantic value of a context-sensitive term relative to context can be an intrinsic property, such as height in the case of *tall*). Moreover, which intrinsic property it will be needn't depend on a comparison with a comparison class of other individuals; for instance, one can count as heavy, or tall, relative to the maximum safety limit for a rollercoaster.

evaluative terms in general.³⁹ Shapelessness and outrunning must thus be understood as features of extension relative to context.

A lot more could be said about how shapelessness and outrunning are related to context-sensitivity. But most of this would be tangential to my main argument, so let me simply note that invoking context-sensitivity isn't meant to suggest that outrunning or shapelessness are nothing but manifestations of context-sensitivity.⁴⁰ It is rather to acknowledge a potential objection. A sophis-

³⁹One might claim that according to McDowell, mastery of evaluative concepts requires a sensitivity to evaluative features which is a form of sensitivity to context. Now determining whether something falls under an evaluative term *E* may well require a sensitivity; perhaps one needs to appreciate that some things are *E* in virtue of being *P*, other things are *E* in virtue of being *Q* & *R*, something that is *P* may not be *E* if it is also *S* and so on. But this may be true even if *E* is a context-invariant term. It is quite another thing to say that the standard for *E*-ness itself varies with context. McDowell incorporates the former point, but not the latter.

⁴⁰Let me nonetheless note that some natural objections to using context-sensitivity to explain outrunning or shapelessness are weaker than they might seem. One might object that the comparatives of evaluative terms (e.g. *morally better than* or *more courageous than*) exhibit outrunning and shapelessness just as much their positive forms – for instance, that what is correctly called *more courageous than* what is no more unified under a non-evaluative similarity relation than what is correctly called *courageous*. The objection then has it that context-sensitivity couldn't be used to explain this because the comparatives of context-sensitive gradable terms, such as *older than* and *taller than* are context-invariant. (This objection is due to Jamie Dreier.) However, the comparatives of “multidimensional” gradable terms (e.g. *more painful than* or *balder than*) are context-sensitive because the different dimensions (e.g. intensity and duration in the case of *painful*) may get weighted differently in different contexts and thus yield different extensions (Benbaji 2009; van Rooij 2011). But since a wide range of evaluative terms (e.g. *courageous*, and perhaps *morally good*) are also multidimensional in this sense, their comparatives will similarly be context-sensitive; this narrows down the range of potential context-insensitive evaluative predicates. For instance, impartiality may be highly relevant to what counts as morally good in allocation of public resources but much less relevant in the context of special relationships. So context-sensitivity could after all be used to explain why the comparatives of these evaluative terms exhibit outrunning and shapelessness insofar as they do so. But the explanation has nothing in particular to do with being evaluative. See Väyrynen (forthcoming) for discussion. The only point I'll argue below in the text is that when an evaluative term isn't multidimensional, it is much harder to show that its extension relative to context is shapeless or exhibits outrunning.

ticated champion of ST might claim that the outrunning point and ST remain plausible even if they are revised into claims about the extensions of evaluative terms relative to context. This would leave the door open for the claim that there is something here that is characteristic of the evaluative in particular. Perhaps the extensions of evaluative terms relative to context exhibit outrunning and aren't unified under non-evaluative similarity relations even though the extensions of non-evaluative context-sensitive predicates relative to context tend to be unified in a way that doesn't generate outrunning.⁴¹ (For instance, letting Y_c be the standard for youngness relative to context c , things that meet the standard are unified by the property of being less than Y in age.)

My reply is twofold: first, the extensions of many evaluative terms relative to context are in fact non-evaluatively shapely; second, in the remaining cases my main argument from §4 (with the potential quietist proviso added in §5) applies with equal force to show that ST understood as a claim about extension relative to context won't have the kinds of metaethical implications that tend to get attributed to ST.

