

Slim Epistemology with a Thick Skin

Pekka Väyrynen

Abstract: The distinction between ‘thick’ and ‘thin’ value concepts, and its importance to ethical theory, has been an active topic in recent meta-ethics. This paper defends three claims regarding the parallel issue about thick and thin epistemic concepts. (1) Analogy with ethics offers no straightforward way to establish a good, clear distinction between thick and thin epistemic concepts. (2) Assuming there is such a distinction, there are no semantic grounds for assigning thick epistemic concepts priority over the thin. (3) Nor does the structure of substantive epistemological theory establish that thick epistemic concepts enjoy systematic theoretical priority over the thin. In sum, a good case has yet to be made for any radical theoretical turn to thicker epistemology.

1. Introduction

The distinction between ‘thick’ and ‘thin’ evaluative and normative concepts, and its importance to ethical theory, has been an active topic in recent meta-ethics. That debate has a reputation for obscurity which isn’t, I think, wholly unearned. But that needn’t be an obstacle in determining whether a similar distinction can be drawn between thick and thin epistemic concepts and what kind of importance such a distinction might have to epistemology.

This paper concerns meta-epistemology. It defends three claims concerning thick and thin epistemic concepts. There is no straightforward way to establish a good, clear distinction between thick and thin epistemic concepts on the basis of an analogy with thick and thin concepts in ethics (§2). Assuming there is such a distinction, its importance to epistemology cannot be established on semantic grounds; there is no semantic case for treating thick epistemic concepts as prior to the thin or otherwise taking a turn to a thicker epistemology (§3). Considerations regarding the structure of substantive epistemological theory also don’t establish that thick epistemic concepts enjoy systematic theoretical priority over the thin (§4). A good

case has yet to be made for a radical theoretical turn to thicker epistemology.

2. The Thick and the Thin in Ethics and Epistemology

Discussions of thick and thin concepts in ethics almost invariably begin with examples. Typical examples of thick evaluative or normative concepts include CRUELTY, BRUTALITY, EXPLOITATION, DECEITFULNESS, GENEROSITY, and GRATITUDE.¹ Typical examples of thin concepts are those expressed by many occurrences of words like ‘good,’ ‘wrong,’ and ‘ought.’²

But, moving beyond examples, it is hard to find an adequate account of what makes an evaluative concept count as thick and distinguishes these from concepts that are thin in some relevant contrasting sense. In an influential discussion, Bernard Williams says that thick concepts express a ‘union of fact and value’ in that their applicability is both ‘action-guiding’ (it at least typically indicates the presence of reasons for action) and ‘world-guided’ (it depends on how the world is in certain non-evaluative respects).³ It has become common to say that a term stands for a thick value concept if it expresses an evaluative concept with significant non-evaluative content, and that thick value concepts differ from the thin at least in the way they seem to combine description and evaluation as a matter of their meaning. It might be thought, for instance, that even if causing gratuitous pain is both cruel and wrong, its wrongness isn’t encoded in the very meaning of ‘wrong’ in the way its cruelty is encoded in that of ‘cruel.’

1 I’ll use small capitals to designate concepts. Words and other linguistic expressions will appear in single quotes when they are mentioned. Properties are designated by italics. Italics are used also for highlighting.

2 Since it’ll make no difference to my discussion whether there is some significant distinction between the evaluative and the normative, I’ll use ‘evaluative’ to cover both.

3 See Williams (1985: 128, 140). For some other ways of developing the general idea behind Williams’s distinction, see e.g., Gibbard (1992), Blackburn (1992), and Dancy (1995).

The existence of some sort of distinction between thick and thin concepts in ethics isn't in doubt. At least the typical examples of thick and thin value concepts seem clearly to differ, along some dimension of specificity, with respect to *how* descriptive they are. So far as their meanings go, 'cruel' and 'deceitful' seem to have richer non-evaluative meaning, and their applicability more robustly world-guided, than 'wrong' or 'bad.' This much fits the observation that the standard way to draw the distinction between thick and thin value concepts marks a difference in degree (of non-evaluative information encoded in the concept) rather than in kind (cf. Scheffler 1987: 417-18). But what remains unclear is how to draw a theoretically significant but more or less neutral distinction between thick and thin evaluative concepts.⁴

It is equally unclear how to distinguish between thick and thin epistemic concepts on the basis of an analogy with thick and thin concepts in ethics. The idea that thick value concepts have significant non-evaluative content whereas thin value concepts have no, or not much, significant non-evaluative content could be applied to epistemic concepts in two different ways.

One option is to take a *direct* analogy with the case of ethics and distinguish between thick and thin epistemic concepts in terms of the same distinction between *evaluative* and *non-evaluative* content. On this view both thick and thin epistemic concepts are evaluative concepts, distinguished by whether or not they have significant non-evaluative content.

The other option is to take a *structural* analogy with the case of ethics and distinguish between thick and thin epistemic concepts in terms of the distinction between *epistemic* and *non-epistemic* content. On this view thick epistemic concepts have some significant non-epistemic content and thin epistemic concepts have no, or not much, such content.

