The Thin/Thick Dilemma

Andrew Sepielli

Introduction

Ethics could not be an objective domain unless some atomic ethical concepts applied to some things in the world. There are two fundamental ways of calling this concept-world relationship into doubt. The first is by arguing that ethical concepts couldn't apply to items in the world because of something about the world; the second is to claim that ethical concepts couldn't apply because of something about the concepts (or, if we suppose that facts about concepts are explanatorily posterior to facts about thinkers and their thoughts — because of something about us.) The former way is characteristic of what Huw Price labels "object naturalism", while the latter gives expression to what he calls "subject naturalism".[1]

Object-naturalist doubts are well-represented in the literature. We are told that ethical concepts could apply only if there were ethical properties, and that it is not clear how the world could contain such properties. It's claimed that we don't need ethical properties to explain our thoughts, or anything else for that matter; that such properties would have to be both motivation-independent and prescriptive, but that

_____

[1] See Price (2011).

1

this is impossible; that the best theories of the successful natural sciences don't mention ethical properties; that there's no way of reducing ethical properties to natural-scientific ones, and that absent such a reduction, we should be unwilling to countenance the former's existence.

Some very thoughtful realists and would-be realists lose sleep over arguments like these, but I can't say I'm among them. My fundamental concerns about ethics are quite different; they may be captured in the form of a dilemma, the first horn of which is an argument of the subject-naturalistic sort. This argument locates the mystery of how concepts like [wrong] could apply to worldly items in the thinness of these concepts' constitutive structures. Such concepts, the story goes, can't be defined in terms of non-ethical concepts, nor do they incorporate non-ethical concepts in a prototype or exemplar structure, nor do they have any constitutive links with raw experience, and so on. There is nothing about these concepts that we might use to explain when they apply and when they don't. There is no plausible theory of mental representation that can account for how thinkers can get the world right in their employment of these concepts. Or so I plan to argue.

It is natural to seek a solution in concepts with thicker constitutive structures. Of course these include so-called "thick concepts" like [courageous]. Perhaps they also include concepts like [wrong] — that is, perhaps my argument rests on underestimating the robustness of such concepts' constitutive structures. The envisaged solution is that we use the added elements of these thicker structures to explain when these concepts apply and when they don't.

But this solution seems to me to trade one difficulty for another, and thus impale

2

us on the dilemma's second horn. For even if we could explain how thick concepts apply, it would remain a mystery how their application could ever settle an ethical conflict — how an action or attitude's having the properties they designated could ever render that action or attitude (as I shall put it) "mandatory". This is because, for any thick ethical concept that guides my doing an action, there are what I shall call its "anti-concept" and its "counter-concept", which would guide me not to do the action, or to do an action in its complement set, respectively. I shall argue that there is no explanation (of a certain desired sort) for privileging any thick concept over its anti- and counter-concepts.

Finally, I claim that ethics can be an objective domain only if there is at least one ethical concept that applies to things such as to make them mandatory. But if the foregoing is correct, it is not clear how there could be such a concept. We can explain, perhaps, how some ethical concepts can apply to things in the world. And we can explain, perhaps, how some concepts could render things mandatory, if only they applied to those things. But we have no explanation of the desired sort of how any one concept can do both. So we have no such explanation of how ethics could be objective.

I call this argument "The Thin/Thick Dilemma". It is hinted at in chapter eight of Bernard Williams's Ethics and the Limits of Philosophy. Something like its "thin" horn has been presented by expressivists like Terry Horgan and Mark Timmons.2 Additionally, it recalls the objection to constitutivist theories of practical reasons that whatever conception of action they adopt will either be too thin for action to have the alleged constitutive aims, or else too thick for action to be unavoidable in the right

---

2   See their (1992).

sense.3 But I have not seen exactly the argument I'm labelling the "Thin/Thick Dilemma" presented explicitly and in its entirety.

I. The Thin/Thick Dilemma: a First Pass

The basic argument is as follows:

1. Every ethical concept is either thin or thick.

2. If a concept is thin, then its constitutive structure cannot provide an explanation of the concept's conditions of correct application.

3. If an ethical concept's structure is insufficient to explain its conditions of correct application, then there is no (other) external explanation of these conditions.

4. If a concept is thick, then there is no external explanation of how something's having the feature the concept designates could ever render doing/having that thing "mandatory".

5. For there to be an explanation of the objectivity of ethics that is sufficient for certain fundamentally important purposes, there would have to be at least one concept such that:

_____

3   See, e.g., Enoch (2011a). I discuss this other "thin/thick dilemma" later on.

(a) there is an external explanation of its conditions of correct application, and

(b) there is an external explanation of how something's having the feature it designates   could ever render doing/having that thing mandatory.

6. Therefore, there is no explanation of the objectivity of ethics that is sufficient for certain   fundamentally important purposes.

Some comments on the argument before I go on to defend its premises:

It should be clear that the notion of an "external" explanation is crucial. An internal explanation of a concept's application conditions appeals to purported facts within the domain to which that concept belongs. An external explanation is one that is not internal; it is what John McDowell would call a "sideways-on" account.4

If I were to explain why the concept [evil] applies to racism by citing the fact that racism is evil, that would be an internal explanation, appealing as it does to an ethical fact to explain the application of an ethical concept. Even an explanation deepened by a moral theory of why racism is evil would count as internal, since such a theory is nothing other than a purported ethical fact (or set of such facts). Similarly, an explanation of [evil]'s application that adverted to the virtuousness or moral sensitivity of the concept's users would be internal, since claims about virtuousness and moral sensitivity also belong to the ethical domain. Basically, if you're doing first-

---

4   See McDowell (1994), esp. Lecture IV.

order ethical theory, at whatever level of abstraction or sophistication, you're in the business of offering internal explanations.

An external explanation of concepts' application conditions is the sort of thing many philosophers are trying to provide through domain-general theories of mental representation. Here is a programmatic explanation of this sort: "Any non-conceptual sensory representation of X that has properties P1...Pn is veridical. A thought that applies concept C to X is true if it is formed based on a veridical experience, in accordance with the correct rule of rational conceptualization. It is possible to form a thought that applies [man] to Barack on the basis of a non-conceptual sensory representation of Barack that has properties P1...Pn, in accordance with the correct rule of rational conceptualization. So [man] applies to Barack." Note that no facts belonging to the domains of biology or gender theory or Presidential history are being marshalled to explain why Barack falls under [man]; that is why this is not an internal explanation.

I take no official stand on whether conditions 5(a) and 5(b) must be met for ethics to be objective. That is, for all I argue, ethics may be brutely objective — objective without there being any explanation of how such objectivity is possible. And I take no stand on whether 5(a) and 5(b) must be met for there to be an explanation of ethical objectivity that achieves some worthwhile purpose or other. That is, maybe internal explanations are good for something. I claim only that there are some purposes for which we need an external explanation, and that some of these are of fundamental importance.

So-called "quietist" philosophers like McDowell deny this. They deny that we

need external explanations for any such purposes, and suggest that our hunger for such explanations betrays an anxiety that is best treated with "therapeutic" philosophy. I take this position seriously — indeed, I lose sleep over it! — but unfortunately I lack the space in this paper to give it a full hearing. But I will at least try to say something about why I think we need external explanations, and thereby place a burden on those who affirm objectivity without external explanations to justify their stance. Specifically, I will argue that external but not internal explanations (a) ensure against what I shall call "representational miracles", (b) unify ethics with other apparently objective domains, and (c) are dialectically effective against a range of opponents who we should be — and as social beings, must be — concerned to address. I will then say a bit about why I am unconvinced by some of the purported grounds for dispensing with external explanations.

   I should also say that I don't regard my insistence on external explanations as exceptional. As I noted earlier, a chief aim of philosophers who offer general theories of mental or linguistic representation — e.g. Jerry Fodor, Ruth Millikan, Christopher Peacocke, and Robert Brandom — is to provide external explanations of how and when concepts apply. To be sure, some philosophers of a quietist bent may doubt the necessity or possibility of developing such a theory,5 but they do not constitute the majority. In the general "objectivity" literature, Crispin Wright suggests as a condition on the objectivity of a domain that it obey what he calls "cognitive command", which it does just in case it is a priori that any difference of opinion in that domain is

_____

5 See especially McDowell's (1987); but it is also worth reading Dummett's anti-quietist response in

   the same volume.

7

attributable to at least one party's "cognitive shortcoming". But, importantly, he and most others in this literature go to great lengths to try to avoid what they call "trivializing" cognitive command, by which they mean imputing a shortcoming to a person simply on account of her accepting P when ~P is the case. What they have in mind by a "non-trivial" characterization of cognitive command is one on which a cognitive shortcoming is characterized in a domain-external way. Additionally, it is clear that many ethicists have theoretical ambitious that I claim cannot be accomplished solely through internal explanations. For example, constitutivists like Christine Korsgaard to Jurgen Habermas are explicit about wanting to convince skeptics and other radical opponents, while realists from Richard Boyd to Ralph Wedgwood are concerned to assimilate their accounts of how ethical thoughts can get the world right to accounts of how thoughts generally can get it right (a causal account in Boyd's case, a conceptual-role account in Wedgwood's).

Most significantly, it seems to me that almost every other argument for ethical skepticism also rests on assumption that only external explanations are fully satisfactory. For consider how many such arguments, whether from causal impotence, the aetiology of our ethical judgments, the persistence of moral disagreement, or the mysteriousness of moral supervenience, could easily be shown to be unsound, if only we could exploit first-order ethical claims in so showing. For instance, rampant moral disagreement is no threat if we can, on the basis of our own ethical judgments, classify some parties to such disagreement as having better-functioning moral sensibilities or better "track records" on moral questions than others. And the charge that natural selection is unlikely to produce true ethical beliefs looks weak once we entitle

ourselves to claims of the form: "But natural selection produced beliefs in ethical propositions P1...Pn, and lo and behold, P1...Pn are all true!" The proponents of these skeptical arguments obviously don't think them unsound, though, because they think there's something illicit about defusing an argument against ethical objectivity by drawing on ethical judgments. I'm just trying to make this demand for external explanations explicit, so that it and the reasons for it may be held up for scrutiny.