Many evaluative terms have “shapely” extensions relative to context. This will be clearest in the case of unidimensional evaluative terms, that is, terms whose application is a function of just one quality of which things can have more or less. For instance, a term like *selfish* will be unidimensional if selfishness is understood as involving preferring one's own happiness or interests over some (to be specified) amount of a greater contribution to the happiness or interests of others.⁴² By contrast, a term like *morally good* will be multidimensional insofar as its correct application relative to context is some function of multiple qualities (such as well-being, desert, justice or whatnot) and some relative weighting of those qualities. Determining the extension of *selfish* relative to context requires that context fix some or other differential in contribution to happiness or interests as the standard of selfishness relative to context. Not all contexts may of course fix a standard; some contexts are defective. It is also possible for people to

⁴¹This point was raised independently by Jamie Dreier and Matti Eklund.

⁴²Here I follow Elstein and Hurka (2009, 522).

dispute the contextually appropriate standard; this will be one way to disagree about the extension of a context-sensitive term like *selfish*.⁴³ Both sorts of cases are irrelevant to ST, since the notion of a shape is undefined for indeterminate extensions. When the context *does* fix what differential in preferring one's own happiness over a greater happiness for others is minimally sufficient to count as selfish, anything that makes the grade is unified by the property of exceeding that threshold. That property seems to be intelligible independently of one's understanding of selfishness; one can grasp what it is to prefer some amount of happiness for oneself over ten units more happiness for others without thinking of such a preference in terms of whether it is selfish or not. Such a property also seems relevantly non-evaluative. Thus the extension of *selfish* relative to context seems to be unified under an independently intelligible non-evaluative similarity.⁴⁴

Is this argument limited to unidimensional evaluative terms or does it generalize some way further to at least some multidimensional evaluative terms? This depends on the difficult general issue of how the extensions of evaluative terms relative to context are determined in the first place. But considering the implications of one candidate view about the matter will illustrate its relevance to ST. To understand what the utterances of many sentences of which *good* is a constituent say, we need to have some idea of what dimensions of evaluation are relevant – are we talking about moral goodness, being good with children or what? As noted earlier, context might fail to fix the relevant dimensions (some contexts are defective) and people may disagree about what the relevant dimensions are. This may happen even in the case of such more specific predicates

⁴³I speculate, but cannot argue here, that much of what gets called the “essential contestability” of evaluative concepts may be explicable in terms of this more general possibility of reasonable disagreements over the contextually relevant standard. This may happen even in such cases as ‘What counts as tall around here?’ (Sundell 2011, 278-9). If so, this spells trouble for the suggestion (made e.g. by Roberts 2011) that essential contestability may combine with shapelessness to yield a hallmark of the evaluative.

⁴⁴The issues that arise in this paragraph are addressed in greater length, in the context of thick terms and concepts, in Väyrynen (forthcoming).

as *a good philosopher*: people may disagree about just which qualities (among such candidates as originality, creativity, clarity, rigor, insight, carefulness and more) are relevant to whether someone counts as a good philosopher relative to context, to what extent they must possess these qualities and how these qualities are to be weighted. The term has no (determinate) extension unless and until these matters are fixed. But at least in some contexts they do get fixed, at least determinately enough for the purposes at hand. (Hiring does, after all, happen in philosophy!) When they are, I can see no reason why at least those things that are determinately in the extension of *a good philosopher* relative to context shouldn't be fairly unified under the set of qualities plus their relative weighting which ended up contextually salient.

The relevance of these observations is that the outrunning point is more difficult to get going once we focus on contexts where the relevant qualities and their weighting are fixed. If these factors are what fix extension relative to context, they can be used to assess what else falls into that extension. If the context determines, for instance, clarity and originality as important, and as more important than rigor and carefulness, we would seem to have a decent shot at projecting from some things in the extension of the term relative to context to others. It is no objection to say that in different contexts a slightly different set of qualities or their relative weighting might be different. This is just to repeat that *a good philosopher* is a context-sensitive term. That fact is irrelevant to whether the extension of *a good philosopher* relative to a given context is shapeless or exhibits outrunning. The parallel line of thought would seem to apply to many other multidimensional context-sensitive terms, whether evaluative or non-evaluative, from *painful* and *bald* to *courageous* to *morally good*.