4 Eklund (MS) argues that several existing accounts of the thick-thin distinction are defective on this score.

Neither way of taking the analogy with the case of ethics is clearly good. To see this, first consider examples typically used to illustrate a distinction between thick and thin epistemic concepts. Standard ‘thin’ epistemic concepts include EPISTEMIC JUSTIFICATION, RATIONALITY, and KNOWLEDGE. Standard ‘thick’ epistemic concepts come in a somewhat more varied list, including INTELLECTUAL CURIOSITY, GULLIBILITY, and the concepts expressed by certain occurrences of ‘conscientious,’ ‘careless,’ ‘lucky,’ and ‘trustworthy.’ The existence of some sort of distinction between these groups of concepts isn’t in doubt. It also seems plausible that in the case of epistemology, too, a distinction between the thick and the thin marks a continuum rather than a binary distinction. Just as concepts like JUSTICE, IMPARTIALITY, EQUALITY, RIGHTS, AUTONOMY, and INTERESTS fit no more obviously with GOOD or RIGHT than with BRUTALITY or KINDNESS, so concepts like DECEIVING, PROPORTIONED TO EVIDENCE, INDUCTIVE SUPPORT, SIMPLICITY, COHERENCE, INSIGHT, and EPISTEMIC RESPONSIBILITY fit no more obviously with JUSTIFICATION or KNOWLEDGE than with INTELLECTUAL CAUTION or GULLIBILITY.

The direct analogy presupposes that all epistemic terms are evaluative. Its whole point is to distinguish thick epistemic concepts from the thin on the basis of whether they have some significant non-evaluative content *in addition to* evaluative content. But, even if *some* epistemic terms are evaluative, it is a substantive and controversial issue whether *all* are. For instance, whether epistemic concepts are normative is a central issue in debates over naturalized epistemology sparked by W.V.O. Quine (1969). The point can also be raised just with examples.

One type of example is that calling someone ‘quick to jump to conclusions’ or ‘intellectually biased’ often conveys some kind of negative evaluation of the person or their belief-formation. But, since words which aren’t as a matter of their meaning evaluative can still be used to communicate evaluations in contexts with suitable common grounds, this negative evaluative content might be merely pragmatically implicated or conveyed. Sometimes such negative evaluation might also not be

appropriate. Some hold that it can be epistemically appropriate for some beliefs to be influenced by certain intellectual biases or heuristics when beliefs formed in these ways tend in fact to be reliable.⁵

Another type of example is that various concepts of reliability often occupy an important role in epistemology, but don't seem to be evaluative. All these concepts concern some or other sort of probabilistic connection to truth, but neither TRUTH nor PROBABILITY is, itself, an evaluative or normative concept. What makes it true that smoke raises the probability of fire, for example, are the worldly nomological connections between the presence of the one and the presence of the other, not anything normative.

Thus the direct analogy from ethics to epistemology seems, at present, too quick. To further assess this charge, we can consider two ways to resist it. One move is to claim that concepts which seem epistemic but non-evaluative in fact are evaluative. For instance, suppose that jumping quickly to conclusions is sometimes but not always epistemically bad. This shows that QUICK TO JUMP TO CONCLUSIONS isn't an evaluative concept only if evaluative concepts cannot license different kinds of evaluation in different contexts, such as positive in some but neutral or even negative in others. Many discussions of thick concepts in ethics challenge just this assumption.⁶ But its truth or falsity is clearly not to be settled by fiat. Nor would its falsity alone establish that concepts like QUICK TO JUMP TO CONCLUSIONS really are evaluative concepts after all. So this first move on behalf of the direct analogy would go only so far.

Another move is to claim that concepts which seem epistemic but non-evaluative in fact aren't epistemic. Just as concepts like PLEASURE and PAIN, for instance, needn't be evaluative concepts to figure in moral

⁵ See e.g., Gigerenzer *et al.* (2000). Bishop and Trout (2005) argue that epistemic excellence doesn't ban ignoring factors that seem evidentially relevant; using simple statistical prediction rules based on a limited number of cues is at least no less efficient and reliable. On epistemic shortcuts generally, see Bach (1984).

⁶ This assumption is challenged by theorists as different as Blackburn (1992) and Dancy (1995). It is a crucial implicit premise in the argument that thick virtue concepts aren't in fact evaluative concepts in Brower (1988).

discourse and play a role in ethical theory, so concepts like RELIABILITY and QUICK TO JUMP TO CONCLUSIONS needn't be epistemic concepts (or evaluative concepts, for that matter) to figure in epistemic discourse or play a role in epistemology. If they were non-epistemic concepts, then they would obviously cut no ice against the direct analogy.

This second move on behalf of the direct analogy is viable only to the extent that there is a good distinction between epistemic and non-epistemic concepts. Obviously, the same is true of the structural analogy. The existence of some sort of distinction isn't in doubt, since there are clear cases of both epistemic and non-epistemic concepts. But notice that the present dialectical context requires a distinction which is also neutral on the issue whether there is a good distinction between thick and thin epistemic concepts. That is precisely the issue at stake. It wouldn't do to say, for instance, that epistemic concepts are those which have some significant conceptual connection to knowledge or epistemic justification, since KNOWLEDGE and EPISTEMIC JUSTIFICATION are among the paradigm examples of thin epistemic concepts. A dialectically kosher distinction between epistemic and non-epistemic concepts requires a basis that falls on neither side of the intuitive distinction between thick and thin epistemic concepts.

But now it has become difficult to see on what basis the requisite kind of distinction could be drawn. Numerous unclear cases reinforce the worry. For instance, are 'intellectually cautious,' 'trusting,' 'discerning,' or 'hasty,' as a matter of their meaning, epistemic terms? The answer is unclear. These terms seem to have no clear connections to concepts which would both suggest classifying them as epistemic and be dialectically admissible. Yet they don't seem clearly non-epistemic either. They can be used to describe things like judgments, inferences, and habits of belief-formation, which all fall under the subject matter of epistemology.