Finally, let me say a few words about the notion of mandatoriness. I use "mandatory" as a term of art, not equivalent in meaning to "obligatory". It is a notion defined in terms of the quality of an agent's functioning. An action or attitude is mandatory if and only if anyone who is in a position to do it/have it and does not must therefore be, ceteris paribus:

(a) at a cognitive deficit relative to[6], or

(b) less rational than[7],

---

[6] S is at a cognitive deficit relative to T just in case S is worse with respect to believing truths than T is. S may be worse with respect to believing truths than T is by believing fewer truths than T does, by *knowing* fewer truths than T does, perhaps by having a worse *understanding* of those truths, perhaps by believing less *important* truths than T believes, and so on. These are difficult issues in epistemic value theory, although as you'll see, we won't need to delve very deeply into them for our purposes. My later claims about cognitive deficits should be uncontroversial. The usefulness of this "better with respect to X" locution was impressed upon me by Temkin (2012).

[7] I mean "rational" in its "internal" or "structural" sense. See, e.g. Broome (1999), Kolodny (2005) and Scanlon (2007).

a person who does do it/have it. Clearly, the fact that an action is obligatory renders it mandatory. But so does the fact that it is what the agent has most reason to do. And for all I've argued so far, so might the fact that it is at the same time the most just and most compassionate action. (I shall argue later, however, that this fact cannot render an action mandatory.)

I will say a good deal in Section V about why ethical objectivity requires mandatoriness. But the basic idea is this: Many of us care about ethical objectivity. We may care whether ethics is objective in the sense that we see it as making a practical difference whether ethics is objective. We may care that ethics is objective in the sense that we long for such things as objective ethical values. Many of these cares are implicit; if we are duly self-conscious theorists, then some are explicit. We should adopt a conception of ethical objectivity on which it makes a practical difference whether ethics is objective, and on which ethical objectivity is an object of our longing — in short, a conception that offers us what we care about when we care about ethical objectivity. And I claim that this is offered only by a conception of ethical objectivity as requiring mandatoriness.

In holding that ethical objectivity requires mandatoriness, I leave open the question of what else it requires, since the answer is irrelevant to the argument of this paper. I also leave open the question of what other sorts of objectivity — e.g. aesthetic, epistemological, — require something like mandatoriness, although I touch on it briefly in Section 5. Finally, it is worth noting that my partial conception of objectivity is by no means proprietary to realist views of ethics. As sophisticated expressivists like Simon Blackburn, Allan Gibbard, and Terry Horgan and Mark Timmons have gone to great

10

pains to show, their views are amenable to certain conceptions of ethical objectivity; among these conceptions, it seems to me, is the one on offer in this paper.

II. In Support of Premise 2: If a concept is thin, then its constitutive structure cannot provide an explanation of the concept's conditions of correct application.

I will say that a thin ethical concept's constitutive structure includes no more than:

(a) its cognitive links with other ethical concepts,

(b) its referential link with an ethical property, either in the actual world or in some possible or conceivable world; and

(c) some bit of conceptual role other than a cognitive or referential link.

I do not rule out that its structure may include less than this, either in the sense that it may include only some of the elements listed above, or in the sense that it may not have an element-constituted structure at all. Nor do I rule out that one or more of (a), (b), and (c) may be explained by one or more of the others. (I find it particularly plausible that (b) must be explanatorily posterior.) But in what follows, I will proceed as though (a), (b), and (c) are equally fundamental aspects of a thin concept's structure, each equally capable in principle of serving as the ultimate ground of the concept's application conditions. This gives my opponent the best chance of explaining thin concepts' application conditions on the basis of their structures, and as such, the

best chance of falsifying this premise.

Relatedly, it may be that some thin concepts are richer with regard to (a), (b), and (c) than others. For example, some thin concepts have a larger array of cognitive links with other ethical concepts than other thin concepts have. Some thin concepts have more and more robust roles of the sort classified under (c). Any concept with more than (a), (b), and (c), however, is thick, with some such concepts being thicker than others.

I draw the "thin/thick" line where I do because, given the argument to come, that provides us with the most efficient way of labelling things. If one were to insist on drawing the line elsewhere — such that, say, some of the concepts I count as "thick" are counted instead as "thin" — then I could simply re-word this paper, and say that what I now call the "thin" horn in fact applies to only the thinner thin concepts, while what I now call the "thick" horn in fact applies to the thicker thin concepts along with the thick ones. But then this paper would have to be called "The Thinner-Thin/Thick-and-Thicker-Thin Dilemma", and that would amount to lingual abuse.

The notion of "cognitive linkage" in (a) is a capacious one. It does not assume the Classical Theory of concepts, on which conceptual structure is represented by necessary and sufficient conditions — i.e. a definition. And the notion is broader than that of, say, an inferential link. Neither analogical reasoning nor pictorial reasoning is inference, exactly, but each is cognitive nonetheless. To say that two concepts are cognitively linked is to say that there are features of either concept that make rational at least one mental transition from a thought in which one is tokened to a thought in

which the other is tokened. So two concepts may be cognitively linked if one may be analyzed in terms of another, or if one stands for a property that figures in the prototype or is possessed by the exemplar of another, and so on.8

In light of the cognitive links mentioned in (a), we will need to sharpen Premise 2. For it would seem that such links furnish us with straightforward explanations of when thin concepts apply. Surely it is built into the structure of either [ought] or [a reason] that if you ought to do something, there's a reason to do it, in which case we can explain quite easily on the basis of conceptual structure what you have a reason to do — specifically, you have a reason to do what you ought to do! So Premise 2 is better stated as: "If a concept is thin, then its constitutive structure is insufficient to explain its non-ethically-described conditions of correct application." In other words, the structures of thin concepts can't take us outside the ethical "circle" or "web". This, it seems to me, better captures the spirit behind ethical skepticism than does the original Premise 2.

I should emphasize that in asking for "non-ethically-described conditions of correct application", we are not asking specifically for a definition of ethical concepts in terms of non-ethical ones. As we will see throughout this section and the next, there are many ways that a concept's structure (broadly understood) might explain its application conditions other than by incorporating those application conditions as part of a definition. To say that a concept's structure explains its application conditions

---

8   For influential discussions of conceptual structure and the reasoning it enables, see Rosch (1978), Smith and Medin (1981), Lakoff (1987) the essays in Laurence and Margolis (1999), Stenning (2002), Machery (2009) and works cited *infra*.

under a non-ethical description is simply to say: (1) that there is a true statement of the form "Ethical concept [E] correctly applies to a just in case C1...Cn", where "C1... Cn" is a non-ethical description; and (2) that [E]'s structure — again, understood broadly as illustrated above — explains the truth of that statement.

With Premise 2 sharpened, we can see that the "(b)" element won't be useful in defeating it, either. For while it is quite plausibly constitutive of an ethical concept that it functions as to pick out an ethical property, this won't help us show how this concept applies to the actual world as non-ethically described. The concept [witch], after all, functions to pick out witches, but it does not pick them out in the actual world, where there are no such things. We would need some further account of why the ethical depends on the non-ethical in the way that it does. This dependence need not be something as metaphysically robust as identity or reduction, but whatever it is, it must be accounted for somehow.

And so it seems that he falsification of the new Premise 2 will require that the "(c)" feature ground thin concepts' conditions of correct application. We can't know whether it can do this without knowing more about "(c)": What might this extra bit of conceptual role be?

I claim that it is some sort of guiding role. What makes a concept an ethical one in my sense is that its deployment in first-personal deliberative contexts generally disposes us towards, or has the function of disposing us towards, or renders rational or intelligible (there are different ways of going here) the action, emotion, or other item to which the thinker applies it. In other work, I defend a view about the precise nature

of this guiding role.9 But for the purposes of this paper, it doesn't matter which

articulation of this relatively uncontroversial thesis we go with.

There are reasons to think that this guiding role cannot ground a concept's

application conditions. Following Sellars (1954), the first major work in Conceptual

Role Semantics, we may think of conceptual roles as specifying three sorts of

transitions: "(intra-conceptual) moves", "entry transitions", and "departure

transitions". Some examples of these, respectively: It is constitutive of the concept

[bachelor] that the thought that X is an unmarried male of a marriageable age renders

intelligible, or guides, the formation of the thought that X is a bachelor. It is

constitutive of the concept [spheroid] that a non-conceptual visual representation of a

spheroid renders intelligible the thought that there is a spheroid before me.10 And as

we have said, it is constitutive of the concept [ought] that the thought that I ought to

do something renders it intelligible that I do that thing.

It is arguable that we may derive some concepts' application conditions from

their constitutive entry transitions, or from the entry transitions constitutive of other

concepts plus some intra-conceptual moves.11 Of course, such a derivation is no piece

of cake. Nobody so far has offered an uncontroversial theory of conceptual

---

9   [OMITTED]

10   Whether *Sellars himself* would accept this as an example of a rational entry transition is, for familiar

     reasons, open to question. See his (1956).

11   Arguable, namely, by those in the quite catholic Conceptual Role Semantics tradition -- e.g. Sellars

     (1954), Field (1977), Block (1986), Harman (1987), Peacocke (1993), and Horwich (1998).

representation. But if thin concepts' structures specified entry transitions or inter-conceptual links (with non-ethical concepts, that is), we would at least be in as good a position to explain their application conditions as we are to explain those of, say, "middle-sized dry good" concepts: [table], [cup], etc. However, since the "guiding role" specifies a departure transition, it is not clear how it could ground the conditions under which a thin-concept-user is getting the world right. If you will: entry transitions are about what is taken up from the world, and inter-conceptual moves are about how this uptaken stuff flows through the mental pipes. Departure transitions are about what is expelled back into the world, and so it is mysterious how they could be used to explain when a person's attempted uptake of the world has succeeded and when it has failed.

The burden, then, seems to be squarely on the person who wants to explain thin concepts' application by appealing to their constitutive structures, rather than the person who doubts that this can be done.

I can think of two strategies by which one might try to meet this burden. One is to tell a story on which guiding roles are capable of explaining concepts' application conditions. The other is to argue that there are aspects of thin concepts' structures in addition to the ones I've claimed, and that these additional aspects can help to explain the concepts' application conditions.

It is easiest to get a feel for the former by first considering how we might explain the application of concepts that are characterized by entry transitions and inter-conceptual moves. It seems as though the way to explain these concepts' application is

to advert to the correctness conditions of the states that either render their tokening intelligible, or whose formation their tokening renders intelligible. If [spheroid] is partly characterized in terms of the entry transition mentioned earlier, then we would explain [spheroid]'s application by adverting to the veridicality of a non-conceptual representation of a spheroid. Or we can explain why [bachelor] does not apply to a thing by adverting to the fact that [married] does apply. Of course, these explanations will prompt further questions — "What makes the non-conceptual state veridical?", "In virtue of what does [married] apply?" — but they seem fine as far as they go (however far that is).