The suggestion I am making is that shapelessness and outrunning are features of at most a limited range of evaluative terms when understood as features of extension relative to context. How limited this range is will depend on how these extensions are determined. If they are fixed by determining what qualities are contextually relevant for falling under the term in question and how context

weights those qualities, then the examples above would seem to favor greater confidence that the extensions of evaluative terms relative to context will generally be non-evaluatively shapely than that they will be shapeless. (This would challenge the epistemic claim that ST is true for all we know.) Whether things are different if the extensions of evaluative terms relative to context are determined in some other way remains open to future work. My aim is simply to note that the issue is highly relevant to the plausibility of ST.

There is probably no *a priori* argument to show whether the extensions relative to context of any evaluative terms are left over as non-evaluatively shapeless once the issue just discussed is settled. It is therefore important to show that my main line of argument from §4 would apply with equal force to context-sensitive terms. The claim to show is that the outrunning point can be got going even regarding extension relative to context already if the extension isn't unified under any similarity relation expressible without the target term, and that this has nothing to do with being evaluative in particular.

Suppose that there is no way of saying what the various different ways of counting as cruel relative to context have in common without helping ourselves to our understanding of cruelty. All this shows is that the extension of *cruel* relative to context cannot be specified in *cruel*-free terms, and this has nothing to do with being evaluative in particular. If *cruel* is evaluative in meaning, then its extension relative to context has an evaluative shape which (if ST is true) corresponds to no classification specifiable in independently intelligible non-evaluative terms – but these would be *cruel*-free terms. If *cruel* isn't evaluative in meaning, then its extension relative to context is unified under the non-evaluative similarity relation of counting as cruel – but whether there is any independently intelligible *cruel*-free way to specify that non-evaluative shape is, again, simply a further issue. And again the argument works the same with respect to thin evaluative terms. If the extension of *morally good* relative to context outruns descriptions given in *morally good*-free terms, this can be explained by noting that many a term *T*, evaluative or otherwise, is such that its extension relative to

context cannot be unified by similarity relations describable in *T*-free terms.⁴⁵ Nothing about being evaluative need come into it. Explaining such phenomena doesn't therefore require invoking any putative characteristic of the evaluative in particular. It is thus doubtful that shapelessness or outrunning, understood as features of extension relative to context, support the sorts of metaethical implications that tend to get attributed to ST.

7. Conclusion

In this paper I have done three things. I clarified the thesis that the extensions of evaluative terms are non-evaluatively shapeless and noted that the thesis is typically motivated by appeal to the outrunning point. I argued that neither the shapelessness thesis nor the outrunning point are characteristic of the evaluative in particular because both can be explained by more general factors that have nothing in particular to do with being evaluative, and that the argument is robust across complications that arise from the context-sensitivity of many evaluative terms. And I argued that if shapelessness isn't characteristic of the evaluative in particular, we shouldn't expect it to carry the sorts of distinctive metaethical implications that tend to get attributed to it.

I'll conclude by considering briefly how my argument bears on thick terms and concepts in ethics, such as *cruel*, *selfish*, *courageous* and *generous*. ST has been invoked to argue that the evaluative and non-evaluative aspects of thick terms and concepts aren't separable.⁴⁶ This argument is increasingly recognized to be ineffective: the "inseparabilist" can accept ST but hold that the extensions of thick terms and concepts are partly determined by evaluation and partly by distinct non-evaluative description.⁴⁷ At stake in this debate is the larger issue

⁴⁵My earlier examples of *painful* and terms standing for morphological concepts in biology would seem to carry over. If these terms aren't context-sensitive, their extension relative to context will be the same in every context, so they can be treated as a limiting case.

⁴⁶See e.g. Williams (1985, 140-1) and Dancy (1995). Contrary to a common assumption, this argument isn't found in McDowell (1981). His notion of "disentangling" is meant to apply to evaluative terms and concepts across the board; see the text surrounding note 11 above).