The structural analogy is thrown into doubt if there is no good basis for distinguishing between epistemic and non-epistemic concepts which falls on neither side of the intuitive distinction between thick and thin

epistemic concepts. The only serious candidate that I can think of, TRUTH, is dialectically inadmissible: it would sit ill with those advocates of thicker epistemology who think that contemporary epistemology is excessively focused on truth.⁷

Nor is the direct analogy rescued from doubt. One procedure for settling unclear cases is to classify a term as epistemic if we are uncertain whether it is epistemic.⁸ The procedure would count such a varied range of concepts as epistemic that it would be implausible that a wide range of concepts which seem epistemic but non-evaluative in fact are evaluative concepts or that they in fact aren't epistemic concepts.⁹ This result would render implausible precisely the claims made by our first and second defensive moves on behalf of the direct analogy.

I conclude that it is far from clear that there is a good, clear way to distinguish between thick and thin epistemic concepts on the basis of an analogy with ethics. The foregoing also gives some evidence that if a good, clear distinction is found, it won't be straightforward. In what follows, I won't try to say more about what would be the best way to understand a distinction between thick and thin epistemic concepts. More clearly remains to be said on the issue.

3. How Not to Strike it Thick: the Semantics of Epistemic Terms

Suppose that a good way to draw a distinction between thick and thin epistemic concepts were eventually found, whether by analogy with ethics or not. (Again, the existence of some sort of distinction isn't in doubt.) It would still be a further claim that this distinction makes some important difference in epistemology. I'll now argue that its importance cannot be established on semantic grounds: there is no systematic

7 See e.g., Elgin (2008). See also §4 below. I'm assuming that TRUTH isn't an epistemic concept.

8 For an analogous procedure for classifying terms into evaluative and descriptive, see Jackson (1998: 120).

9 My point here is dialectical. I do not mean to rule out the possibility that some terms which might seem epistemic (or of whose classification we are uncertain) in fact aren't, as a matter of meaning, epistemic terms.

semantic difference between thick and thin epistemic concepts which as such makes the distinction important to epistemology. My case is that the basic semantic treatment which is appropriate to (most) epistemic terms treats terms expressing thick concepts and terms expressing thin concepts fundamentally in the same way.

The semantic treatment I have in mind is the standard sort of formal semantics for gradable adjectives, such as ‘tall,’ ‘young,’ and ‘expensive.’ Gradable adjectives are well studied in linguistics. Thus, if what we know about their semantics is applicable to epistemic terms, then this will provide some independent basis for assessing whether the importance of a distinction between thick and thin epistemic concepts can be given a semantic basis.

The class of words which express epistemic concepts is grammatically heterogeneous. It includes adjectives, gerund and participle constructions, nouns, and verbs: consider ‘gullible’ and ‘coherent,’ ‘discerning’ and ‘proportioned to evidence,’ ‘justification’ and ‘intellectual curiosity,’ and ‘trusts’ and ‘knows,’ respectively. This isn’t a deep problem for my strategy. Most of these words are gradable or have a cognate gradable expression. So most epistemic concepts have linguistic expressions with just the syntactic features which are sufficient for gradability. First, they admit of comparatives: Things can be *more* rational or *better* proportioned to evidence *than* others, or *less* gullible or *less* epistemically responsible *than* others. Second, they take on degree modifiers like ‘very’ or ‘somewhat’: Things can be *quite* quick to jump to conclusions, *somewhat* intellectually curious, *totally* reliable, or *very* well justified.¹⁰ So, what we know about gradable expressions seems applicable to a wide range of epistemic terms.

What comparatives do semantically is to establish an ordering between two things regarding some common feature that can be enjoyed

¹⁰ Most epistemic terms also exhibit a further syntactic feature of gradable expressions: they tend to take on certain other modifier phrases, such as ‘for’ and ‘to’ prepositional phrases and ‘enough to’ adverbial phrases. Consider ‘reliable for a heuristic,’ ‘too gullible to be a detective,’ ‘incoherent enough to be delusional,’ etc.

to a greater or lesser extent. The standard semantics for gradable adjectives takes them to locate objects on a *scale* (such as tallness in the case of ‘tall’), defined as a set of *degrees* (or intervals) ordered with respect to some *property dimension* (such as *height* in the case of ‘tall’).¹¹ So the basic semantic value of a gradable adjective is a *function* from objects to values (degrees or intervals) on a scale.¹²

The interpretation of the comparative form is simple. (1a) is analyzed as (1b-c):

- 1a. Stan is taller than Buster.
- 1b. The value Stan takes on a scale of tallness is greater than the value Buster takes on a scale of tallness.
- 1c. **tall(Stan) > tall(Buster)**¹³

An adequate treatment of gradable adjectives must explain how comparatives are semantically related to their non-comparative positive correlates and capture the context-sensitivity which the latter exhibit in many cases. For instance, someone can count as tall relative to a context where silent film comedians are under discussion but not tall when basketball players are.

According to the scalar analysis, ‘tall’ can express different properties relative to different contexts of utterance because ‘tall’ allows context to determine some particular degree value *d* as the minimal point which an object must take on the scale of tallness—the scale along which ‘is taller than’ compares objects—to count as satisfying the predicate. So here is the positive form:

11 Kennedy (2007) requires that the ordering be a total ordering. But many gradable adjectives seem more plausibly linked to partial orderings rather than total orderings. In at least some cases the ‘scale’ isn’t linear.