The corresponding treatment of [ought] would have us explain its application by adverting to correctness conditions of the action that an ought-thought renders intelligible. [Ought] is correctly applied if and only if the action that it renders intelligible is correct. Now, it does not make much sense to speak of actions as being veridical or true, and so it seems that we are left with two views about when an action is "correct". The first is that an action is correct insofar as its performance bears some law-like relation to the fulfillment of our aims. This is just the notion of a "successful action" that shows up in the contemporary literature on success semantics,12 and of course, the work of the classical pragmatists like James and Dewey. (This is why Dewey eventually stopped calling himself a pragmatist and started calling himself an "instrumentalist".) The second view is that an action is correct insofar as it meets some substantive standard that is independent of our aims. This is the notion of correctness that Ralph Wedgwood has in mind when he argues that: (a) thinking you

_____

12 See, e.g. Whyte (1990).

ought to make it the case that P rationally commits you to making P part of your "ideal plan", and so (b) the ought-thought is true if and only if that plan is correct.13

Both glosses on "correct action" are reasonable enough. But neither provides the materials for a vindication of ethical objectivity. If a correct action is conceived of in the first way, then the correctness of actions will depend on the agent's aims. This would imply that the proper application of [ought] depends on the agent's aims. But objectivity seems quite clearly to be incompatible with aim-dependence of this sort. (For what it's worth, my yet-to-be-discussed requirement of mandatoriness rules out aim-dependence.) And if correctness is conceived of in the second way, then the characterization of an action as correct is itself an ethical characterization. But the new formulation of Premise 2 states only thin concepts' structures cannot ground non-ethically-described conditions of correct application. One cannot defeat Premise 2 by citing a constitutive link between the correctness of ought-thoughts and the correctness of action, and then giving an account of correct action that is, in its essence, nothing more than an internal explanation of when ought-thoughts are correct.

The second strategy for defeating Premise 2 is to posit richer structures for thin concepts than I've allowed, and then to argue that we can use the elements of these richer structures to explain when such concepts apply. There are many ways of carrying this out. The most extreme way is through simple definition. For example, we might try to define a right action as one that maximizes utility. The popular writer

_____

13 Wedgwood (2007), p. 97.

Sam Harris takes essentially this route in his bestseller The Moral Landscape. A less extreme way is to posit an exemplar or prototype structure that links ethical concepts with non-ethical ones. We might think of the structure of a concept like [wrong] in terms of a set of paradigmatically wrong actions (e.g. killing, stealing, lying...), and a similarity metric that specifies how similar an action must be to one or more of these paradigms to count as wrong. Or we might think in terms of a set of paradigmatically wrong-making features, and a metric that specifies how many of these features an action must have, and to what degree it must have them, to count as wrong. In forthcoming work, Terrence Cuneo and Russ Shafer-Landau pursue an idea along these lines, using the term "moral fixed points" to designate more or less what theorists of concepts call "exemplars".[14] A much more modest route is to think of an ethical concept's structure as serving to rule out items from its extension rather than to rule them in. This seems to be Philippa Foot's proposal when, in her seminal "Moral Beliefs", she suggests that "No one can call a man good because he clasps and unclasps his hands".[15] While Foot by no means offering a definition, she is making a (quite modest) claim about what I would call conceptual structure, at least if we assume a relatively tight connection between the intelligibility of the speech-act of "calling" someone "good" and structure of the concept [good].

Now, one way to attack any of these proposals is to claim that the "additional" elements of structure are not, in fact, sufficient to determine thin concepts' application conditions. But there is no way to make this charge stick without considering each of

---

[14] See Cuneo and Shafer-Landau (ms). On "exemplar theory" in general, see Medin et. al. (1982).

[15] See Foot (1958).

the proposals in greater detail than I wish to do here. Rather, my response to strategies of this sort is a classificatory one. On my scheme, an ethical concept with structure beyond the (a), (b), and (c) elements mentioned above counts as a thick concept rather than a thin one. That it may be expressed with an apparently thin term like "good" or "right" does not disparage this fact. Again, some with different theoretical goals may wish to classify concepts differently; I would respond to this re-classification with nothing other than a re-classification of my own: I would say say that Premises 2 and 3 apply to all thin concepts as defined earlier but either to only some or to no thin concepts on this new understanding; however, Premise 4 applies to all of those concepts, thin or thick, to which Premises 2 and 3 do not.

III. In Support of Premise 3: If an ethical concept's structure is insufficient to explain its conditions of correct application, then there is no (other) external explanation of these conditions.

The most obvious way of rejecting Premise 3 would be to argue that thin concepts' application conditions can be explained by the structures of non-ethical concepts. Suppose, for simplification's sake, we hewed to the Classical Theory of concepts. Then, instead of defining [wrong] partly in terms of [autonomy] or [utility] or some such, we might define [autonomy] in terms of [wrong]: it is part of the very definition of autonomy, maybe, that it is prima facie wrong to restrict or circumvent it. Mutatis mutandis for all of the other theories of concepts.

But this strategy seems unpromising. [Autonomy]'s constitutive structure

doesn't include [wrong], and it is joined in this regard by every other concept we'd designate as non-ethical. If we defined autonomy as something it is prima facie wrong to restrict, then we'd have to say that an ethical nihilist could not both (a) fully grasp the concept of autonomy, and (b) think that there could ever be restrictions of autonomy. But this is on the face of it, absurd.

Now, some philosophers are fond of the idea that a virtuous person could not fully grasp certain non-ethical concepts while refusing to apply certain ethical ones.16 I needn't deny this. I would claim only that this story, relying as it does on the idea of virtue, provides an internal rather than an external explanation of the ethical concepts' application conditions. But that is not the sort of explanation needed to defeat Premise 3.

There is a related way of rejecting Premise 3 that is more subtle. Perhaps we can explain thin concepts' application conditions by appealing to thick ethical concepts. We might say, for example, that it if [unjust] applies to an act, then there is reason not to do that act, or that if [callous] that if it applies to a person, that person's character is thereby worse.

We cannot dismiss this strategy as quickly as we did the previous one. For as I'll explain later, it is highly plausible that thick ethical concepts are cognitively linked with thin ones. And if we can further give an external explanation of how thick concepts apply to the world, we've got all of the components of a fully external "two-

---

16 E.g. McDowell (1979). Again, *he* would not be troubled by the fact that his is not a fully external explanation.

stage" explanation of how thin concepts apply to the world.

But I do not think this strategy will succeed. It founders on what I shall call the "Problem of Explanatory Symmetry". Before we see how, we'll first need to distinguish between two ways of thinking about thick concepts. The first way is to think of them as, basically, thin concepts plus some further structural ingredient that makes them thick. Assuming this picture, a thick concept's structure consists in:

(a) its cognitive links with other ethical concepts;

(b) its referential link with an ethical property, either in the actual world or in some possible or conceivable world;

(c) one or more "thickening elements" — features that make it thick (e.g. cognitive links with non-ethical concepts — which, again, may involve a prototype or exemplar structure rather than a definitional one), and:

(d) what I shall call "valences", which may consist in the incorporation of thin ethical concepts, or the incorporation of one or more non-inferential roles like the "guiding role" discussed earlier. It is not hard to pick up the idea of a valence from examples. We typically think of [courageous], [chaste], and [just] as having a positive valence, and [cowardly], [lewd], and [unjust] as having a negative one.

Now, although this "thin-plus" way of thinking about thick concepts is "reductionist" in a rather straightforward sense, it is compatible with many claims about thick concepts that so-called "non-reductionists" have made. It is compatible with the suggestion made by Jonathan Dancy, Simon Kirchin, and Debbie Roberts that

22

thick concepts may have multiple or variable valences.17 It is compatible with the idea that what I call the "thickening element" is nothing more robust than what Dancy calls its "domain of operation".18 It is compatible with the claim that thick concepts are not composed of elements to which they are posterior in the order of understanding; for just as a cognitive link with [colored] is constitutive of [red] despite [red]'s neither being composed of nor posterior to [colored], the concept [courageous] may be cognitively linked with a non-ethical concept without being, in any intuitive sense, composed of it or posterior to it. It is compatible with the view that the various elements, including the valence, interact so as to explain the concept's application conditions (rather than, say, the "thickening elements" doing all the work). And it strikes me as compatible with the "shapelessness" hypothesis on most understandings thereof,19 both in virtue of the interaction effects just mentioned, and because of the possibility that (c) takes the form of prototype or exemplar structures.20

---

17 Dancy (1995), Kirchin (2013), Roberts (2013).

18 Dancy (1995), p. 277.

19 For a helpful distinction between stronger and weaker interpretations of this hypothesis, see Kirchin (2010).

20 On Vayrynen (2012)'s characterization, for example, a thick concept is shapeless with respect to other concepts if the items in thick concepts' extensions are not "unified under relations of real similarity that are intelligible without the use of that concept". It is well-known that these non-Classical structures are capable of generating extensions for a concept that are not so unified. Indeed, the dissimilarity of the items in concepts' extensions is often cited *in support of* the prototype and exemplar views and against the Classical view.

The other way of thinking about thick concepts is as unstructured — or at least, as not having a structure in which there are separate (c) and (d) elements (i.e. "the part that makes it evaluative" and "the part that makes it thick"). On this picture, the deployment of a thick concept is not evaluation plus non-evaluative description; it is simply a different kind of evaluation — a kind that can't fully be specified in terms of the sorts of structural features just mentioned.

If we hew to the first, "thin-plus" picture, then I shall want to claim that, for any set of thick concepts C1...Cn with valences V1...Vn and thickening elements E1...En, there is

1) A set of anti-concepts that also have E1...En but have opposing valences; and
2) A set of counter-concepts that also have V1...Vn but have complementary thickening elements.21

------

[21] When I say that there "is" an anti-concept for every concept, I do so under the understanding that a concept may exist even if it is not expressed in language, even if it is defective (like [tonk] or [boche]), even if it applies to nothing (like [witch]), and even if it is inconceivable that it could apply to anything (like [square circle]). Put simply, I will say that there is a concept, or idea, just in case there *could be* a thought with a content partly constituted by that concept or idea. I consider this a terminological point, and I am just as comfortable re-characterizing my general point as "For any actual thick concept, there is no satisfactory explanation of why its *possible* or *conceivable* anti-concept or counter-concept does not *actually* apply."