⁴⁷See e.g. Elstein and Hurka (2009) and Roberts (2011).

whether the meanings of thick terms and concepts can be understood reductively as some combination of non-evaluative description plus thin evaluation (such as *good for its own sake* or *right*) or whether the evaluative character of thick terms and concepts is somehow irreducibly thick.⁴⁸

The debate just described assumes that thick terms and concepts are evaluative in meaning and focuses on disputing the more specific sense or way in which they are evaluative, in part by reference to what the non-evaluative shapelessness of thick terms and concepts would show in this respect.⁴⁹ What my argument implies is that this way of framing the debate cannot be justified by considerations of shapelessness, and indeed that my terminological convenience of writing throughout the paper as if thick terms and concepts fell into the category of evaluative ones is controversial. For if shapelessness isn't characteristic of the evaluative in particular, then thick terms and concepts could well be shapeless, in the sense that their extensions relative to context aren't unified under similarity relations specifiable in "thickness-free" terms, even if they weren't evaluative in meaning. All this applies to the outrunning point. Treating any particular notion as evaluative in meaning requires support that is independent of the outrunning point or ST. We should therefore focus on settling whether thick terms and concepts are evaluative in meaning before considering whether ST tells us anything interesting about them. It is an open possibility that careful attention to whether it is coherent to grant that something is selfish (lustful, lazy, courageous, generous and so on) but suspect the presumption that it is thereby bad (or good) in any particular way will show that thick terms and concepts aren't evaluative in meaning.⁵⁰ In that case ST wouldn't even apply to thick terms and concepts; they would rather be closer to the non-evaluative instances of the

⁴⁸See Elstein and Hurka (2009), Kirchin (2010), Roberts (2011) and Parfit (2011, 315-23).

⁴⁹I mean 'evaluative in meaning' in the sense that, for instance, *A is selfish* entails, as a conceptual or semantic matter, that A is bad in a certain sort of way (where the particular way of being bad may or may not be analyzable in thin evaluative terms). See Väyrynen (2012) for a further elaboration.

⁵⁰I discuss this issue extensively in Väyrynen (2012). For some earlier work, see Hare (1952, 121-2; 1981, 17-18, 73-75), Blackburn (1992) and Richard (2008, 32-33).

more general shapelessness phenomenon described in this paper. In this modest way my argument bears on the question of what terms or concepts count as evaluative in the first place.

References

- Benbaji, Yitzhak (2009), “Parity, Intransitivity, and a Context-Sensitive Degree Analysis of Gradability”, *Australasian Journal of Philosophy* 87, 313-35.
- Blackburn, Simon (1981), “Rule-Following and Moral Realism”, in *Wittgenstein: To Follow A Rule*, ed. S. Holtzman and C. Leich (London: Routledge), 163-87.
- (1984), *Spreading the Word* (Oxford: Clarendon Press).
- (1992), “Through Thick and Thin”, *Proceedings of the Aristotelian Society*, supplementary volume 66, 284-99.
- Block, Ned (1980), “Troubles with Functionalism”, in *Readings in the Philosophy of Psychology*, volume 1, ed. N. Block (Cambridge, MA: Harvard University Press).
- Brigandt, Ingo and Love, Alan (2008), “Reductionism in Biology”, *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, ed. E. N. Zalta.
(plato.stanford.edu/archives/fall2008/entries/reduction-biology/)
- Child, William (1993), “Anomalism, Uncodifiability, and Psychophysical Relations”, *Philosophical Review* 102, 215-45.
- Dancy, Jonathan (1993), *Moral Reasons* (Oxford: Blackwell).
- (1995), “In Defence of Thick Concepts”, *Midwest Studies in Philosophy* 20, 263-79.
- Davidson, Donald (1970), “Mental Events”, reprinted in his *Essays on Actions and Events* (Oxford: Clarendon Press, 1980), 207-24.
- Elstein, Daniel Y. and Hurka, Thomas (2009), “From Thick to Thin: Two Moral Reduction Plans”, *Canadian Journal of Philosophy* 39, 515-36.
- Fodor, Jerry (1974), “Special Sciences and the Disunity of Science as a Working