12 See especially Kennedy (2007) and Glanzberg (2007). For earlier work developing this approach, see Cresswell (1977), Klein (1980), and von Stechow (1984). On the use of intervals instead of degrees, see Schwarzschild and Wilkinson (2002). Degrees (and intervals) are here understood as abstract representations of measurement.

13 I will use boldface, like **tall**, to designate semantic values, interpreted in this case as degrees on a scale.

- 2a. Stan is tall.
- 2b. The value Stan takes on a scale of tallness is greater than (or, at least as great as) the minimal value d required for counting as satisfying ‘tall’ in context c .
- 2c. **tall(Stan)** $> d$

The value of d (determined for the adjective by context) is the *standard* for that adjective and context. To count as satisfying ‘tall’ is to meet the contextually determined standard for tallness. Degree modifiers can then be understood as modifying the required degree value on the relevant scale. For instance, ‘very tall’ can be taken to increase the required degree of height.

We already saw syntactic evidence for counting most epistemic terms of varying degrees of thickness and thinness as gradable and thereby (following standard methodology in linguistics) for treating these terms as semantically linked to scales in the way outlined above. This latter claim has also some non-syntactic evidence behind it. Many examples suggest that whether something counts as satisfying one of these terms depends on whether it takes a value on some scale which is greater than the minimal value required for counting as satisfying the adjective in the context. One can count as epistemically responsible relative to a context where epistemic couch potatoes are under discussion, but not count as epistemically responsible relative to scientific or legal contexts, or gullible relative to members of the Skeptical Society but not five-year-olds. And epistemological contextualists have argued at length that a belief can count as epistemically justified (or as knowledge) relative to a ‘low-stakes’ context, but not count as epistemically justified (or as knowledge) relative to a ‘high-stakes’ context.¹⁴

¹⁴ See e.g., Cohen (1999) and DeRose (1999). Note that if the standards for counting as satisfying some term vary with context, it doesn’t directly follow that a contextualist semantics is correct for the term in question.

One might object that this semantics fails to generalize properly because some epistemic concepts have no gradable expression. For instance, ‘know’ clearly doesn’t behave syntactically like a gradable expression.¹⁵ This seems to be no mere syntactic fluke. Certain expressions which are syntactically similar to ‘know’ behave like gradable expressions; consider ‘believe,’ ‘regret,’ or ‘trust.’ But this objection has at most limited force.¹⁶ Even if not all epistemic concepts have gradable expressions, and even if this turns out to be epistemologically important, it still fails to track a distinction between thick and thin epistemic concepts. Many typical examples of thin epistemic terms, such as ‘epistemically justified’ and ‘rational,’ are gradable.

This semantics for epistemic terms doesn’t support the claim that a distinction between thick and thin epistemic concepts makes an important difference to epistemology. Terms that express thick epistemic concepts get fundamentally the same semantic treatment as terms that express thin ones: both are analyzed along the lines of (1)-(2). Just as ‘Maria is intellectually curious’ is true just in case the value Maria takes on a scale of intellectual curiosity is greater than the minimal value required for counting as satisfying ‘intellectually curious’ in context, so ‘Jim’s belief that p is epistemically justified’ is true just in case the value Jim’s belief takes on a scale of epistemic justification is greater than the minimal value required for counting as satisfying ‘epistemically justified’

15 For an early case, see Dretske (1981: Ch. 5). For a more recent and complete case, see Stanley (2004, 2005).

16 The objection may in fact have no force at all. Since it takes the fact that certain epistemic terms don’t behave like gradable expressions to show that the appropriate semantics for these terms isn’t scalar, it presupposes that an expression e is semantically linked to a scale s only if e is gradable along s . But the scalar analysis requires only that epistemic terms have semantic links to scales, not that they behave syntactically like gradable expressions. For instance, even if ‘know’ isn’t gradable, it may still be semantically linked to a scale of epistemic strength along which knowledge states can be ordered, and may still bear some other syntactic marks of semantic links to scales, such as adverbial modifiability (see e.g., Blome-Tillman 2008: 43-47). Extending the scalar analysis to ‘know’ would require no more. Whether all non-gradable epistemic terms in fact bear syntactic marks of semantic links to scales is, of course, an empirical question. But at least all of the epistemic terms that are mentioned above are either gradable or adverbially modifiable.

in context. Moreover, this semantics doesn't determine which property something enjoys to a certain extent when it counts as satisfying a given term, but only requires that *there be* a property dimension, along which objects can be ordered, to be associated with the term in context.¹⁷ In this respect of their semantics, terms which express epistemic concepts are treated the same. If terms that express epistemic concepts get the same basic semantic treatment regardless of how thick or thin they are, then the importance of a distinction between thick and thin epistemic concepts cannot be based on their basic semantics.

For all this, epistemic terms might still differ semantically in some respect other than their formal semantics. For instance, their semantic interpretation requires fixing a scale and the values of contextual parameters like the standard. These are set by a computation based on a wide range of inputs from context.¹⁸ So, thick and thin epistemic terms might still exhibit some systematic differences—differences that might have importance to epistemology—with respect to factors which determine the semantic interpretation of the occurrences of these terms.¹⁹

To cut a long story short, among the relevant factors only the conventional meanings of epistemic terms can hope to exhibit

17 Often the property dimension won't be as fully determined by the conventional meaning of the word as with 'tall' (height), 'young' (age), and 'cheap' (cost). Semanticists skirt the issue of what the relevant dimension is.