For example, suppose [chaste]'s constitutive structure is such that: (a) it has a positive valence, and (b) it has as its prototype applicanda dressing in body-obscuring clothing and refraining from sex. Then there is also a concept [anti-chaste] -- perhaps heretofore unemployed, perhaps common as corn -- with a negative valence and the same prototype applicanda. And there is a concept [counter-chaste] with a positive valence just like [chaste] but with prototype applicanda of dressing in body-revealing clothing and indulging in sex.

While the terms "anti-concept" and "counter-concepts" are neologisms, my intention in introducing them is to sharpen a familiar idea. We know that different cultures traffic in different thick concepts. Some cultures' conceptual schemes encode a generally negative stance towards violence and an indifference towards glory in battle; others endorse a very positive stance towards glory in battle, and an indifferent or even positive stance towards violence. In Ethics and the Limits of Philosophy, Williams claims that this cultural variability makes some sort of trouble for moral objectivity, but he doesn't go into much detail. My aim in this paper has been to fill in that detail. Ultimately, I think sort of argument at which Williams was gesturing works only if there are (conceivable) conceptual schemes that differ precisely and in totality. So, for example, while we might say that [chaste] has something of an anti-concept in the modern-day [sex-positive], it is not precisely an anti-concept, and it is far from clear that every familiar thick concept has an anti- or counter-concept that has actually found expression within a culture. So ultimately, I think the right way to develop a Williams-style argument involves looking at conceivable conceptual schemes, involving anti- and counter-concepts, rather than at actual conceptual

schemes, involving rough approximations of the same.

With the apparatus of anti- and counter-concepts on the table, we can present what I'd called the "Problem of Explanatory Symmetry". First, let us assume that "thin-plus" picture of thick concepts. On this picture, thick concepts have the constitutive aspects listed above, and I claim that, so far, we have no reason to suppose that these aspects cannot be mixed and matched like doll parts. That there are these interchangeable parts makes it difficult to see how these concepts can be used to explain the application conditions of thin concepts. For consider a toy example. Suppose one wanted to explain on the basis of the constitutive structures of [just], [efficient], and [honorable] how, if these concepts all apply to an action A, then so does the thin concept [what one ought to do all things considered]. I claim that if we have grounds for accepting such an explanation, then we also have grounds for accepting that the structures of [anti-just], [anti-efficient], and [anti-honorable] can explain how, if these concepts apply to A, then so does [what one ought NOT to do]. (Mutatis mutandis for [counter-just] etc, [what one ought to do], and a complementary action ~A.)

But of course it can't be the case that one both ought and ought not to do A all things considered, nor can it be that each of A and ~A is the (single) thing one ought to do. Explanations have to be factive, though, in the sense that if there is a satisfactory explanation of P, there can't be a satisfactory explanation of anything that is inconsistent with P. So we must conclude that the original purported explanation in terms of the structures of [just], [efficient], and [honorable] could not have been

26

satisfactory.

It is a mystery, then, how the structures of thick concepts might explain the application conditions of the thin ones. To clear up this mystery, we would have to somehow break the "symmetry" of explanatory power that exists between thick concepts and their anti- and counter-concepts — to somehow privilege the former over the latter. Later, in our more detailed discussion of thick concepts, we will have a look at some proposals for doing this. However, I shall argue that no such proposal succeeds.

And now let us assume the second, less reductionist picture of thick concepts. It is true that adopting this view gives us fewer resources for a positive explanation of how there could be anti- and counter-concepts that apply just as the "base" concepts do. We cannot say, as we did before, "Just mix-and-match the parts!", because there aren't parts — or at least, there aren't parts (c) and (d) — to mix and match. Instead, we will simply have to put the challenge in this way: why could there not be, say, concepts with the same extension as [chaste], but with opposing sets of valences, and as such, opposing explanatory powers?

This is weaker argumentative posture for us, as we are not giving the opponent a positive reason to think that there are applicable anti- and counter-concepts, and consequently, a prima facie explanatory symmetry. But our opponent, too, is in a weaker argumentative position, for now she cannot appeal to thick concepts' constitutive structures to break the alleged explanatory symmetry. She cannot say, "Here is how we get from the fine, intricate structure of [chaste] to chaste's application

conditions, but lo, there is no similar path from the fine, intricate structure of [anti-chaste] to its application conditions." Without the aid of something like conceptual structure, one is left without the resources to provide an external explanation of why a concept applies while its anti- or counter-concept does not. One would have to resort to explanations from "inside" the practice of thick concept use.

Now, for quite fundamental reasons, those on the less reductionist side of the "thick concepts" debate tends to be more comfortable with purely internal explanations, so I take it many of them would see what I've called a weaker argumentative position as a perfectly strong one. That just means that my agitation for external explanations at the end of this paper will be all the more important.

Constitutivism

The aforementioned are perhaps the most obvious ways of rejecting Premise 3, but the way best represented in the literature is via the theory of practical reasons known as "constitutivism".

The constitutivist believes we can derive practical normativity from what is constitutive not of some concept or other, but of agency. Of course, many things are arguably constitutive of agency. The most effective versions of constitutivism hold that the norm-grounding feature is an aim common to all agency as such. For instance, on David Velleman's (1989 and 2009) view, self-understanding is an aim of all action; on Christine Korsgaard's (2009) view, self-constitution is.

Agency is unavoidable. We are, as Korsgaard puts it, condemned to act. It follows that if there are any agency-constituting aims or ends, they are unavoidable, too.

Additionally, we are committed in two senses to take the means to our ends or to do what constitutes their fulfillment. First, we are rationally committed. There are both narrow-scope and wide-scope readings of the means-ends rule, but for the constitutivists' purposes, it doesn't much matter which is correct. If I cannot but have some end, then I must take the means to that end on pain of being irrational. Second, we are in a certain way psychologically committed. Perhaps we always take the believed means to our (genuine, current) ends unless there is external interference. At the very least, we will always see taking the believed means to our ends as intelligible. I cannot conceive of an agent who persists in having an end while genuinely asking whether/why to take the means to it.

We are committed, then, in all important senses, to do what it takes to achieve our unavoidable ends. A constitutivist hoping to vindicate a normative rule will say that following that rule is what it takes. It is Korsgaard's position that if we persistently fail to guide our behavior by the rules of morality, we will in some sense "come apart" as agents. It is Velleman's that our behavior will fail to make sense to us.22

The capstone of the constitutivist position is a claim about reasons: If one is committed by one's unavoidable ends to following some norm, then one has reason to act in accordance with that norm. So, of course, the concept [what one has reason to do] applies to acts that accord with that norm.

There is no way I can do justice to the constitutivist position here; mostly I'll just

---

22 Korsgaard (1996) and (2009); Velleman (1989) and (2009).

tip my hat to others who've criticized it at greater length. But briefly, I find this reply to Premise 3 unsatisfying because constitutivism falls prey to a "thin/thick" dilemma of its own. It demands that there be something that satisfies two conditions: (1) being constitutive of agency; and (2) committing us to the acceptance of the (usually quite substantive) norms that the constitutivist is trying to vindicate. The problem is that the more plausible it is that something satisfies (1), the less plausible it is that it satisfies (2), and vice versa. Hegel argued that even if we could not help but act "under the ideal of a free will", Kant was wrong to think that this committed us to any substantive morality23. The starting point, he said, was simply too thin. Later constitutivists, anxious to avoid a similar criticism, have seemed to me to err in (at the very least) the other direction. Aiming at self-constitution or self-understanding does not seem to me to be constitutive of action, if by "action" we mean the sort of thing I've been engaged in continuously since I woke this morning. Rather, action with such an aim strikes me as no less optional than action with the aim of helping others or of unseating the Prime Minister — and for that matter no less optional than the employment of some thick concept or other.

Of course there is more to say, and more that has been said, on this score. But

---

23 Hegel (1820).

since my response here is intended to be merely programmatic, let me move on. 24

Connections with Experience and (Non-Ethical) Reality

One might also try to reply to Premise 3 by associating with thin concepts what Sellars called "entry transition" rules. On such a proposal, [good] is like [yellow]: it's the concept it is because of its causal or rational connections with (take your pick): raw experience, non-conceptual content, retinal stimulation, sense data, etc. A thin concept would apply, then, when the experiential state with which it is associated applies — that is to say, when it accurately represents the world, when it is veridical.

Or, we might say that thin concepts are partly constituted by their links with the non-ethical parts of the extra-mental world. A theory of the sort I'm imagining is Jerry Fodor's Asymmetric Dependence Theory, on which the concept C is correctly applied to a worldly item I if and only if (a) I causes tokenings of C, and (b) Things other than I would not cause tokenings of C unless I did, and (c) I would still cause such tokenings

---

24   This criticism differs slightly from David Enoch's influential "schmagency" objection to constitutivism. See his (2006) and (2011a). Enoch seems to think the problem for the constitutivist is that *even if* there is a proprietary norm of agency as such (as there are proprietary norms of basketball), we would not have reason to behave in accordance with such a norm unless we had reason to be agents. But to my ears, there is no difference between "what we have reason to do" and "what would accord with a proprietary norm of agency". I think the problem is just that we can't derive proprietary norms of agency from what is constitutive of agency — not unless we mean "agency" in an overinflated sense. Enoch endorses this objection, too — see his 2011a — but it is not his primary argument.

even if these other things didn't.25 Such a theory purports to explain the conditions under which concepts apply by adverting only to such mind-world causal relations, without relying on what I've been calling a concept's "constitutive structure". (Indeed, Fodor does not think that concepts have constitutive structures so understood.)

If either of these sorts of pictures is right for ethical concepts, then we can give external explanations of how such concepts apply to the world even if these concepts lack definitions, or constitutive prototypes or exemplars.

My reply to both of these proposals is like my reply to the second strategy from the previous section. I would want to count constitutive connections both with experience and with the extra-mental world as part of a concept's structure. This would mean that concepts with these constitutive connections would could as thick concepts, whatever terms we used to express them.

As a purely classificatory maneuver, this reply shouldn't arouse any objections, and if it does, I am of course prepared to re-word the argument to get around them. (For what it's worth, it's not out of keeping with the very broad notion of "conceptual role" favored by some writers on conceptual role semantics.26) What it may arouse is the suspicion that I plan to employ this maneuver in response to any claim that the structure of thin concepts is thicker than I've let on. This suspicion is entirely correct. What this means, though, is that my arguments against the possibility of using thick

---

25 See Fodor (1987) Chapter 4; other sophisticated causal theories include Dretske (1981) and Millikan (1984). The best-known meta-ethical application of this view (although it is put in terms of linguistic rather than mental representation) is of course Boyd (1988).