- Hypothesis”, *Synthese* 28, 77-115.
- Gibbard, Allan (2003), *Thinking How to Live* (Cambridge, MA: Harvard University Press).
- Glanzberg, Michael (2007), “Context, Content, and Relativism”, *Philosophical Studies* 136, 1-29.
- Hare, R. M. (1952), *The Language of Morals* (Oxford: Clarendon Press).
- (1981), *Moral Thinking* (Oxford: Clarendon Press).
- Hurley, S. L. (1989), *Natural Reasons* (Oxford: Oxford University Press).
- Kaplan, David (1989), “Demonstratives”, in *Themes from Kaplan*, ed. J. Almog, J. Perry, and H. Wettstein (Oxford: Oxford University Press), 481-563.
- Kennedy, Christopher (2007), “Vagueness and Grammar: The Semantics of Relative and Absolute Adjectives”, *Linguistics and Philosophy* 30, 1-45.
- Kirchin, Simon (2010), “The Shapelessness Hypothesis”, *Philosophers’ Imprint* 10(4), 1-28. www.philosophersimprint.org/010004/
- Lang, Gerald (2001), “The Rule-Following Considerations and Metaethics: Some False Moves”, *European Journal of Philosophy* 9, 190-209.
- Little, Margaret O. (2000), “Moral Generalities Revisited”, in *Moral Particularism*, ed. B. Hooker and M. O. Little (Oxford: Clarendon Press), 276-304.
- McDowell, John (1981), “Non-Cognitivism and Rule-Following”, reprinted in his *Mind, Value, and Reality* (Cambridge, MA: Harvard University Press, 1998), 198-219.
- (1985), “Values and Secondary Qualities”, reprinted in his *Mind, Value, and Reality* (Cambridge, MA: Harvard University Press, 1998), 131-50.
- McNaughton, David (1988), *Moral Vision* (Oxford: Blackwell).
- Miller, Alexander (2003), *Introduction to Contemporary Metaethics* (Cambridge: Polity Press).
- Moore, G. E. (1903), *Principia Ethica* (Cambridge: Cambridge University Press).
- Parfit, Derek (2011), *On What Matters*, volume 2 (Oxford: Oxford University Press).
- Railton, Peter (1986), “Moral Realism”, *Philosophical Review* 95, 163-207.
- Richard, Mark (2008), *When Truth Gives Out* (Oxford: Oxford University

- Press).
- Roberts, Debbie (2011), “Shapelessness and the Thick”, *Ethics* 121, 489-520.
- van Rooij, Robert (2011), “Measurement and Interadjective Comparisons”, *Journal of Semantics* 28, 335-58.
- Sreenivasan, Gopal (2001), “Understanding Alien Morals”, *Philosophy and Phenomenological Research* 62, 1-32.
- Sturgeon, Nicholas L. (2003), “Moore on Ethical Naturalism”, *Ethics* 113, 528-56.
- Sundell, Tim (2011), “Disagreements about Taste”, *Philosophical Studies* 155, 267-88.
- Väyrynen, Pekka (2012), “Thick Concepts: Where’s Evaluation?”, in *Oxford Studies in Metaethics*, volume 7, ed. R. Shafer-Landau (Oxford: Oxford University Press), 235-70.
- (forthcoming), “Thick Concepts and Underdetermination”, in *Thick Concepts*, ed. S. Kirchin (Oxford: Oxford University Press).
- Wiggins, David (1993), “A Neglected Position?”, in *Reality, Representation, and Projection*, ed. J. Haldane and C. Wright (Oxford: Oxford University Press), 329-36.
- Williams, Bernard (1985), *Ethics and the Limits of Philosophy* (Cambridge, MA: Harvard University Press).