18 What contextual factors qualify as inputs, what computational rules must be taken into account, and how these must be weighed in semantic interpretation are complex issues grouped together as 'metasemantics.' For a brief overview of the small amount of systematic work done to date on metasemantics, see Glanzberg (2007).

19 Differences in what the occurrences of one and the same epistemic term express in different contexts won't suffice to make this case. For instance, if the contextualist thesis that 'know' is context-sensitive is true, then in principle any factor of the relevant kind can generate important differences in what relation is expressed by different utterances of 'know'; and similarly for contextualism about any other epistemic term. But this is a point about different tokens of *one and the same* epistemic term. It wouldn't support the claim that there are systematic epistemologically important differences between those contextual factors which determine the interpretation of pairs of *different* epistemic terms, such as 'intellectually curious' and 'epistemically justified.'

systematic semantic differences between them.²⁰ (It is plausible that to compute what a sentence expresses in context one should make as much use as possible of the conventional meanings of expressions.²¹) Terms like ‘intellectually curious’ and ‘epistemically justified’ seem thick to different degrees. And this seems to reflect precisely a difference in their conventional meanings, namely in the amount of information which they encode concerning the property dimensions of their associated scales. If I call someone intellectually curious, then competent speakers will take me to have made certain fairly specific claims about the cast of her intellectual character. But if I say that some belief is epistemically justified, then competent speakers can only take me to have made a certain relatively generic claim about the degree of epistemic strength enjoyed by the belief, unless they or the context supply some further, more specific substantive ideas about what epistemic strength is.

But why think that this sort of difference in the conventional meanings of terms like ‘intellectually curious’ and ‘justified’ establishes the importance of a distinction between thick and thin epistemic concepts to epistemology? It is highly contingent that words have the conventional meanings they do. Why should the factors on which conventional meaning is contingent, or even the particular factors which have made certain epistemic terms thicker than others, be factors which generate or even track epistemologically important distinctions?

No doubt there is no easy answer here. But consider this. Saying that someone’s way of forming beliefs is conscientious or careless, for

20 Among other potentially relevant factors, (a) computational rules tend to be insufficiently word specific, (b) discourse structure is highly contingent and may be determined by non-epistemic or epistemically objectionable factors, and (c) the intentions, interests, and other salient properties of speakers and hearers or context which have parameter-setting force can seemingly differ systematically between different epistemic terms only if the roles of these terms in epistemic discourse already differ in ways one would expect their conventional meanings to reflect. This list of factors borrows from Kennedy (2007) and Glanzberg (2007). I don’t claim that it is comprehensive.

21 Kennedy (2007: 36) calls this rule the ‘principle of interpretive economy’ and defends it at length.

instance, doesn't impart very specific information about how they form their beliefs. They might be conscientious in collecting evidence but careless in checking that the samples are representative, or *vice versa*. These claims are similarly indeterminate with respect to many other aspects of belief formation. For all that has been shown so far, these features of carelessness or conscientiousness might be epistemologically important only insofar as they bear on such more general factors as the quality of one's evidence, the reliability of one's beliefs, or the like. Differences in the conventional meanings of epistemic terms of varying degrees of thickness, such as the sort of more specific information which the conventional meanings of terms like 'conscientious' and 'careless' encode in comparison to terms like 'justified,' don't seem sufficiently deep and systematic to establish that a distinction between thick and thin epistemic terms as such makes an important difference to epistemology. I take this as evidence that the importance of this distinction must be established on some other than semantic grounds.

4. The Thick and the Thin in Substantive Epistemology

The moral of my story so far is that, assuming the possibility of making some good distinction between thick and thin epistemic concepts, one can establish that it makes a significant difference to epistemology only through substantive epistemology. My aim in this final section is to argue that the case is yet to be made for anything that could truthfully be advertised as a radical turn to thicker epistemology. My strategy is to argue that there is no good reason to think that thick epistemic concepts enjoy any systematic theoretical priority over the thin. Since the issue concerns whether epistemology has a certain general and systematic kind of theoretical shape, establishing my conclusion requires only that some wide enough range of thick epistemic concepts play no privileged theoretical role. If some individual thick concepts play important roles fully in their own right, this doesn't yet establish any systematic priority or establish that a distinction between thick and thin epistemic concepts as such is important.

One way to defend my claim is to argue that thick and thin epistemic concepts are mostly similar in respects which could systematically make a substantive difference to epistemology. One such respect is that many concepts which play a role in epistemological theory do so only because they either are partly defined or explicated in terms of truth or else bear some other kind of important connection to truth.²² KNOWLEDGE, JUSTIFICATION, and RELIABILITY are often taken to satisfy the former condition. GULLIBILITY and REPRESENTATIVE SAMPLE, among others, are often taken to satisfy the latter. For instance, even many true beliefs of gullible people are connected to truth in a way that many epistemologists regard as defective. Gullible people tend to form beliefs on the basis of factors that have no stable connection to truth. Even many of their true beliefs tend not to be beliefs one would have only if they were true; even if they are in fact true, they could all too easily have been false.²³ This kind of connection to truth seems to be too unstable to make the true beliefs in question count as knowledge. A similar indirect connection to truth underpins the theoretical role of many other concepts in epistemology. So, it seems legitimate to defend or dispute assigning or withholding some concept a theoretical role in epistemology on the basis of whether it bears some important connection to truth.²⁴

22 I mean 'only because' to state a necessary constitutive condition which may not be sufficient.

23 That is, a gullible person's beliefs tend to fail the 'safety' condition on knowledge advocated in Sosa (1999). Even Elgin (2008), in a discussion otherwise aimed against a supposed hegemony of truth in epistemology, argues that TRUSTWORTHINESS bears just this kind of indirect but important counterfactual connections to truth.