26 See Greenberg and Harman (2008).

concepts to vindicate ethical objectivity will have to bear a lot of weight.[27] Let us turn to these now.

IV. In Support of Premise 4: If a concept is thick, then there is no external explanation of how something's having the feature the concept designates could ever render doing/having that thing "mandatory".

At the heart of my argument for this premise is the claim from the previous section that for any thick concept, there is an anti- and a counter-concept. I will argue that: (1) The existence of anti- and counter-concepts explains why thick concepts cannot render things "mandatory" in our sense, and that: (2) the fact that thick concepts cannot render things mandatory means that we won't be able to vindicate ethical objectivity by appealing to them. The rest of this section is devoted to (1). (2) will be established in Section V.

Recall that for an action or an attitude to be mandatory is for it to be the case that one who is in a position to do it/have it and does not is thereby ceteris paribus at a cognitive deficit relative to, or else less rational than, one who does do it/have it. That an action is what one all-things-considered ought to do always renders it mandatory. Someone who can do it and does not must either lack the belief that he ought to do it, or else fail to do as he believes he ought. So either he's at a cognitive deficit compared

---

27   As it would if we took *all* ethical concepts to be thick. This is more or less the position taken by

   Timothy Chappell (2013), although he admits that we could artificially define thin concepts (in

   much the same way we artificially define some of the anti- and counter-concepts, I suppose).

to an otherwise similar person who believes this truth, or else he's less practically rational than such a person who does as he believes he ought. We might say that the facts that render things mandatory settle questions about what to do. These facts are the "shot-callers" of the evaluative universe.

"Thick" facts are lower in the chain of command. This is because, for any set of thick concepts the application of which would guide one to do A, there is a set of their anti- and counter-concepts the application of which would guide one to avoid doing A, or to do something in A's complement, respectively. For this reason, the fact that A falls under a set of thick concepts can never render it mandatory: one who fails to do A needn't be ceteris paribus cognitively or rationally worse off than someone who does A. She may simply be applying be applying the anti- or counter-concepts and acting on those.

There are affinities between this argument and Tristram McPherson's (2011) "schmeasons" argument against meta-ethical quietism — with, e.g., anti-justice standing to justice as schmeasons stand to reasons, and mandatoriness playing a role in my argument akin to that played by the idea of "robust normativity" in McPherson's. One point of dissent: if we take [reason] to be a thin concept, we have good reason to think that there cannot be a separate concept [schmeason]. Thin concepts can't have anti-concepts, for their sparse structure positively precludes there being enough elements to "mix and match". To put the point otherwise: There is no equally-thin alternative to the concept [reason]; someone who asks why we should do what we have reason to do instead of what we have schmeason to do is simply asking about what we have reason to do, and is meta-semantically confused (albeit on theoretically

sophisticated grounds, rather than silly ones).

So imagine two people; call them "Val" and "Mal". Val and Mal may live in very different societies or simply have very different political orientations, or each may be the other's counterpart in another possible world. Suppose Val does A on the basis of her (correct) application to A of the concepts C1...Cn. And now suppose Mal omits to do A on the basis of her (correct) application to A of D1...Dn, which are C1...Cn's anti-concepts. Notwithstanding this formal symmetry, mightn't these facts suffice to make it mandatory for, say, Val to do A?

There are two arguments that they may not so suffice. One argument is predicated on the assumption that it is possible for a possessor of a thick concept to also possess its anti-concept; the other argument is not.

Suppose first that Mal, just like Val, correctly applies C1...Cn to A. This is, after all, not ruled out by the aforementioned facts. The difference is that Mal is utterly unmoved by these thoughts. C1...Cn are not among "her concepts". "Her concepts" — the ones that get practical traction on her — are D1...Dn. We can't say so far that Mal is at a cognitive deficit relative to Val, for there is nothing we've said that Val knows that Mal doesn't. And we can't say so far that Mal any less rational. She's acting on the correct application of "her concepts" (but not Val's) just as Val is acting on the correct application of "hers" (but not Mal's).

But it may strike you as implausible that an "outsider" like Mal could really apply C1...Cn. Maybe she could know the conditions under which Val correctly applies C1...Cn, but that's not the same thing. Actually being able to apply the concepts C1...Cn

herself would require being guided by their first-personal application to her present options — "walking the walk", as it were. And we've said that she is not so guided. So let us say that Mal is ignorant of what falls under C1...Cn. This is a cognitive shortcoming of sorts, to be sure. But then we must say, by the same token, that Val is ignorant of what falls under D1...Dn. Ceteris paribus, then, we have nothing that places Mal at a cognitive deficit relative to Mal (or vice versa). There is one part of the world that each carves up using a set of concepts that, at least on this version of our story, precludes application of her counterpart's concepts. Nor is Mal less rational than Val; she is guided to omit A by her application of D1...Dn just as Val is guided to do A by her application of C1...Cn. If it helps, contrast this case with one in which Val does A because she knows she ought to, while Mal fails to do A due to her ignorance that she ought to. Then Mal most certainly is, ceteris paribus, cognitively worse off than Val. Val knows something that Mal does not, and Mal knows nothing (so far) that Val does not.

Now, at this point, one might object that some thin concepts do not pick out "mandating" features, either, and so it can't be something about thick concepts specifically which explains their failure to do so. For example: I may fail to do what I have a reason to do without thereby displaying a cognitive or rational deficit vis a vis those who do the action in question. After all, I have a reason to shoot heroin or to vote for Michele Bachmann. Perhaps what explains the failure of thick facts to render anything mandatory is that these facts have the same broadly pro tanto character as do facts about reasons. And what about the category of the permitted, which delimits

what we may do rather than what we must do or what is better to do? It may seem borderline incoherent to say that something is mandatory because it is permitted.

But this is too quick. For while facts about reasons and permissions do not always render actions mandatory, they sometimes can. For suppose I must choose between two options: O and P. If I have a reason to choose O and no reason to choose P, or simply stronger reason to choose O than to choose P, this would suffice for the choice of O to be mandatory. If I chose P, then I would either have to be ignorant of the relative strength of reasons, or akratic, and as such, practically irrational. And suppose the choice of O was permitted but the choice of P was not. Then, again, the choice of O would be mandatory. Reasons, and other features with a pro tanto character, can settle questions about what to do when there are no reasons or weaker reasons on the other side. And facts about what's permissible can settle such questions, in a sense by default, when only some of the options are permissible. So in general, any thin fact can render an option mandatory when there's no sufficiently "counterbalancing" thin fact on the other side. But with thick facts, there is always something on the other side — anti- and counter-facts — and hence, no circumstances in which they can settle a conflict in the relevant sense.

Extensionless Concepts

It may be objected that there's an obvious difference between concepts like [courageous] and their anti- and counter-concepts: While some actions are courageous, none are anti-courageous; [anti-courageous] is an extensionless concept. And so someone who is acting on beliefs about what's anti-courageous is ceteris paribus

operating at a cognitive deficit compared to someone who is acting on true beliefs about what's courageous. So it would seem that, at least in some circumstances, the fact that an act is courageous renders it mandatory in our sense.

From the point of view of first-order ethical theory, this response seems right. Nothing is anti-courageous. But what's needed to defeat Premise 4 are external explanations of why thick facts can't render things mandatory, not explanations that draw on claims internal to first-order ethical theory. And I claim that such explanations are not to be had. They are not to be had insofar as such explanations of thin concepts' application conditions are not to be had.

For suppose we accept the "thin plus" view of thick concepts, on which [anti-courageous] might mean something like [bad in virtue of facing up to danger]. Then we might argue that this concept is extensionless because something's being an instance of facing up to danger is just not a reason why it would fall in the extension of the concept [bad]. But there could not be a properly external explanation of this unless there were already an external explanation of [bad]'s extension. The previous two sections were devoted to showing that there is none. Generally, the reason there can be no external explanation of why anti- or counter-concepts do not apply is that such an explanation would require the concepts' valences, in conjunction with their other structural features like the "thickening elements", to determine the concepts' extensions. But it is unclear how valences could do this, for as I argued at length earlier, they play no extension-determining role on their own. They are concept-individuating, but not extension-determining.

Nor does it help to accept the less reductive picture of thick concepts. As I argued

in the previous section, this picture fails to provide us with a fully external way of explaining which thick concepts apply and which do not. The non-reductivist's explanations must always take place from within evaluative practice.

Metaphysical Realism

But perhaps there are cognitive shortcomings other than having false beliefs and lacking true ones. A certain sort of metaphysical realist would say that a person whose concepts limn the "fundamental structure" of the world — its "joints", as Quine put it — is capturing the world better than someone whose concepts fail to do so, even if the second person applies his concepts correctly.28 Someone sympathetic to this view may want to say that some thick concepts limn the structure of the world better than others do, and so users of the latter are ceteris paribus at a cognitive deficit relative to competent users of the former.

I have two objections to this sort of argument. First, I join others in finding it difficult to make sense of the idea of structure. Of course, some concepts will be more useful than other concepts in navigating the world, in offering explanations, in tracing out what strike us as patterns. But the metaphysical realist wants to go further than this. She will claim, instead, that the idea of limning worldly structure is not identical to any of these or their combination. Perhaps she will say that concepts with these

---

28 For example, Sider (2011) introduces the idea of structure by asking us to imagine a group of people who use strange concepts and characterizing them as follows: "[T]hey make no mistakes when they apply their own concepts....The problem is that they've got the wrong concepts; they're carving the world up incorrectly." p. 3

other features have them at least partly in virtue of the fact that they limn structure.29 My reason for being confused is roughly this: Metaphysical realists claim that failure to limn structure is a failure to get the world right. To put things in "direction of fit" language, such failure is a failure of thoughts to fit with the world. As I understand the notion of representation, any failure of this sort is representational failure. Now, realists like Sider claim that I may apply a concept correctly and yet fail to limn structure. But I cannot see how, as far as that concept goes, I exhibit any failure of representation when I apply it correctly. So while the notion of liming the world's structure allegedly has features that would make it a representational notion, I cannot see how it could be representational if it goes beyond simply applying concepts correctly.

More conciliatorily: Insofar as I can get a rough idea of "structure" through examples, it doesn't seem to me that structure would be a useful tool in privileging some thick concepts over others. Many of the realists' examples involve contrasts between concepts in different domains. Fundamental physical concepts are alleged to carve nature at the joints; racial concepts clearly do not. But of course, thick concepts like [courageous] and [anti-courageous] are drawn from the same domain. And many realist examples invoke contrasts between concepts that seem to pick out "simple" properties and those that seem to pick out "complex" properties; [tall] carves closer to the joints than does, say, [taller than at least one of the seventeen most famous Bollywood actors who nearly moved to Anaheim], for instance. Now, it is true that some thick concepts seem to pick out simpler properties than others. Going with a

---

29 See Putnam (1975).

reductionist view for a moment, it seems natural to say that [good in virtue of facing up to danger] is simpler than [good in virtue of contributing to the reduction of melancholy brought on by co-dependent relationships with parakeets]. However, the former does not seem simpler than [bad in virtue of facing up to danger]. The class of privileged concepts cannot, then, be winnowed down very much on the basis of simplicity.

Pragmatism

Like the realist, the pragmatist wishes to privilege some concepts over others. But she does not claim that the former enable us to better represent the world than the latter do; rather, their advantages are distinctly practical. So a pragmatist reply to Premise 4 is best read not as a claim that [anti-courageous] users exhibit cognitive shortcomings, but as a suggestion that an action may also count as mandatory when failure to do it is due to one's using a concept that is worse in some non-representational way.

I find the pragmatist position plausible just for the banal reason that using concepts is something we do, and it is trivial that we should do what we have stronger practical reasons to do. Indeed, I would have no problem suggesting that we opt for some non-ethical concepts rather than others on pragmatic grounds.

Things are dicier, though, when it comes to ethical concepts. Claims about what we have practical reason to do are claims within the ethical domain, and so privileging some thick concepts over others on pragmatic grounds is a kind of within-the-domain maneuver. In valorizing the use of concepts on pragmatic grounds, we are taking sides

41

within the very domain to which those concepts belong. For example: Should we structure our normative thinking around concepts like [distributive justice]? A Nietzsche would say "no", on the the grounds that doing so would not create a fertile climate for the ascension of "free spirits" and "overmen". But it is difficult to imagine many friends of distributive justice giving up their favored concept on these grounds.

I should say: it is not obvious what practical effects using various concepts would have. It may be that the use of certain concepts would be bad by almost everyone's lights. It may even be that the use of certain concepts would be bad by those whose evaluative outlooks are captured by those concepts! All the same, it does seem that pragmatic vindications of ethical concepts are internal vindications, and, therefore, vindications which will be rejected on the basis of judgments employing the very concepts whose use is up for debate.

V. In Support of Premise 5: For there to be an explanation of the objectivity of ethics that is sufficient for certain fundamentally important purposes, there would have to be at least one concept such that: (a) there is an external explanation of its conditions of correct application, and (b) there is an external explanation of how something's having the feature it designates could ever render doing/having that thing mandatory.

To defend this premise, we'll need to answer two pressing questions:

(1) Why should we think that ethical objectivity requires at least one concept that picks out a   feature that renders actions or other things mandatory?

(2) Which fundamentally important purposes are served by external explanations but not   internal ones? For what, in other words, do we need external explanations?

Why Require Mandatoriness?

One might think that ethics is objective just so long as there are objective ethical truths, which there are just in case there are ethical truths that are, in the appropriate sense, stance-independent. But I deny this. There is a gap between a domain's containing objective ethical truths, and the domain itself's being objective. For example, it is an objective truth that my act was courageous just in case it is true independently of anyone's evaluative stance that it was courageous. However, this truth would not suffice for ethical objectivity unless the act's status as courageous could render the act mandatory. I have argued that it could not.

So why set the "bar" for objectivity where I do? Why insist on mandatoriness?

The starting point for addressing these questions should be our interests in objectivity. We might have, prior to any attempt to explicitly lay out the criteria for ethical objectivity, a practical interest in whether ethics is an objective domain. That is, we might want to act in one way if it is, and a different way if it is not. We should set out criteria for ethical objectivity such that it ends up having roughly these practical implications, and look less kindly upon characterizations on which it makes no difference whether ethics is an objective domain. Additionally, we might have, again prior to any theorizing, an attitudinal stake in believing that ethics is objective or not objective. Put simply, we may want ethics to be, or to not be, an objective domain. We

should prefer characterizations of ethical objectivity such that, if it obtains, we get what we wanted. My methodology says that ethical objectivity is just "what we care about when we care about ethical objectivity".30

The methodology suggested here cannot be fully pursued in such a short space. There are lots of reasons why we might care about ethical objectivity. But let me give just give one reason we might care whether ethics is objective, and one reason we might want ethics to be objective. I'll argue that the conception of objectivity I offer in this paper is a better candidate than the "stance-independence" conception to serve as the object of these interests.

Why We Care Whether Ethics is Objective

David Enoch (2011b) expresses support for the methodology above, and suggests that ethical objectivity matters in the following way: In a "significant, roughly recognizable class of cases" in which I must make an ethical decision together with a

---

[30] A referee suggests that there is a tension between my insistence on explanations that do not themselves depend on evaluation, and my urging a criterion of objectivity on grounds of what we value. Now, I think that insofar as it's a problem for me, it is — or rather, *should* be — a problem for everyone who insists on non-evaluative explanations in ethics, because I see no satisfactory grounds for adopting any criterion of objectivity that are not, ultimately, evaluative grounds. I do, however, think that the referee's observation is astute. Indeed, the vindicatory project for which I intend this paper to serve as an argumentative foil is centered on the thought that we cannot truly step outside of the evaluative standpoint, even while engaged in areas of philosophy that do not present themselves as purely evaluative.

conflicting party, I ought to:

      (a) Opt for a conciliatory solution (e.g. flipping a coin, taking turns, letting an
impartial third   party decide) if ethics IS NOT objective; but

      (b) "Stand my ground" if ethics IS objective.[31]

Now, Enoch takes it as a further premise that I ought indeed to stand my ground in cases of this recognizable class. If, for example, we find a stray dog and must decide what to do with it, I should respond to your preference for torturing it with the resolute insistence that we look for its owner instead. Coupled with the conditional claim above, this leads him to the conclusion that ethics is objective.

My concern here is not with this further premise, but with the conditional claims. We can see their plausibility by focusing, as Enoch does, on two sorts of non-ethical conflicts: those of mere preference and those based on beliefs about obviously factual domains, respectively. A conflict of mere preference: Suppose you and I are passengers in our mutual friends car. I would most prefer to listen to Stereolab; you would most prefer to listen to Grizzly Bear. It would seem completely inappropriate for me to insist on Stereolab over Grizzly Bear. This seems to follow from the natural thought that there is no objective fact about which of these groups is better. And now, a

---

[31] (2011b), Ch. 2. Those inclined to dismiss the view are urged to read Enoch first and to note in advance: (a) the importance of his restriction of the relevant "class" to cases of parties making a decision *together*, and (b) unobviousness in many cases of what constitutes the best interpretation of a "conciliatory" solution.

conflict based on beliefs about a factual domain: Suppose our driver is trying to navigate his way out of Port Moresby. If I am a Port Moresby native and you are an outsider, it seems right to insist that the driver follow my directions rather than yours. This seems to follow from the fact that domain of Port Moresbian geography is objective, that the world settles our dispute in favor of me over you.

But why are these conditional claims plausible? Why should ethical objectivity have such practical significance? The answer seems to me to be this: In a conflict of our wills, I have a stronger claim to mine prevailing if mine is a purer expression of my self than yours is of your self. This happens when yours is more infected by certain sorts of mistake than mine is. There seem to me to be two types of mistake that can so contaminate a will. First, the direction of the will may be based on a failure to represent the world accurately. This is what is going on, for instance, in Mill's well-known "bridge" example of justified paternalism.[32] Second, its direction may be based on a failure to act in a way that is rendered rational by one's representation of the world. The legitimacy of standing one's ground depends on objectivity because objectivity means that at least some differences evince mistakes. In a non-objective domain, it may be that every difference is... well, just a difference.

So if we want a conception of ethical objectivity on which it has Enoch's upshot, then we should favor "objectivity as requiring mandatoriness" over "objectivity as requiring mere stance-independence". For suppose ethics is objective in the first sense. Then there are some Enochian conflicts where one solution is mandatory. So if one party embraces that solution, then the other party must be doing worse either

----

[32] Mill (1859).

cognitively or rationally. In other words, this other party must be, to a greater degree, mistaken in one of the aforementioned senses. And we've just seen that such mistakes have practical consequences: One is more justified in overcoming or circumventing the wills they influence. But now suppose objectivity is just stance-independence. This is consistent with there being no mandatory actions, and thus with no party being any more convictable of mistake than any other, and thus with no will being any more liable to silencing than any other, and thus with no ethical conflicts in which "standing your ground" is right. Basically, if ethical objectivity doesn't require mandatoriness, then it's just not at all clear why it should have the practical implications Enoch attributes to it. Prima facie, this speaks in favor of the "mandatoriness" conception.

Why We Want Ethics to Be Objective

Richard Rorty's lecture "Human Rights, Rationality, and Sentimentality" aims to convince us that the issue of ethical objectivity is not practically significant. His efforts of course prompt the psychological question of why, in spite of this, so many of us take the issue so very seriously. Rorty's Nietzschean answer to this question is evocative enough to be worth excerpting in substantial part:

> "To rely on the suggestions of sentiment rather than on the commands
> of reason is to think of powerful people gradually ceasing to oppress
> others, or ceasing to countenance the oppression of others, out of mere
> niceness, rather than out of obedience to the moral law. But it is
> revolting to think that our only hope for a decent society consists in

softening the self-satisfied hearts of a leisure class. We want moral

progress to burst up from below, rather than waiting patiently upon

condescension from the top. The residual popularity of Kantian ideas of

"unconditional moral obligation" – obligation imposed by deep

ahistorical noncontingent forces – seems to me almost entirely due to

our abhorrence for the idea that the people on top hold the future in

their hands, that everything depends on them, that there is nothing

more powerful to which we can appeal against them.