24 I have *not* said or implied that the epistemic status of a belief is some direct function of its truth or that having true beliefs is the most important epistemic challenge or aim, let alone the only fundamental one. I can happily allow that not all epistemic challenges concern truth or even knowledge. Even in relation to truth, the challenges of avoiding false beliefs, avoiding beliefs that could all too easily have been false, and so on, are distinct from, but no less important than, the challenge of having true beliefs, but may require different strategies. And clearly the full epistemic status of a belief is only an indirect and partial function of its truth. I also haven't said that the relevant connection to truth is

Many debates in epistemology are further evidence for this claim. Coherence theories of epistemic justification are often criticized on the grounds that coherence isn't truth indicative: a set of beliefs could be highly coherent even if most of them were false or their coherence and truth were only accidentally connected. Basically the same objection is sometimes raised against internalist theories of knowledge or epistemic justification: the fact that a belief has a certain internalistic feature is only accidentally connected to truth and so isn't a good indicator of truth. Also note that opponents of coherentism (or internalism) often agree that the lack of coherence (or the relevant internalistic feature) can weaken or defeat epistemic justification, precisely when beliefs which lack the feature in question are to that extent less likely to be true. Many other debates seem to replicate basically the same structure.

It seems clear that if a particular concept bears some important connection to truth, it does so independently of whether it is thick or thin. So there is no difference in this respect between thick and thin epistemic concepts which could make a substantive epistemological difference.²⁵ The range of thick concepts which bear an important connection to truth also seems wide enough to rule out any systematic theoretical priority of thick concepts over the thin. The natural objection is that these concepts might yet differ systematically in some other respect that *is* important to epistemology. I can offer two points to preempt this objection.

Firstly, thick epistemic concepts which make a difference to epistemology independently of their connections to truth or such paradigmatically thin notions as knowledge seem too few to establish that thick concepts enjoy systematic theoretical priority over the thin.

promoting true belief (so that, e.g., the importance of avoiding being gullible is that this promotes true belief). I only require *some* important connection to truth.

²⁵ Although many thin epistemic concepts seem to be partly defined in terms of truth whereas most thick ones have only some less direct connection to truth, this difference seems insufficiently systematic or categorical to show that the distinction as such makes a substantive and significant difference to epistemology.

Many theories of knowledge incorporate various other epistemic notions that are often cited as thick in the literature. Some hold that the concept of knowledge is used to flag *trustworthy* informants (Craig 1991). Others hold that it is *reasonable* or *appropriate* to rely on a proposition p as a premise in practical reasoning and act on its basis only if one knows that p (Hawthorne 2004, Stanley 2005). Many hold that a belief which might easily have been false doesn't count as knowledge, so that knowledge requires basing one's beliefs on stable and safe reasons. But if one's beliefs regarding some phenomenon are of this kind, then typically one appreciates explanatory connections regarding it in a way that at least partly constitutes *understanding*. I see nothing in these connections to suggest that the thick epistemic notions in question would be theoretically prior to knowledge. What they do suggest is that many contemporary epistemologists already do incorporate various thick epistemic notions in their theories.

Secondly, the claim that thick epistemic concepts don't enjoy systematic priority over the thin fits the structure of epistemological theories better than the claim that they do. For instance, many theories of knowledge attempt to analyze knowledge in terms of belief, truth, plus further, epistemic conditions, such as justification or warrant (and, often, some 'fourth condition' to handle Gettier cases). No such theory will tell us which beliefs count as knowledge until supplemented with a further substantive theory of the conditions in virtue of which a belief is epistemically justified or warranted when it is. Issues at stake include: whether or not these conditions comprise only factors which indicate truth and also are 'internal' in some sense; what the more specific such factors are, such as whether the belief is proportioned to evidence, or coheres with one's other beliefs, or was formed responsibly or reliably, or what have you; and so on. Many of these issues at stake between these substantive theories are formulated in terms that are thicker than

notions of justification, warrant, or knowledge.²⁶ But nothing in this theoretical structure shows that thick epistemic concepts enjoy systematic theoretical priority over the thin.

A comparison with the case of ethics may, in this case, prove illuminating. Theories of moral rightness attempt to specify the conditions under which an action is morally right: maximizing the good, treating no one as a mere means, violating no stringent rights, being what a virtuous person would do, and whatnot. These different substantive theories are structurally on a par, in that none of them will tell us which actions are morally right until supplemented with a substantive account of their central notion: what the good is, what it is to treat someone as a mere means, what it is to violate a stringent right, what a virtuous person is, or whatnot.²⁷ Such accounts are typically formulated in terms of thicker ethical notions.

But a thick ethical concept seems to play this theoretical role only when its role can be characterized partly in thinner terms, in a way that locates and structures the relevance of the thick concept in the ethical domain.²⁸ For instance, even if WELL-BEING is a thick value concept, it plays an important role in ethical theory not only insofar as considerations of well-being are legitimately invoked in ethical assessment, have some connection to motivation, and so on, but also only insofar as they are relevant to what is right or good. Otherwise WELL-BEING would seem to belong together with CHIVALRY and CHASTITY in the theoretical scrap heap.