[W]hy does this preference make us resist the thought that

sentimentality may be the best weapon we have? I think Nietzsche gave

the right answer to this question: We resist out of resentment. We resent

the idea that we shall have to wait for the strong to turn their piggy

little eyes to the suffering of the weak. We desperately hope that there is

something stronger and more powerful that will hurt the strong if they

do not – if not a vengeful God, then a vengeful aroused proletariat, or, at

least, a vengeful superego, or, at the very least, the offended majesty of

Kant's tribunal of pure practical reason. The desperate hope for a

noncontingent and powerful ally is, according to Nietzsche, the common

core of Platonism, of religious insistence on divine omnipotence, and of

Kantian moral philosophy."

Whether or not this describes "we", it most certainly describes me. Of course I

grant that the strong will not lose anything material by virtue of failing before the "tribunal of pure practical reason". But I find that it matters to me that they will lose something in the way of status — that they will eo ipso count as irrational or morally blind — that they will be robbed of a certain kind of, well, legitimacy, that would otherwise cause me to take their cruelty and indifference to heart.

We should also say: there's no guarantee that an objective tribunal will find for you rather than the other guy. And while I suppose you might welcome the prospect of ethical reality "calling you on your shit", as it were, the desire Rorty's articulating seems to be one that sustains itself against the background of a conviction that you are on the side of the angels — that the positively-valenced concepts you're applying to your own behavior really do so apply; mutatis mutandis for the negatively-valenced ones and the oppressors' behavior. The desire for objectivity, then, is rooted in a desire that the oppressors be mistaken.

All else being equal, we should prefer a characterization of ethical objectivity that makes it the object of this yearning. And here, again, it seems that the "mandatoriness" conception fares better than the "mere stance-indendepdence" conception. The latter conception, again, would allow for objectivity even if nothing turns out to be mandatory in my sense. But if nothing is mandatory, then there is nothing that one is (comparatively) cognitively deficient or irrational simply for not doing. Objectivity on this latter conception needn't label anyone any more morally blind or irrational than anyone else. By contrast, if ethics is objective on the "mandatoriness" conception, then some possible actions are indeed mandatory. Those who fail to do them are, ceteris paribus, worse at representing the world, or else at

transitioning from representation to action, than those who do them. Objectivity of this sort at least has the chance of giving the "resenters" what they want — a diagnosis of the powerful as mistaken if they abide or perpetrate oppression. So far, then, it seems that mandatoriness is what we care about when we care about ethical objectivity.

This puts us in a position to answer a question I broached earlier: Do the arguments of this paper, if sound, imperil objectivity in significantly many non-ethical domains? It seems to me that they do not. What's special about ethics is that the possibility for conflict outruns the possibility of theoretical disagreement. Specifically, my use of a thick concept and your use of its anti- or counter-concept can embroil us in a practical clash even where there is no theoretical one. I've argued in this section that we care about objectivity in large part because we care about there being a right side and a wrong side in practical conflicts. We care about this for the two reasons I've outlined. So objectivity in ethics would seem to require some way of privileging certain thick concepts over their anti- and counter-concepts, but I've claimed there is no such way that is fully domain-external.

However, in most other domains, there can be no conflict without genuine disagreement. Our attitude towards those with partially different conceptual schemes tends to be "Live and let live!" rather than "There can be only one!". This is why most of us are perfectly comfortable with productive stipulation in the non-ethical context but would be taken aback were an interlocutor to try to stipulate, say, the truth of utilitarianism. There is no need to privilege some non-ethical concepts over any

others, since we cannot be brought into conflict merely by dint of our using different non-ethical concepts.

There are difficult liminal cases, though, because we as yet lack a satisfactory theory of what a conflict is, as distinguished conceptually from a disagreement. If I regard a piece of music as brutal and am moved to a state of cathartic nihilism by it, while my wife regards it as bludgeoning and is moved to state of depressive nihilism by it, are we in a conflict? As you might guess, my approach to developing such a theory would be to try to identify a conflict with what we care about insofar as we pre-theoretically care about conflicts.

Why External Explanations?: Dialectical Effectiveness, Avoiding Representational Miracles, and Unifying Domains

Some readers may want to question the insistence on "external" explanations. In doing so, they would be joining forces with so-called "philosophical quietists", who typically regard such explanations both as impossible and as unnecessary for any fundamentally important purposes. I take the quietist challenge seriously — too seriously to suppose that I can defang it in such a short space. But let me at least make explicit what I think external explanations can do that internal ones cannot, and in the process defuse some of the positive reasons for rejecting external explanations.

First, external explanations can do dialectical work that internal ones cannot. They are the only explanations capable of convincing the would-be ethical skeptic, whether metaphysical or epistemological, of one's ethical claims. And they may be capable of convincing certain non-skeptical opponents who would not be swayed by

ordinary, first-order argument. These include radical foes whose views we might classify as unreasonable, but also many people who we'd generally regard as our allies. The deep disputes between, say, Frances Kamm and Peter Singer may well be intractable though ethical argument alone, despite the fact that each can see how a nice, well-brought-up person could support the other side, and despite the fact that they'd generally agree at the level of practical solutions to real-world problems. Finally, external explanations may be capable of emboldening those who agree with us, by showing them that — and how — those on the other side are mistaken. Preaching to the converted is valuable when their faith in the cause wavers.

Quietists may argue that this dialectical is not a significant advantage. They'll note at the outset that we needn't be capable of convincing others in order to be justified ourselves. Additionally, they'll say, it isn't clear that there is any other value in being able to convince opponents or wavering supporters of the sort I've described. We may deeply yearn to do so — a yearning admirably laid bare in the Rorty passage above and indeed in much of Rorty's work33 — but the most thoroughgoing quietists will claim that we oughtn't to try to satisfy such yearnings; rather, we should regard them as post-Enlightenment anxieties to be dissolved by philosophy-as-therapy (and maybe just therapy).

---

33 E.g. Rorty (1992): "I wanted very much to be some kind of Platonist, and from 15 to 20 I did my best…The first goal [of the Platonist] is to achieve argumentative power over others — e.g., to become able to convince bullies that they should not beat one up, or to convince rich capitalists that they must cede their power to a cooperative, egalitarian commonwealth. The second goal is to enter a state in which all your own doubts are stilled, but in which you no longer wish to argue."

But it is not clear to me, in the first place, that my own justification can float free of my ability to convince others, for justification may well depend on context, and persistent and salient disagreement and doubt may well give rise to contexts in which only external explanations can provide justification.34 I also think we should care about dialectical power for reasons quite apart from the justificatory status of our own beliefs. We are inevitably faced with practical conflicts. Either the United States will continue to allow private citizens to own semi-automatic weapons or it will not; either Britain will stay in the European Union or it will not. The members of a diverse society cannot all get what they want. Now, we can respond to these conflicts either by withdrawing to the sidelines and effectively ceding victory to our opponents, or by trying to gain ground from those opponents. It strikes me as less-than-honourable for any individual on the correct side of a conflict to respond in the first way, and positively disastrous for most or all of those on the right side to respond in this way. That leaves the second, more steadfast response. There are many ways of gaining ground from an opponent — force, fraud, bargaining, shaming, playing on sentiment, and rational argument that carries the opponent from premises she already accepts to the conclusion you endorse. Of course there is much to say about the ethics of political engagement, but it seems to me that only the last of these treats the opponent in a way that is both non-coercive and respectful of her rational faculties. This is all to say: we should care about dialectical power because should want to enter the moral and political fray — but not only that, also to win — but not only that, also to win cleanly.

---

34 On contextualism in epistemology, see Lewis (1996) and, of course, many others.

We should also want external explanations because they are required to eliminate the appearance of what we might call representational miracles — episodes of putatively accurate representation that are not explained by substantial accounts of representation. One way for there to be representational miracles is for there to be accurate representations that are not amenable to any explanation at all — i.e. brutely accurate representation. Another, perhaps less obvious way is for there to be accurate representations that are amenable only to explanations from outside the theory of representation. All else being equal, we should want to avoid positing representational miracles just as we want to avoid positing medical miracles — e.g. a patient's kidney's failing for no medically articulable reason.

Both internal and external explanations of a representational episode can rescue it from being brute. But only an external explanation can rescue it from being miraculous in the second way. This is because internal explanations of accurate representation are not explanations from within the domain of representation itself, but rather from within the domain(s) of the concept(s) involved in the representing, augmented by something like conceptual ascent. They are not the sorts of explanations proffered by theorists of mental representation like Fodor or Millikan, or for that matter, theorists of pictorial or graphical or poetical representation. Equipped only with explanations of accurate representation internal to the domains from which the representing concepts are drawn, our account of representation will be nothing but a patchwork of miracles: "Even numbers are divisible by 2, so the thought that even numbers are divisible by 2 accurately represents; subversion of others' autonomy is prima facie wrong, so the thought that subversion of others' autonomy is prima facie

wrong accurately represents; when you play with fire you get burnt, so the thought that when you play with fire you get burnt accurately represents...". Such a patchwork fails to unify representational facts, and reveals nothing about the notion or nature of representation. We would at least hunger for more than this in any other theoretical domain; it is unclear why we should be more easily contented when it comes to the representational one.

Finally, an external vindication of ethical objectivity would serve to unify ethics with other objective domains like physics, biology, and history, which seem to admit of such vindications. This is a comparative point — not simply that ethics is suspect unless we can vindicate it from "sideways-on", but that ethics looks all the more suspect if we can vindicate other domains from sideways-on but not ethics. As I've noted, one way to vindicate a domain from the outside is to offer a robust account of representation, and then use it to show that certain thoughts about that domain can accurately represent the world.

Now, some quietists may simply deny that there are any robust accounts of representation capable of vindicating these other fields. I have two quick arguments against such blanket skepticism. The first is to note an interesting sort of tension in the ethical quietists' thought: Most of these quietists do not think that there can be no robust explantions within first-order ethics. Even if they reject comprehensive, exceptionless ethical theories, they typically do not deny that we can explain, e.g., why theft is wrong, or why taxation is not tantamount to theft. They don't think that the tendency towards rigorous ethical thought betrays an "anxiety" that is better "treated" than satisfied. Indeed, they want to steer people away from any sort of constructive

meta-ethics and towards the careful, confident sort of normative ethics one sees in, e.g., Aristotle. But the philosophy of mental representation is an evaluative domain just as ethics is. One concerns questions of right action; the other concerns questions of right thought. So if we can avoid positing bare miracles in ethics, why not in the theory of representation?35 The second is an argument-from-authority. Such a blanket skepticism must deny not only that generic accounts of representation like Fodor's may be successful, but also that, e.g., accounts of specifically scientific representation — penned by people with a concern for fidelity to scientific practice — may be successful. Here I am thinking of work like: Hesse (1962), Giere (1988), Morgan and Morrison (1999), Suarez (2003), Contessa (2007), and van Fraassen (2008). It seems to me that the quietist is fighting an uphill battle here against philosophers who are closer to being insiders to the relevant first-order domains than the quietists themselves are.

It is common for quietists to reject the insistence on external explanations because they think the search for such explanations is futile, and such, will lead us to skepticism. Since skepticism is surely mistaken, our only option is to learn to content ourselves with internal explanations. Now, of course, if the skepticism at issue is

---

35 The alert quietist will try to turn it around on me: "Well, if ethics and the theory of representation are so similar, and we need domain-external vindications of ethics, then don't we need domain-external vindications of the theory of representation -- even for the negative claims you make in this paper? And where are you gonna get *those*?" A good question; something like it should trouble any (as McDowell would say) "constructive" philosopher. I offer a response in my [OMITTED]. See also Dewey (1911).

specifically ethical, and the external explanations in question are specifically those external to ethics, then this simply begs the question against an argument like mine. But I do not think that is what quietists have in mind. I think they worry that a demand for external explanations is, rather, a recipe for a more global skepticism — that just as we cannot explain from outside ethics how ethical concepts apply, we cannot explain from outside of any domain how its concepts apply. This is just to say: they reject the presuppositions behind the third advantage of claimed for external explanations. This, then, seems to me to be a fruitful locus for future meta-ethical debate, but I won't pursue that any further here.

Conclusion

It strikes me that there are, broadly speaking, two very different lines of response open to the ethical objectivist. First, he try to show, on the basis of a domain-external theory of mental representation, how one can accurately represent the world in ethical thought. But there' s a hitch: if he does this by imputing content or structure to ethical concepts beyond the bare minimum, he opens himself up to the challenge: "Why couldn't there be alternative ethical concepts, with different content or structure? And if there could be, then why think your concepts are in good order while these alternative concepts are not?" I've expressed doubt that it is possible to meet this challenge through an argument that is domain-external in the way that a general theory of mental representation is supposed to be.

Second, the objectivist may explain why ethics does not stand in need of vindication by a domain-external theory of mental representation. This strikes me as

the more promising tack. But it is as yet the minority position among those working in meta-ethics. Many so-called quietists think of the archetypal skeptical challenge as what I'd earlier called an "object naturalist" one, which alleges that it is mysterious how there can be such things as ethical properties in the world as described by science. My hope is that quietists will start to turn their attention to "subject naturalist" challenges like the one I've outlined here, which may succeed even if we entirely put aside worries about how such things as ethical properties might find a place amidst the furniture of the world.

REFERENCES

Block, Ned. "Advertisement for a Semantics for Psychology." Midwest Studies in Philosophy 10.1 (1987): 615-678.

Boyd, Richard N. "How to be a Moral Realist." Essays on Moral Realism, Geoff Sayre-McCord (ed.), Ithaca, NY: Cornell University Press, pp. 181-228.

Broome, John. "Normative Requirements." Ratio 12.4 (1999): 398-419.

Chappell, Timothy. "There Are No Thin Concepts," Thick Concepts, Simon Kirchin (ed.), Oxford University Press, 2013.

Contessa, Gabriele. "Scientific Representation, Interpretation and Surrogative Reasoning." Philosophy of Science 74 (2007): 48–68.

Cuneo, Terence and Russ Shafer-Landau. "Moral Fixed Points: New Directions for Moral Nonnaturalism." (ms).

Dancy, Jonathan. "In Defense of Thick Concepts," Midwest Studies in Philosophy (1995).

Dewey, John. "The Problem of Truth." (1911) The Essential Dewey, Volume II: Ethics, Logic, and Psychology. Larry A. Hickman and Thomas M. Alexander (eds). Bloomington: Indiana University Press, 1998.

Dretske, Fred. Knowledge & the Flow of Iinformation. The MIT Press, 1981.

Enoch, David. "Agency, Shmagency: Why Normativity Won't Come from What Is Constitutive of Action." The Philosophical Review 115.2 (2006): 169-198.

Enoch, David. "Shmagency Revisited." New Waves in Metaethics, Michael Brady (ed.) Palgrave Macmillan, 2011.

Enoch, David. Taking Morality Seriously: a Defense of Robust Realism. Oxford University Press, 2011.

Field, Hartry. "Logic, Meaning, and Conceptual Role." The Journal of Philosophy (1977): 379-409.

Fodor, Jerry A. Psychosemantics: The Problem of Meaning in the Philosophy of Mind. The MIT Press, 1987.

Foot, Philippa. "Moral beliefs." Proceedings of the Aristotelian Society. Vol. 59. Blackwell Publishing, 1958.

Giere, Ronald. Explaining Science: A Cognitive Approach. University of Chicago Press, 1988.

Greenberg, Mark, and Gilbert Harman. "Conceptual Role Semantics." Oxford Handbook of the Philosophy of Language. Ernie Lepore and Barry Smith (eds.), 2008.

Hare, R. M.. The Language of Morals. Oxford University Press, 1972.

Harman, Gilbert. "(Non-Solipsistic) Conceptual Role Semantics." New Directions in Semantics. Ernie Lepore (ed.) London: Academic Press, 1987.

Harris, Sam. The Moral Landscape: How Science Can Determine Human Values. Simon and Schuster, 2011.

Hegel, G.W.F. Elements of the Philosophy of Right (1820). Allen Wood (ed.), H.B. Nisbet (trans.) Cambridge University Press, 1991.

Hesse, Mary. Models and Analogies in Science. South Bend, IN: Notre Dame University Press, 1962.

Horgan, Terence, and Mark Timmons. "Troubles on Moral Twin Earth: moral queerness revived." Synthese 92.2 (1992): 221-260.

Horwich, Paul. Meaning. Oxford University Press, 1998.

Kirchin, Simon. "The Shapelessness Hypothesis," Philosophers' Imprint (2010).

Kirchin, Simon. "Thick Concepts and Thick Descriptions," in Kirchin (ed.), 2013.

Kolodny, Niko. "Why be Rational?." Mind 114.455 (2005): 509-563.

Korsgaard, Christine M.The Sources of Normativity. Cambridge University Press, 1996.

Korsgaard, Christine M. Self-Constitution: Agency, Identity, and Integrity. Oxford University Press, 2009.

Lakoff, George. Women, Fire, and Dangerous Things: What Categories Reveal About The Mind. University of Chicago Press, 1987.

Lewis, David. "Elusive Knowledge." Australasian Journal of Philosophy 74.4 (1996): 549-567.

Machery, Edouard. Doing Without Concepts. Oxford: Oxford University Press, 2009.

Margolis, Eric, and Stephen Laurence (eds.) Concepts: Core Readings. The MIT Press, 1999.

Mackie, J.L. Ethics: Inventing Right and Wrong. London: Pelican Books, 1977.

McDowell, John. "Virtue and Reason." The Monist 62.3 (1979): 331-350.

McDowell, John. "In Defence of Modesty", Michael Dummett: Contributions to Philosophy, Barry Taylor (ed.) Martinus Nijhoff, Dordrecht, 1987.

McDowell, John. Mind and World. Harvard University Press, 1994.

McPherson, Tristram. "Against Quietist Normative Realism." Philosophical Studies 154.2 (2011): 223-240.

Medin, Douglas L., et al. "Correlated Symptoms and Simulated Medical Classification." Journal of Experimental Psychology: Learning, Memory, and Cognition 8.1 (1982): 37-50.

Mill, John Stuart, On Liberty. (1859) Hackett Publishing, 1978.

Millikan, Ruth G. Language, Thought, and Other Biological Categories: New Foundation for Realism. The MIT Press, 1984.

Morgan, Mary and Margaret Morrison (eds.) Models as Mediators. Cambridge University Press, 1999.

Peacocke, Christopher. A Study of Concepts. The MIT Press, 1992.

Peirce, Charles Sanders. "How to Make Our Ideas Clear." Popular Science Monthly 12.286-302 (1878).

Price, Huw. Naturalism Without Mirrors. Oxford University Press, 2011.

Putnam, Hilary. Mathematics, Matter, and Method. Vol. 1. Cambridge University Press, 1979.

Roberts, Debbie. "It's Evaluation, Only Thicker," in Kirchin (ed.), 2013.

Rorty, Richard. "Trotsky and the Wild Orchids." (1992) Philosophy and Social Hope. Penguin, 1999.

Rorty, Richard. "Human Rights, Rationality and Sentimentality." (1993) Truth and Progress: Philosophical Papers. Cambridge University Press, 1998.

Rosch, Eleanor. "Principles of Categorization." Concepts and Categorization. Rosch and Barbara Lloyd (eds.) Hillsdale, NJ: Lawrence Erlbaum, 1978.

Scanlon, T. M. "Structural Irrationality." Common Minds: Essays in Honor of Philip Pettit. Geoffrey Brennan, Robert Goodin, and Michael Smith (eds.) Oxford University Press, 2007.

Sellars, Wilfrid. "Some Reflections on Language Games." Philosophy of Science 21.3 (1954): 204-228.

Sellars, Wilfrid. Empiricism and the Philosophy of Mind (1956), Harvard University Press, 1996.

Sider, Theodore. Writing the Book of the World. Oxford University Press, 2011.

Smith, Edward E., and Douglas L. Medin. Categories and Concepts. Harvard University Press, 1981.

Stenning, Keith. Seeing Reason: Image and Language in Learning to Think. Oxford University Press, 2002.

Suárez, Mauricio. "Scientific Representation: Against Similarity and Isomorphism." International Studies in the Philosophy of Science 17 (2003): 225–44.

Temkin, Larry. Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning. Oxford University Press, 2012.

Väyrynen, Pekka. "Shapelessness in Context." Noûs (2012), online at: http://

onlinelibrary.wiley.com/doi/10.1111/j.1468-0068.2012.00877.x/abstract

Van Fraassen, Bas. Scientific Representation. Oxford University Press, 2008.

Velleman, J. David. Practical Reflection. Princeton University Press, 1989.

Velleman, J. David. How We Get Along. Cambridge University Press, 2009.

Wedgwood, Ralph. The Nature of Normativity. Oxford University Press, 2007.

Whyte, J.T. "Success Semantics." Analysis 50 (1990): 149–57.

Williams, Bernard. Ethics and the Limits of Philosophy. London: Taylor & Francis, 1985.

Wright, Crispin. Truth and Objectivity. Harvard University Press, 1993.