In the epistemic case, a similar line seems plausible. Thicker concepts like INTELLECTUAL CAUTION and GULLIBILITY seem to play an important theoretical role in epistemology only insofar as they are legitimately invoked in epistemic assessment and have some connection, however

26 It would, however, be worth exploring when accounting for a thin epistemic notion in terms of a thicker one succeeds *because of* the latter's thickness.

27 For this point about the structure of moral theories, see especially Hursthouse (1996).

28 This should seem especially reasonable if there is anything to the idea that value concepts get their content at least partly from the role they play in (some idealized version of) ethical thought. See Jackson (1998: Ch. 4).

indirect, to considerations of rationality, justification, or knowledge. Circumstances in which no such actual connection exists are easy to imagine. If, for instance, we inhabited a highly uniform world, then whether one were intellectually cautious or drew hasty generalizations wouldn't affect the actual reliability of one's beliefs. Neither one's way of selecting the samples from which one generalizes nor their size would have such an effect. Similarly, if we inhabited a world of compulsive truth tellers, gullibility would have no effect on the actual reliability of our beliefs. Given just the actual facts about these worlds, intellectual caution and gullibility have no effect on the epistemic status of one's beliefs. They can count as epistemic defects in these sorts of worlds only insofar as epistemic assessment doesn't depend solely on facts about the actual world but has also some kind of counterfactual dimension.

But note that the relevant counterfactual dimension, and thus the epistemic defect involved in gullibility or lack of caution, is usually analyzed partly in thin epistemic terms. Some theorists say that these beliefs aren't robustly reliable because they would no longer have been reliable had the world been different even in fairly minor ways; others, that they fail to count as knowledge because they were formed without ruling out certain relevant alternatives which were epistemically possible to the agent; and so on.²⁹ If the theoretical role of many thick epistemic concepts in epistemology is best characterized partly in thin terms, then we have no reason to think that thick epistemic concepts enjoy any systematic theoretical priority over the thin. (It doesn't follow that thin concepts enjoy any systematic theoretical priority over the thick. It could be that there is no general priority either way. This would still allow for local priorities between individual thick and thin notions in either direction.)

29 As in the ethical case, these characterizations locate and structure the relevance of a thick notion in the epistemic domain. For instance, they can be helpful in articulating and explaining sources of defeat: reasons why something may fail to possess the kind of general positive epistemic status which is expressed by such thin epistemic concepts as RELIABILITY or JUSTIFICATION.

I conclude that if a distinction between thick and thin epistemic concepts as such makes a significant difference to epistemology, this isn't because thick concepts enjoy some systematic theoretical priority over the thin. The cumulative force of the challenges I have sought to raise for this hypothesis is that it is difficult to see on what other substantive epistemological grounds the distinction as such could plausibly be held to have importance to epistemology, beyond the possibility (which I can happily allow) that some individual thick epistemic notions turn out to have been insufficiently appreciated in substantive epistemological theory.³⁰ Of course, for all I have said, such grounds may yet exist. But finding them remains a challenge to the advocates of a greater focus on thick epistemic concepts. The case has yet to be made for anything that could truthfully be advertised as a radical turn to thicker epistemology.³¹

5. Conclusion

The general position defended in this paper could perhaps be thought of as slim epistemology with a thick skin or exterior. But I suspect that many contemporary epistemological theories may already be thick enough that it would be unsound to charge them with neglecting thick epistemic concepts in favor of the thin in some way that is distorting or too narrow.³² In conclusion, I wish briefly to address this charge against contemporary epistemology.

30 Note that everything that I have said allows that thick epistemic concepts raise theoretically important issues in substantive epistemology. For instance, the issue whether thick epistemic concepts can be informatively decomposed into thin epistemic and purely non-epistemic components need be no less relevant to the prospects for virtue epistemology or expressivist treatments of epistemic discourse, and can be given their full due, even if a distinction between thick and thin epistemic concepts as such makes no significant difference to epistemology.

31 I applaud Elgin (2008) for recognizing that such a case must be made within substantive epistemology. Still, I think her paper fails to expose any deep gap in mainstream epistemology. My case for this is largely implicit in my discussion in this section. But this isn't the occasion to present an explicit account of those reasons.

32 Williams (1985) is the *locus classicus* of this charge in the case of contemporary ethical theories. My argument against the parallel charge I describe in the text against contemporary epistemological theories is parallel to the argument against Williams's charge concerning ethics in Scheffler (1987: 417).

We saw earlier that the distinction between thick and thin epistemic concepts is a matter of degree rather than binary and that many epistemic concepts seem to fall somewhere between such typical examples of thick and thin epistemic concepts as GULLIBILITY and INTELLECTUAL CURIOSITY, and JUSTIFICATION and KNOWLEDGE, respectively. Consider, just as a sample selection, BASING, PROPORTIONED TO EVIDENCE, SELF-EVIDENCE, INDUCTIVE SUPPORT, SAFETY, RELIABILITY, COHERENCE, EPISTEMIC LUCK, and EPISTEMIC RESPONSIBILITY. Are the concepts on this list thick or thin?

If these concepts are thick, then contemporary epistemology is far more concerned with thick epistemic concepts than the wholesale charge against contemporary epistemology allows. After all, many contemporary theories are built around concepts on this list. But if these concepts are thin, then the class of thin epistemic concepts is far more rich and diverse than the typical examples of such concepts suggest and it is far from clear that focusing on thin concepts distorts or oversimplifies. After all, the concepts on this list are familiar currency in ordinary epistemic discourse. And if some of the concepts on the list are thick while the others are thin, that gives some support to each horn of this dilemma for the charge that contemporary epistemology neglects thick epistemic concepts over the thin

It is a complicated issue whether the list of concepts which could be compiled to show that contemporary epistemology is more rich and diverse than some people claim would supply *enough* richness and diversity. That depends, among other things, on what substantive connections obtain between the notions which figure on the list and the notions which don't. It would be premature to make any general pronouncement on these issues here.

But even if such a list of concepts still lacked enough richness and diversity, this would only establish the piecemeal result that certain individual thick epistemic concepts have greater theoretical importance than has previously been appreciated. It wouldn't show that the distinction between thick and thin epistemic concepts as such has the

sort of theoretical importance which could truthfully be advertised as supporting a radical turn to thicker epistemology. Thus, far from checking myself in for rehab with Anorexic Epistemologists Anonymous, I remain yet to be convinced that systematically favoring thick epistemic concepts over the thin would lead to a desirable change in epistemology.³³

University of Leeds

References

- Bach, Kent (1984), 'Default Reasoning: Jumping to Conclusions and Knowing When to Think Twice,' *Pacific Philosophical Quarterly* 65: 37-58.
- Bishop, Michael A. and Trout, J. D. (2005), *Epistemology and the Psychology of Human Judgment* (Oxford: Oxford University Press).
- Blackburn, Simon (1992), 'Through Thick and Thin,' *Proceedings of the Aristotelian Society Supplementary Volume* 66: 284-99.
- Blome-Tillman, Michael (2008), 'The Indexicality of 'Knowledge',' *Philosophical Studies* 138: 29-53.
- Brower, Bruce (1988), 'Virtue Concepts and Ethical Realism,' *Journal of Philosophy* 85: 675-93.
- Cohen, Stewart (1999), 'Contextualism, Skepticism, and the Structure of Reasons,' *Philosophical Perspectives* 13: 57-89.
- Craig, Edward (1991), *Knowledge and the State of Nature* (Oxford: Clarendon Press).
- Cresswell, Max (1977), 'The Semantics of Degree,' in *Montague Grammar*, ed. B. Partee (New York: Academic Press), 261-92.
- Dancy, Jonathan (1995), 'In Defence of Thick Concepts,' *Midwest Studies in Philosophy* 20: 263-79.
- DeRose, Keith (1999), 'Contextualism: An Explanation and Defense,' in

33 Thanks to Ben Kotzee and Jeremy Wanderer for helpful comments and suggestions concerning this paper.

- The Blackwell Guide to Epistemology*, ed. J. Greco and E. Sosa (Oxford: Blackwell), 185-203.
- Dretske, Fred (1981), *Knowledge and the Flow of Information* (Cambridge: MIT Press).
- Eklund, Matti (MS), 'Some Problems Regarding Thick Concepts,' unpublished typescript.
- Elgin, Catherine Z. (2008), 'Trustworthiness,' this volume.
- Gibbard, Allan (1992), 'Thick Concepts and Warrant for Feelings,' *Proceedings of the Aristotelian Society Supplementary Volume* 66: 267-83.
- Gigerenzer, Gerd *et al.* (2000), *Simple Heuristics That Make Us Smart* (Oxford: Oxford University Press).
- Glanzberg, Michael (2007), 'Context, Content, and Relativism,' *Philosophical Studies* 136: 1-29.
- Greco, John (2008), 'What's Wrong With Contextualism?,' *Philosophical Quarterly* 58: 416-36.
- Hawthorne, John (2004), *Knowledge and Lotteries* (Oxford: Oxford University Press).
- Hursthouse, Rosalind (1996), 'Normative Virtue Ethics,' in *How Should One Live?*, ed. R. Crisp (Oxford: Clarendon Press), 19-36.
- Jackson, Frank (1998), *From Metaphysics to Ethics* (Oxford: Clarendon Press).
- Kennedy, Christopher (2007), 'Vagueness and Grammar: The Semantics of Relative and Absolute Adjectives,' *Linguistics and Philosophy* 30: 1-45.
- Klein, Ewan (1980), 'A Semantics for Positive and Comparative Adjectives,' *Linguistics and Philosophy* 4: 1-45.
- Quine, W. V. O. (1969), 'Epistemology Naturalized,' in *Ontological Relativity and Other Essays* (New York: Columbia University Press), 69-90.
- Scheffler, Samuel (1987), 'Morality Through Thick and Thin,' *Philosophical Review* 96: 411-34.
- Schwarzschild, Roger and Wilkinson, Karina (2002), 'Quantifiers in Comparatives: A Semantics of Degree Based on Intervals,' *Natural*

Language Semantics 10: 1-41.

Sosa, Ernest (1999), 'How to Defeat Opposition to Moore,' *Philosophical Perspectives* 13: 141-53.

Stanley, Jason (2004), 'On the Linguistic Basis for Contextualism,' *Philosophical Studies* 119: 119-46.

Stanley, Jason (2005), *Knowledge and Practical Interests* (Oxford: Oxford University Press).

von Stechow, Armin (1984), 'Comparing Semantic Theories of Comparison,' *Journal of Semantics* 3: 1-77.

Williams, Bernard (1985), *Ethics and the Limits of Philosophy* (Cambridge: Harvard University Press).